Near-Optimal Multi-Robot Motion Planning with Finite Sampling

Thesis submitted in partial fulfillment of the requirements for the M.Sc. degree in the

Blavatnik School of Computer Science, Tel-Aviv University

by

Dror Dayan

This work has been carried out under the supervision of
Prof. Dan Halperin

October 2021

Acknowledgments

I thank Nir Goren for providing a Python motion-planning framework implementation, and Jiaoyang Li for providing an implementation for `MC-CBS`. I also thank the entirety of our robotics lab for fruitful discussions.

I deeply thank Kiril Solovey for his tremendous help with this research. From instrumental ideas and concepts, through helping with cleaner proofs and text to guidelines for presenting scientific papers.

Finally, I thank very deeply Dan Halperin, for the huge help with the whole process. From the very first steps, though every idea and every part of the text. For helping me understand when to let go and when to try harder, for teaching me how to write my ideas and proofs legibly, and for insisting on the mathematical rigor of this thesis. Without your assistance this thesis would not have been as good as it is.

Abstract

An underlying structure in several sampling-based methods for continuous multi-robot motion planning (MRMP) is the *tensor roadmap*, which emerges from combining multiple probabilistic roadmap (PRM) graphs constructed for the individual robots via a tensor product. We study the conditions under which the tensor roadmap 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 conflict-based search, 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 tensor roadmap. 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.

Contents

1 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 [20, 22, 48, 53]

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. This thesis focuses on centralized approaches for multi-robot motion coordination where the planning is carried out by a single entity that knows the state of the entire system. In the decentralized setting, which is outside the scope of this work, each robot has autonomy in planning its motion typically based on local knowledge about the system (say the state of its few neighboring robots) [4, 33, 37, 39, 60]. A common approach to centralized 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 [15], path-based search methods [59], and conflict-based search [42]. Those have been successfully applied to a variety of problems involving multiple robots—from warehouse management [34] to multi-drone package delivery [7]. 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 [19, 32, 45]. 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 [1, 46, 52, 56]. 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 the 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` [24, 25, 31], and `RRT` [28, 30] 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 [17, 21, 23, 27, 49].

Consequently, a variety of SB-MRMP approaches have emerged, from methods that aim to apply `PRM`-based solutions [41, 58], to techniques that sample local instances of the MAPF problem [29, 47]. A different method, termed discrete RRT (`dRRT`) [51], 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 [44], 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.

Another topic we discuss in this thesis is that of finding an economical cover of a cube of arbitrary dimension using (small) spheres. This problem is of interest to us as we aim to sample points that represent the space with an accuracy of our choosing. In order to do so, we use an ε-cover, which is the concept of covering a set with balls of radius ε [5, 35, 36].

This concept is closely related to another problem, which is covering the whole domain with spheres; in this context, the problem is referred to as *cover density* Cover density is the ratio between the volume of the spheres and the volume of the covered object. In the unbounded case (e.g., the entire plane) it is the limit of covering a ball with a radius tending towards infinity [5, 8]. Finding a set that faithfully represents the underlying hypothesis space is also related to the concept of ε-nets, namely approximating a general set by a collection of simpler subsets [3, 38]. The latter is frequently used in AI to bound the size of the hypothesis space [6, 12, 18, 40, 57].

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` [32] (a continuous-space extension of `CBS`), and `M*` [59] (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`[1] 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 [44] (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 [55], to achieve near-optimality for the single-robot case. We mention that we prove this result for holonomic robotic systems, where the configuration space of each robot is $[0,1]^d$, for some $d \geqslant 2$.

The organization of this thesis is as follows. In Section 2 we provide basic definitions, which we will use later throughout the thesis. In Section 3 we describe the single-robot problem, introduce the staggered grid, and study the theoretical properties of `PRM` using this sampling scheme. In Section 4 we describe our central contribution, namely the extension of this theoretical result to the multi-robot setting. We provide experimental results in Section 5. We then expand in Section 6 on the staggered grid as a covering for a cube, and conclude with an outline of future work in Section 7. In the appendix we consider the special case of systems with planar configuration spaces, i.e., $\mathbb{R}^2$, and describe a sampling scheme that requires fewer samples than our staggered grid.

The main results of this thesis were published in IEEE ICRA 2021 [10], and can also be found publicly online [9].

[1] Throughout the thesis we slightly abuse the term `PRM` to refer generally to maps of individual robots, even when they are deterministic.

2

Preliminaries

In this section we will formally define all the basic terminology, which will be used throughout this thesis. These definitions are all elementary and are here mainly for completeness: metric spaces, the space of free configurations, and trajectories of robots. First let us define a metric function.

Definition 2.1. *Given a set X, a function $d : X \times X \to [0, \infty)$ is called a metric function if the following conditions hold:*

(i) For all $x, y \in X$, $d(x, y) = 0 \Leftrightarrow x = y$.

(ii) For all $x, y \in X$, $d(x, y) = d(y, x)$.

(iii) For all $x, y, z \in X$, $d(x, z) \leqslant d(x, y) + d(y, z)$.

A set for which a metric function exists is called a metric space. One of the common metric spaces is $\mathbb{R}^n$ which is a metric space for all $n \in \mathbb{N}^+$, we will use this space in this thesis.

A common metric function for this space is the *p-norm* which is defined in the following manner.

Definition 2.2. *Let $p \in \mathbb{N}^+$, we define the distance function $d : \mathbb{R}^n \times \mathbb{R}^n \to [0, \infty)$ to be:*

$$d(x, y) = \left(\sum_{i=1}^{n} |x_i - y_i|^p \right)^{1/p}.$$

Note that for $p = 2$ we get the *Euclidean distance* which we will use throughout this paper. In particular we will use $d(x, y) = \sqrt{\sum_{i=1}^{n} (x_i - y_i)^2}$. We will also use the concept of

norm. The norm of a vector x is its distance from the origin, formally $\|x\| = d(x, 0)$, for the Euclidean distance this yields $\|x\| = \sqrt{\sum_{i=1}^{n} x_i^2}$. We will sometimes use $\|x - y\|$ to the denote Euclidean distance between x and y, this is true as $\|x - y\| = \sqrt{\sum_{i=1}^{n}(x_i - y_i)^2} = d(x, y)$.

We denote by $B_r(c)$ the open ball with radius r centered at c. Formally $B_r(c) = \{x : d(x, c) < r\}$. As stated before, we are using the Eclidean distance thus yielding $B_r(c) = \{x : \sum_{i=1}^{n} |x_i - c_i|^2 < r^2\}$.

We next define the configuration space and the free space. For a single robot R, its configuration space $\mathcal{C}$ is the space of parametric representations of all the possible configurations of the robot, namely every point in $\mathcal{C}$ represents one placement of the robot in the physical world. Given a set of obstacles $\{\mathcal{O}_1, \ldots \mathcal{O}_k\}$, we define the free space $\mathcal{C}^f$ to be the set of all configurations in $\mathcal{C}$ in which the robot does not collide with the obstacles. Formally, $\mathcal{C}^f = \{c \in \mathcal{C} : R(c) \cap (\bigcup_{i=1}^{k} \mathcal{O}_i) = \varnothing\}$, where $R(c)$ represents the physical placement of R at configuration c. For multi-robot scenarios, the configuration space is simply the Cartesian product of the configuration spaces of the individual robots. For the free space, an additional constraint is added, namely that the robots must not collide with one another; this is formally defined in Section 4.1.

Next we define a trajectory. Given a configuration space $\mathcal{C}$, a trajectory σ is a function $\sigma : [0, 1] \to \mathcal{C}$, where $\sigma(t)$ is the configuration of the robot at time $t \in [0, 1]$. A trajectory will be called *collision-free* if it is contained in the free space. Formally, given a free space $\mathcal{C}^f \subset \mathcal{C}$, σ will be called collision-free if $\sigma(t) \in \mathcal{C}^f$ for all $0 \leqslant t \leqslant 1$. We note that this definition holds for both single and multi-robot systems, where the main difference between the two cases is the underlying free space.

3

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 follow with our main theoretical result for the single-robot case, pertaining to the quality of solutions obtained using PRM with staggered-grid sampling. Finally we conclude with a comparison with previous work.

3.1 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 $\sigma : [0,1] \to \mathcal{C}^f$ 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 $\bigcup_{0 \leqslant t \leqslant 1} B_\delta(\sigma(t)) \subseteq \mathcal{C}^f$, for $B_\delta(\sigma(t))$ being the d-dimensional closed Euclidean ball with radius δ centered at $\sigma(t)$. We say that $\mathcal{M}$ is δ-clear if there exists a trajectory σ with clearance δ that solves $\mathcal{M}$.

3.2 Probabilistic roadmaps and sample sets

We provide a formal definition of the Probabilistic Roadmap (PRM) method [25], 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., [50, 55]). PRM is also instrumental to our result both on single-robot motion-planning in this section and on multi-robot motion-planning in Section 4.

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

$$\begin{aligned} V :=& (\mathcal{X} \cup \{x^s, x^g\}) \cap \mathcal{C}^f, \text{and} \\ E :=& \left\{ \{v, u\} \in V \times V : \|v - u\| \leqslant r, \mathrm{CH}(\{v, u\}) \subset \mathcal{C}^f \right\}, \end{aligned}$$

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 [55].

Definition 3.1 (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*

$$d(G_{\mathcal{M}(\mathcal{X},r)}, x^s, x^g)) \leqslant (1 + \varepsilon) OPT_\delta,$$

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 $\mathcal{M}$.

3.3 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 3.1 for an illustration. An important property of this staggered grid, which we prove in Lemma 3.5 below, 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 [55]. 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 3.2 (Staggered grid)**.** *For given $\beta > 0, \gamma > 0$ the staggered grid $\mathcal{X}_{\beta,\gamma}$ is the union of two point sets $\mathcal{X}^1_{\beta,\gamma}, \mathcal{X}^2_{\beta,\gamma}$ in $[0,1]^d$, where*

$$\mathcal{X}^1_{\beta,\gamma} = \left\{(p_1, p_2, ..., p_d) : p_i = \gamma + (2k-1)w,\ 1 \leqslant k \leqslant \left\lceil \frac{1-2\gamma}{2w} \right\rceil,\ 1 \leqslant i \leqslant d \right\},$$
$$\mathcal{X}^2_{\beta,\gamma} = \left\{(p_1, p_2, ..., p_d) : p_i = \gamma + 2kw,\ 0 \leqslant k \leqslant \left\lceil \frac{1-2\gamma}{2w} \right\rceil,\ 1 \leqslant i \leqslant d \right\},$$

and $w = \frac{\beta\sqrt{2}}{\sqrt{d}}$.

Notice that both $\mathcal{X}^1_{\beta,\gamma}$ and $\mathcal{X}^2_{\beta,\gamma}$ are square grids of side length $2w$, and the total number of points in the construction is $\left(\left\lceil \frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta} \right\rceil\right)^d + \left(\left\lceil \frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta} \right\rceil + 1\right)^d$. We illustrate the staggered grid in Fig 3.1. A detailed comparative size analysis for the staggered grid can be found in Section 6.

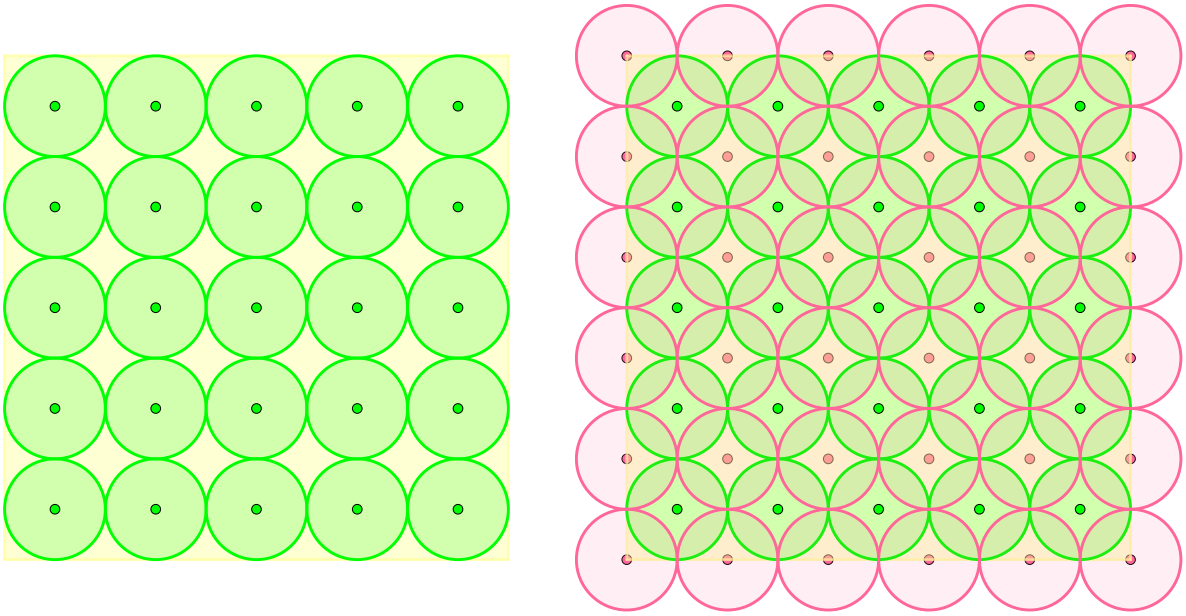

Figure 3.1: Illustration of Definition 3.2 in two dimensions, for $\beta = 0.08$, $\gamma = 0.1$. On the left we visualize the first layer $\mathcal{X}^1_{\beta,\gamma}$ of the staggered grid (green discs). On the right we add the second layer $\mathcal{X}^2_{\beta,\gamma}$ (red discs). The centers of the discs are the points of $\mathcal{X}_{\beta,\delta}$.

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 ($\varepsilon$, δ)-complete, where each of β and r depends on both ε and δ, and γ is equal to δ.

Theorem 3.3 (Sufficient conditions for (ε, δ)-completeness)**.** *Fix a stretch parameter $\varepsilon > 0$ and clearance $\delta > 0$. For a sampling distribution $\mathcal{X} = \mathcal{X}_{\alpha\delta,\delta}$, 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.*

3.4 Proof of Theorem 3.3

In this section we provide a proof of Theorem 3.3. First, we introduce the concept of β-cover[1] [18]. Informally, a β-cover of a domain is a set $\mathcal{X}$ of points such that no point of the domain is too far from some point in $\mathcal{X}$ (defined formally below). Next we show that the staggered grid $\mathcal{X}_{\beta,\delta}$ is a β-cover of $[\delta, 1-\delta]^d$. Then we exploit this property in the proof for Theorem 3.3.

Definition 3.4. *For a given $\beta > 0$, a set $\mathcal{X} \subset \mathbb{R}^d$ is a β-cover for a set $A \subset \mathbb{R}^d$ if for every $a \in A$, there exists $s \in \mathcal{X}$ such that $\|a - s\| \leqslant \beta$*

First, we prove that $\mathcal{X}_{\beta,\gamma}$ forms a β-cover over $[\gamma, 1-\gamma]^d$.

Lemma 3.5. *Let $\beta > 0, \gamma > 0$ and define $A_\gamma = [\gamma, 1-\gamma]^d$. Then the staggered grid $\mathcal{X}_{\beta,\gamma}$ is a β-cover for A_γ.*

Proof. Recall that the width w of a cell in each of the two grids that form the staggered grid $\mathcal{X}_{\beta,\gamma}$ is $\beta\sqrt{2}/\sqrt{d}$. We divide the d-dimensional $(1-2\gamma)$-hypercube, A_γ, by hyperplanes as follows. For each coordinate x_i we define the hyperplanes $x_i = \gamma + 2kw$, for $k \in \{0, 1, \ldots, \lceil \frac{1-2\gamma}{2w} \rceil\}$. Jointly, those hyperplanes induce a partition of A_γ into a set of hypercubes H: each hypercube $h \in H$ is of edge length $2w$, its center is a point from $\mathcal{X}^1_{\beta,\gamma}$, and its vertices are points from $\mathcal{X}^2_{\beta,\gamma}$. See Figure 3.2 for an illustration in $\mathbb{R}^2$.

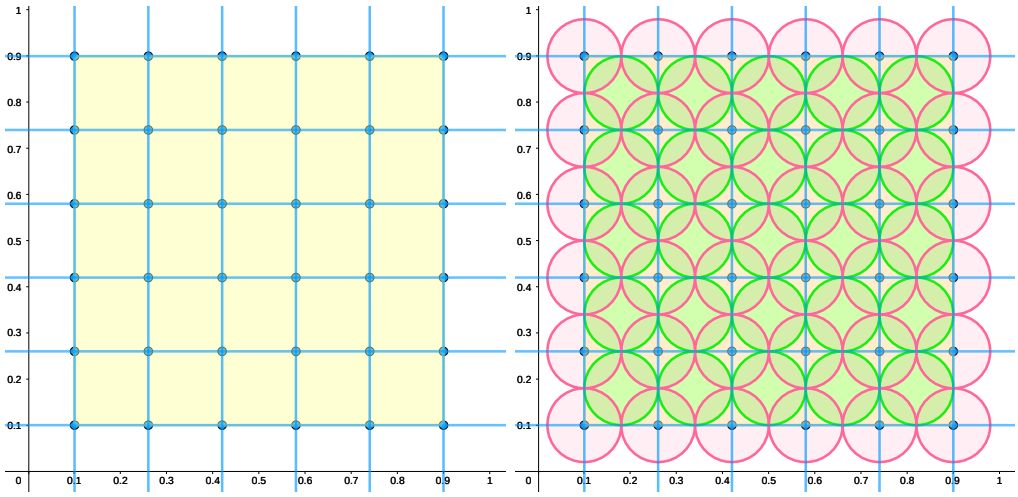

Figure 3.2: Illustration of the hyperplanes partition for Lemma 3.5 in two dimensions, for $\gamma = 0.1$, $\beta = 0.08$, and $w = 0.08$. In both sub-figures the yellow cube represents $[\gamma, 1-\gamma]^d$, and the blue lines are the hyperplanes orthogonal to x_i. On the left we demonstrate how the cube is split into smaller $2w$-cubes. On the right we demonstrate that these $2w$-cubes are covered with β-balls. The green balls are β-balls centered at points of $\mathcal{X}^1_{\beta,\gamma}$, while the red balls are β-balls centered at points of $\mathcal{X}^2_{\beta,\gamma}$.

[1] The common naming for this term is ε-cover. We use the letter β to avoid confusion with the stretch factor which is called ε in this paper.

We now prove that each $h \in H$ induced by our hyperplanes is fully covered using the β-balls centered at its vertices and at its center. Formally, for a given $h \in H$, let S_h be the set of points comprising the vertices of h and the center point of h (i.e., $|S_h| = 2^d + 1$). We show that S_h is a β-cover for h.

Let $s \in S_h$ be the center point of h. Without loss of generality, assume that s is at the origin. Let $B_\beta(s)$ be the β-ball centered at s. Let $p = (p_1, .., p_d)$ be a d-dimensional point such that $p \in h$ and $p \notin B_\beta(s)$. We further assume that $\forall i \in [1..d], 0 \leqslant p_i \leqslant w$; all the other cases can be handled symmetrically. See Figure 3.3 for an illustration in $\mathbb{R}^2$. Thus, since $p \notin B_\beta(s)$ we have that $\|s - p\| > \beta$, which implies that $\beta^2 < \sum_{i=1}^{d} p_i^2$.

Next, recall that $s_w := (w, ..., w) \in S$ since the hypercubes edge length is $2w$. We upper bound the distance of p from s_w in the following manner:

$$\begin{aligned}\|s_w - p\|^2 &= \sum_{i=1}^{d}(w - p_i)^2 = \sum_{i=1}^{d}(w^2 - 2wp_i + p_i^2) \\ &\leqslant \sum_{i=1}^{d}(w^2 - 2p_i^2 + p_i^2) = \sum_{i=1}^{d} w^2 - \sum_{i=1}^{d} p_i^2 \\ &< \sum_{i=1}^{d} w^2 - \beta^2,\end{aligned}$$

where the first inequality follows from the fact that $p_i \leqslant w$, and the second inequality follows from $\beta^2 < \sum_{i=1}^{d} p_i^2$.

Finally, by subsisting $w = \beta\sqrt{2}/\sqrt{d}$, we obtain

$$\|s_w - p\|^2 < \sum_{i=1}^{d} \frac{2\beta^2}{d} - \beta^2 = \beta^2,$$

which implies that p is covered by the β-ball centered at s_w.

We have proven that $\mathcal{X}_{\beta,\gamma}$ provides a β-cover for each $h \in H$, which also guarantees that $\mathcal{X}_{\beta,\gamma}$ is a β-cover for the larger hypercube $[\gamma, 1-\gamma]^d$. Thus we conclude that $\mathcal{X}_{\beta,\gamma}$ is a β-cover for A_γ. □

Next, we exploit the above property for showing that by appropriately setting the values of β, γ and r with respect to ε, δ, we obtain a good approximation of any collision-free trajectory, which implies that $(\mathcal{X}_{\beta,\gamma}, r)$ is (ε, δ)-complete. To do so, we prove the following lemma, which is an extended version of [55, Theorem 2]. This lemma, will also be instrumental in extending the theory we develop for the single-robot case, to the multi-robot setting. In order to state the theorem, we introduce additional notation regarding a sequence of points along a given trajectory of certain spacing.

Definition 3.6. *Let σ be a collision-free trajectory, and let $\rho > 0$ be a step size such that $\|\sigma(0) - \sigma(1)\| \geqslant \rho$. Then the time sequence $T_\sigma^\rho = (\tau_0, \tau_1, \ldots, \tau_\ell)$, where $\tau_0 < \tau_1 \ldots < \tau_\ell$, is defined in the following manner for some $\ell \in \mathbb{N}_{>0}$: $\tau_0 = 0, \tau_\ell = 1$, for all $1 \leqslant i \leqslant \ell - 1$, $\|\sigma(\tau_i) - \sigma(\tau_{i-1})\| = \rho$, and $\|\sigma(\tau_\ell) - \sigma(\tau_{\ell-1})\| \leqslant \rho$.*

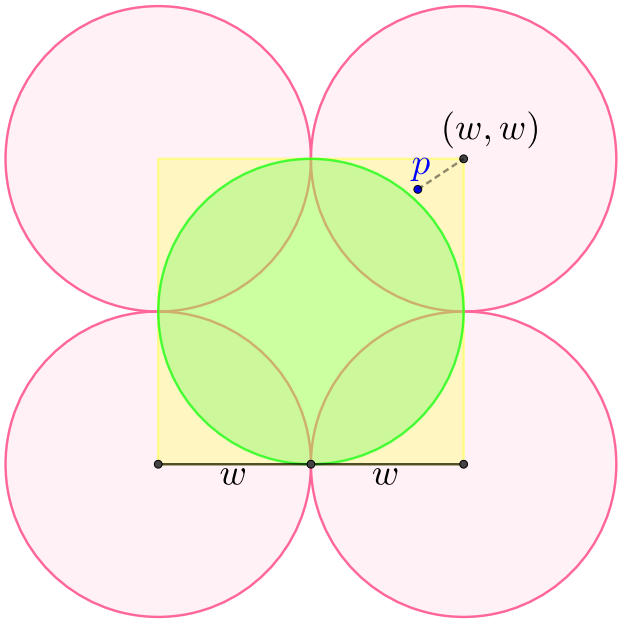

Figure 3.3: Illustration of the covering of a $2w$-hypercube of Lemma 3.5 in two dimensions. The point p is the point tested in the lemma. The red balls are β-balls centered on the cube vertices. The green ball is the β-ball centered in the cubes center called $B_\beta(s)$. The dashed line represents the distance between p and (w, w). We note that in higher dimensions the red balls would overlap as w becomes smaller in relation to β.

Lemma 3.7. *For given $\beta, \gamma, \rho > 0$, such that $\beta^2 + (\rho/2)^2 \leqslant \gamma^2$, let σ be a γ-clear solution for $\mathcal{M} = (\mathcal{C}^f, x^s, x^g)$, and let $T^\rho_\sigma = (\tau_0, \tau_1, \ldots, \tau_\ell)$ be a time sequence as in Definition 3.6. Then there exists a point set $Z^\rho_\sigma = (z_0, \ldots, z_\ell)$, where $z_i \in \mathcal{X}_{\beta,\gamma} \cup \{x^s, x^g\}$ for all $0 \leqslant i \leqslant \ell$, such that the following properties holds:*

(i) $\|z_i - \sigma(\tau_i)\| \leqslant \beta$, *for all* $0 \leqslant i \leqslant \ell$;

(ii) $\|z - \sigma(\tau_{i-1})\| \leqslant \beta + \rho$, *for all* $1 \leqslant i \leqslant \ell$, $z \in CH(z_{i-1}, z_i)$;

(iii) $\|z - \sigma(\tau_i)\| \leqslant \beta + \rho$, *for all* $1 \leqslant i \leqslant \ell$, $z \in CH(z_{i-1}, z_i)$;

(iv) $\|z_i - z_{i-1}\| \leqslant 2\beta + \rho$, *for all* $1 \leqslant i \leqslant \ell$;

(v) $CH(z_{i-1}, z_i)$ *is collision free, for all* $1 \leqslant i \leqslant \ell$. *That is* $CH(z_{i-1}, z_i) \subset \mathcal{C}^f$, *for all* $1 \leqslant i \leqslant \ell$;

(vi) $\sum_{1 \leqslant i \leqslant \ell} \|z_i - z_{i-1}\| \leqslant \left(1 + \frac{2\beta}{\rho}\right) \|\sigma\|$.

Proof. Define, as before, $A_\gamma = [\gamma, 1 - \gamma]^d$. Notice that since σ is γ-clear then $\sigma(t) \in A_\gamma$ for all $0 \leqslant t \leqslant 1$ (and in particular for all $\tau_i \in T^\rho_\sigma$). Next, we address each individual component of the theorem.

Property (i): This follows directly from Lemma 3.5 as $\mathcal{X}_{\beta,\gamma}$ is a β-cover for A_γ and $\sigma(\tau_i) \in A_\gamma$.

Properties (ii) and (iii): First, we upper bound the expression $\|z_i - \sigma(\tau_{i-1})\|$ (the proof for Property (iii) is symmetric). By definition of T^ρ_σ we have that $\|\sigma(\tau_i) - \sigma(\tau_{i-1})\| = \rho$, for $1 \leqslant i \leqslant \ell - 1$. Thus, using Property (i), we have $\|z_i - \sigma(\tau_i)\| \leqslant \beta$. Using the triangle inequality it follows that,

$$\begin{aligned}\|z_i - \sigma(\tau_{i-1})\| &\leqslant \|z_i - \sigma(\tau_i)\| + \|\sigma(\tau_i) - \sigma(\tau_{i-1})\| \\ &\leqslant \beta + \rho.\end{aligned}$$

Let $z \in \mathrm{CH}(z_{i-1}, z_i)$, that is $z = \lambda z_{i-1} + (1-\lambda) z_i$ for some λ such that $0 \leqslant \lambda \leqslant 1$. Thus,

$$\begin{aligned}\|z - \sigma(\tau_{i-1})\| &= \|\lambda z_{i-1} + (1-\lambda) z_i - \sigma(\tau_{i-1})\| \\ &= \|\lambda z_{i-1} - \lambda\sigma(\tau_{i-1}) + (1-\lambda) z_i - (1-\lambda)\sigma(\tau_{i-1})\| \\ &\leqslant \|\lambda z_{i-1} - \lambda\sigma(\tau_{i-1})\| + \|(1-\lambda) z_i - (1-\lambda)\sigma(\tau_{i-1})\| \\ &= \lambda \|z_{i-1} - \sigma(\tau_{i-1})\| + (1-\lambda) \|z_i - \sigma(\tau_{i-1})\|,\end{aligned}$$

where the triangle inequality was used in the second-to-last transition. Due to Property (i) and the fact that $\|z_i - \sigma(\tau_{i-1})\| \leqslant \beta + \rho$ it follows that

$$\begin{aligned}\|z - \sigma(\tau_{i-1})\| &\leqslant \lambda\beta + (1-\lambda)(\beta + \rho) \\ &= \beta + (1-\lambda)\rho \\ &\leqslant \beta + \rho.\end{aligned}$$

Property (iv): This follows from properties (i) and (ii), and triangle inequality:

$$\begin{aligned}\|z_i - z_{i-1}\| &\leqslant \|z_i - \sigma(\tau_{i-1})\| + \|\sigma(\tau_{i-1}) - z_{i-1}\| \\ &\leqslant 2\beta + \rho.\end{aligned}$$

Property (v): Fix $1 \leqslant i \leqslant \ell$. For any $\eta \in [0,1]$ define $u_\eta = \eta z_{i-1} + (1-\eta) z_i$. To prove that Property (v) holds we will show that $u_\eta \in \mathcal{C}^f$ for every $\eta \in [0,1]$.

Fix $\eta \in [0,1]$. Since $\sigma(\tau_{i-1}), \sigma(\tau_i)$ are points on a γ-clear path, the γ-balls $B_\gamma(\sigma(\tau_{i-1}))$, $B_\gamma(\sigma(\tau_i))$ are collision-free. Thus it is sufficient to prove that $u_\eta \in B_\gamma(\sigma(\tau_{i-1})) \cup B_\gamma(\sigma(\tau_i))$. Next let v_η be the closest point to u_η on the straight line between $\sigma(\tau_{i-1})$ and $\sigma(\tau_i)$. That is, $v_\eta = \arg\min_{v \in \mathrm{CH}(\{\sigma(\tau_{i-1}), \sigma(\tau_i)\})} \|u_\eta - v\|$.

We also define $v'_\eta := \eta\sigma(\tau_{i-1}) + (1-\eta)\sigma(\tau_i)$, and prove that $\|u_\eta - v'_\eta\| \leqslant \beta$. This implies that $\|u_\eta - v_\eta\| \leqslant \beta$ as v_η is at least as close as v'_η to u_η. Indeed, using Property (i) and the triangle inequality we have

$$\begin{aligned}\|u_\eta - v'_\eta\| &\leqslant \|\eta z_{i-1} + (1-\eta) z_i - (\eta\sigma(\tau_{i-1}) + (1-\eta)\sigma(\tau_i))\| \\ &\leqslant \eta \|z_{i-1} - \sigma(\tau_{i-1})\| + (1-\eta) \|z_i - \sigma(\tau_i)\| \\ &= \beta.\end{aligned}$$

Building upon this result, we consider several cases with respect to the position of v_η. If v_η is one of the endpoints, that is $v_\eta = \sigma(\tau_{i-1})$ or $v_\eta = \sigma(\tau_i)$, we have that u_η is in a β-ball around this point. Since $\beta^2 + (\rho/2)^2 \leqslant \gamma^2$, and all are positive, we know that $\beta \leqslant \gamma$. Thus if v_η is one of the endpoints we are done as $u_\eta \in B_\beta(\sigma(\tau_{i-1}))$ or $u_\eta \in B_\beta(\sigma(\tau_i))$.

Assuming v_η is not one of the endpoints we can write $v_\eta = \eta^* \sigma(\tau_{i-1}) + (1-\eta^*)\sigma(\tau_i)$ where

$$\eta^* = \arg\min_{\lambda \in (0,1)} \|u_\eta - \lambda\sigma(\tau_{i-1}) + (1-\lambda)\sigma(\tau_i)\| .$$

In this case, the point v_η is the projection of u_η onto the line through $\sigma(\tau_{i-1})$ and $\sigma(\tau_i)$. Therefore,

$$\langle \sigma(\tau_i) - \sigma(\tau_{i-1}), u_\eta - v_\eta \rangle = 0.$$

Since $\{\sigma(\tau_{i-1}), \sigma(\tau_i), v_\eta\}$ are collinear we get that,

$$\langle v_\eta - \sigma(\tau_{i-1}), u_\eta - v_\eta \rangle = \langle v_\eta - \sigma(\tau_i), u_\eta - v_\eta \rangle = 0. \tag{3.1}$$

Those equations are illustrated in Fig 3.4.

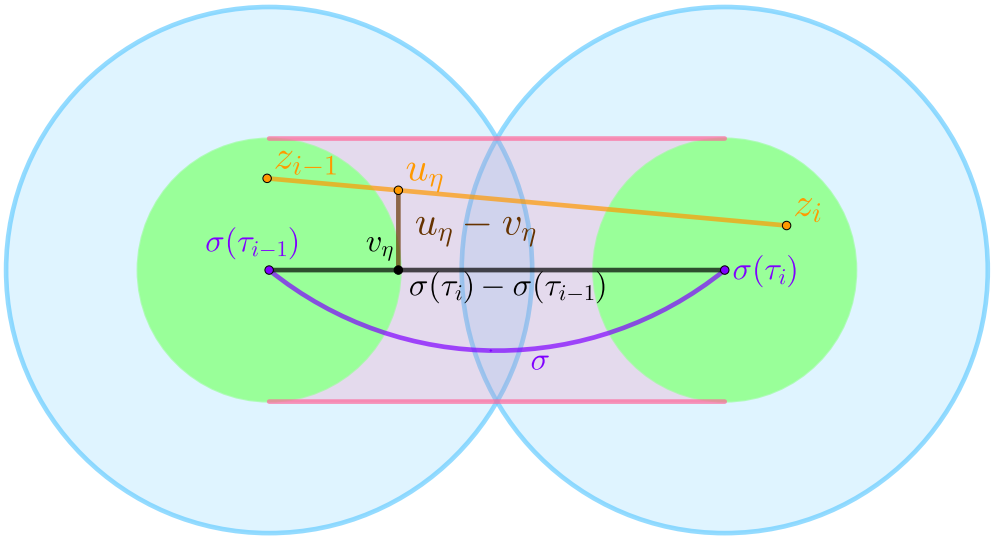

Figure 3.4: Illustration for the proof of Lemma 3.7. The purple curve represents the trajectory σ. The green balls are the β-balls around $\sigma(\tau_{i-1})$ and $\sigma(\tau_i)$. The convex hull of the green balls (namely, the green balls together with the red region in-between) contains the line segment (in orange) $\mathrm{CH}(\{z_{i-1}, z_i\})$, and we show that this region is collision-free. The blue balls are collision-free γ-balls around $\sigma(\tau_{i-1})$ and $\sigma(\tau_i)$. The black line is $\mathrm{CH}(\{\sigma(\tau_{i-1}), \sigma(\tau_i)\})$.

We now consider two separate cases with respect to η^*. First assume $\eta^* \in (0, 1/2]$. Using

Equation (3.1) and Property (i), it follows that

$$\begin{aligned}
\|u_\eta - \sigma(\tau_i)\|^2 &= \|(u_\eta - v_\eta) + (v_\eta - \sigma(\tau_i)\|^2 \\
&= \|u_\eta - v_\eta\|^2 + \|v_\eta - \sigma(\tau_i)\|^2 \\
&\leqslant \beta^2 + \|\eta^* \sigma(\tau_{i-1}) + (1-\eta^*)\sigma(\tau_i) - \sigma(\tau_i)\|^2 \\
&= \beta^2 + (\eta^*)^2 \|\sigma(\tau_{i-1}) - \sigma(\tau_i)\|^2 \\
&= \beta^2 + (\eta^*)^2 \rho^2 \\
&\leqslant \beta^2 + (1/2)^2 \rho^2,
\end{aligned}$$

where the last inequality follows from $\eta^* \leqslant 1/2$. Finally by plugging-in $\beta^2 + (\rho/2)^2 \leqslant \gamma^2$ we get $\|u_\eta - \sigma(\tau_i)\|^2 \leqslant \gamma^2$, which implies that $u_\eta \in B_\gamma(\sigma(\tau_i))$.

Due to symmetry, we can show that for $\eta^* \in (1/2, 1)$, it follows that $\|u_\eta - \sigma(\tau_{i-1})\|^2 \leqslant \gamma^2$, which again implies $u_\eta \in B_\gamma(\sigma(\tau_{i-1}))$. To conclude, we showed that for all $\eta \in [0,1]$ it holds that $u_\eta \in B_\gamma(\sigma(\tau_{i-1})) \cup B_\gamma(\sigma(\tau_i))$, which implies that $\mathrm{CH}\,(z_{i-1}, z_i) \subset \mathcal{C}^f$.

Property (vi): Due to Property (iv) we have that, $\|z_i - z_{i-1}\| \leqslant 2\beta + \rho$ for all $1 \leqslant i \leqslant \ell$. By construction, our solution path through the z_i's coincides with the starting and end points of σ, namely, $z_0 = \sigma(0), z_1 = \sigma(1)$. Hence, for the first and last segment of the motion we have $\|z_i - z_{i-1}\| \leqslant \beta + \|\sigma(\tau_i) - \sigma(\tau_{i-1})\|$. Using this we get,

$$\begin{aligned}
\sum_{1 \leqslant i \leqslant \ell} \|z_i - z_{i-1}\| &= \|z_1 - z_0\| + \|z_\ell - z_{\ell-1}\| + \sum_{2 \leqslant i \leqslant \ell-1} \|z_i - z_{i-1}\| \\
&\leqslant \beta + \rho + \beta + \|\sigma(\tau_\ell) - \sigma(\tau_{\ell-1})\| + \sum_{2 \leqslant i \leqslant \ell-1} (2\beta + \rho) \\
&= 2\beta + \rho + (\ell - 2) \cdot (2\beta + \rho) + \|\sigma(\tau_\ell) - \sigma(\tau_{\ell-1})\| \\
&= (\ell - 1) \cdot (2\beta + \rho) + \|\sigma(\tau_\ell) - \sigma(\tau_{\ell-1})\|.
\end{aligned}$$

Additionally, due to equality $\|\sigma(\tau_i) - \sigma(\tau_{i-1})\| = \rho$ for all $1 \leqslant i \leqslant \ell - 1$, we have that

$$\|\sigma\| \geqslant (\ell - 1) \cdot \rho + \|\sigma(\tau_\ell) - \sigma(\tau_{\ell-1})\|.$$

Thus,

$$\begin{aligned}
\frac{\sum_{1 \leqslant i \leqslant \ell} \|z_i - z_{i-1}\|}{\|\sigma\|} &\leqslant \frac{(\ell - 1) \cdot (2\beta + \rho) + \|\sigma(\tau_\ell) - \sigma(\tau_{\ell-1})\|}{(\ell - 1) \cdot \rho + \|\sigma(\tau_\ell) - \sigma(\tau_{\ell-1})\|} \\
&= \frac{(2\beta + \rho)}{\rho} \cdot \frac{(\ell - 1) + \|\sigma(\tau_\ell) - \sigma(\tau_{\ell-1})\|/(2\beta + \rho)}{(\ell - 1) + \|\sigma(\tau_\ell) - \sigma(\tau_{\ell-1})\|/\rho} \\
&\leqslant \frac{(2\beta + \rho)}{\rho} \\
&= 1 + \frac{2\beta}{\rho},
\end{aligned}$$

thus concluding this proof. □

It remains to prove Theorem 3.3. To use Lemma 3.7, we set $\gamma = \delta, \beta = \alpha\delta, \rho = \frac{2}{\sqrt{1+\varepsilon^2}}\delta$, and $r = \frac{2(\varepsilon+1)}{\sqrt{1+\varepsilon^2}}\delta$, for $\alpha = \varepsilon/\sqrt{1+\varepsilon^2}$, so that the condition $\beta^2 + (\rho/2)^2 \leqslant \gamma^2$ holds.

By setting $r = 2\beta + \rho = 2(\alpha + \frac{1}{\sqrt{1+\varepsilon^2}})\delta = \frac{2(\varepsilon+1)}{\sqrt{1+\varepsilon^2}}\delta$, it follows from Lemma 3.7 (iv) that for all $1 \leqslant i \leqslant \ell$ the PRM edge connecting z_i and z_{i-1} is considered in the PRM construction. Additionally, since $\frac{2\beta}{\rho} = \varepsilon$, it holds that $\|\widehat{\sigma}\| \leqslant (1+\varepsilon)\|\sigma\|$ where $\widehat{\sigma}$ denotes the piece-wise linear trajectory induced by Z^{ρ}_{σ}, which concludes the proof.

3.5 Comparison with previous work

In this section we compare our results with the previous results, by Tsao et al [55]. We first compare our result with their upper bound in the asymptotic case; in this comparison we observe an exponential improvement by a factor of 1.3687^d. We follow by comparing our result with their lower bound for $\varepsilon = \infty$ (again in the asymptotic case); in this comparison we observe that our result is exponentially worse than the lower bound by a factor of 1.4611^d. Finally, we present these comparisons for concrete values of δ and ε in Table 3.1.

We now compare the minimal sample size required to guarantee (ε, δ)-completeness for the same single-robot setting, between the previous approach [55], and the staggered grid as proposed here. We denote by $\mathcal{X}_{\text{prev}}$ the size of the minimal sample set needed using the bound from [55], namely

$$\mathcal{X}_{\text{prev}} \approx \sqrt{\pi d}\left(\sqrt{\frac{2d}{\pi e}} \cdot \frac{1-(2-\alpha)\delta}{\alpha\delta}\right)^d,$$

where, as above, $\alpha = \frac{\varepsilon}{\sqrt{1+\varepsilon^2}}$. We denote by $\mathcal{X}_{\text{curr}} = |\mathcal{X}_{\alpha\delta,\delta}|$ the size of the minimal sample set needed using the bound for the staggered grid, as stated in Theorem 3.3, namely,

$$\mathcal{X}_{\text{curr}} = \left(\left\lceil \frac{(1-2\delta)\sqrt{d}}{\sqrt{8}\alpha\delta} \right\rceil\right)^d + \left(\left\lceil \frac{(1-2\delta)\sqrt{d}}{\sqrt{8}\alpha\delta} \right\rceil + 1\right)^d.$$

We compare the quantities $\mathcal{X}_{\text{prev}}$ and $\mathcal{X}_{\text{curr}}$ for small values of ε and δ (we mention that the radius derived in [55] is equal to the one we are using here). We consider the ratio $\frac{\mathcal{X}_{\text{prev}}}{\mathcal{X}_{\text{curr}}}$ in the asymptotic regime where the values ε and δ tend to zero in the following equation (we ignore the rounding errors for both results as they do not change the asymptotic results):

$$
\begin{aligned}
\frac{\mathcal{X}_{\text{prev}}}{\mathcal{X}_{\text{curr}}} &= \frac{\sqrt{\pi d}\left(\sqrt{\frac{2d}{\pi e}}\cdot\frac{1-(2-\alpha)\delta}{\alpha\delta}\right)^d}{\left(\frac{(1-2\delta)\sqrt{d}}{\sqrt{8}\alpha\delta}\right)^d+\left(\frac{(1-2\delta)\sqrt{d}}{\sqrt{8}\alpha\delta}+1\right)^d} \\
&= \frac{\sqrt{\pi d}\left(\sqrt{\frac{2d}{\pi e}}\cdot\frac{1-2\delta+\boldsymbol{\alpha\delta}}{\alpha\delta}\right)^d}{\left(\frac{(1-2\delta)\sqrt{d}}{\sqrt{8}\alpha\delta}\right)^d+\left(\frac{(1-2\delta)\sqrt{d}}{\sqrt{8}\alpha\delta}+\mathbf{1}\right)^d} \\
&\overset{(a)}{\approx} \frac{\sqrt{\pi d}\left(\sqrt{\frac{2d}{\pi e}}\cdot\frac{1-2\delta}{\alpha\delta}\right)^d}{2\left(\frac{(1-2\delta)\sqrt{d}}{\sqrt{8}\alpha\delta}\right)^d} \\
&= \frac{\sqrt{\pi d}}{2}\left(\sqrt{\frac{16}{\pi e}}\right)^d \\
&\approx \frac{\sqrt{\pi d}}{2}1.3687^d\,,
\end{aligned}
\tag{3.2}
$$

where the approximation in (a) is due to ignoring the highlighted $\alpha\delta$ factor in the numerator of the previous expression, and also ignoring the highlighted unit value 1 in the denominator (the influence of both values tends to 0 as δ or ε tend to 0). This implies that our result yields samples sets that are smaller by an exponential factor in d than the previous work.

Next, we compare between the number of points we need in order to find a (∞,δ)-complete trajectory and the lower bound derived in [55, Theorem 1] for this quantity, which was only derived there for the case $\varepsilon=\infty$. We denote the latter by $\mathcal{X}_{\text{LB}}$, which is proved in [55] to be

$$
\mathcal{X}_{\text{LB}} \approx \sqrt{\frac{e}{2}}\left(1-\frac{2\delta}{1-2\delta}\right)^2\left(\sqrt{\frac{d-1}{2\pi e}}\cdot\frac{1-2\delta}{\delta}\right)^d\,. \tag{3.3}
$$

We compare the quantities $\mathcal{X}_{\text{LB}}$ and $\mathcal{X}_{\text{curr}}$ for small values of δ and using $\varepsilon=\infty$. We consider the ratio $\frac{\mathcal{X}_{\text{curr}}}{\mathcal{X}_{\text{LB}}}$ in the asymptotic regime where the value of δ tends to zero in the following equation (we ignore the rounding errors for both results as they do not change the asymptotic results):

$$
\begin{aligned}
\frac{\mathcal{X}_{\text{curr}}}{\mathcal{X}_{\text{LB}}} &= \frac{\left(\frac{(1-2\delta)\sqrt{d}}{\sqrt{8}\alpha\delta}\right)^d + \left(\frac{(1-2\delta)\sqrt{d}}{\sqrt{8}\alpha\delta} + 1\right)^d}{\sqrt{\frac{e}{2}}\left(1-\frac{2\delta}{1-2\delta}\right)^2\left(\sqrt{\frac{d-1}{2\pi e}}\cdot\frac{1-2\delta}{\delta}\right)^d} \\
&= \frac{\left(\frac{(1-2\delta)\sqrt{d}}{\sqrt{8}\delta}\right)^d + \left(\frac{(1-2\delta)\sqrt{d}}{\sqrt{8}\delta} + \mathbf{1}\right)^d}{\sqrt{\frac{e}{2}}\left(\mathbf{1}-\frac{\boldsymbol{2\delta}}{\mathbf{1-2\delta}}\right)^{\mathbf{2}}\left(\sqrt{\frac{d-1}{2\pi e}}\cdot\frac{1-2\delta}{\delta}\right)^d} \\
&\overset{(a)}{\approx} \frac{2\left(\frac{(1-2\delta)\sqrt{d}}{\sqrt{8}\alpha\delta}\right)^d}{\sqrt{\frac{e}{2}}\left(\sqrt{\frac{d-1}{2\pi e}}\cdot\frac{1-2\delta}{\delta}\right)^d} \\
&= \sqrt{\frac{8}{e}}\left(\sqrt{\frac{\pi e}{4}}\right)^d\sqrt{\left(\frac{d}{d-1}\right)^d} \\
&\approx \sqrt{\frac{8}{e}}\sqrt{\left(\frac{d}{d-1}\right)^d}\cdot 1.4611^d\,,
\end{aligned}
\tag{3.4}
$$

where the approximation in (a) is due to ignoring the highlighted unit value 1 in the numerator of the previous expression, whose influence tends to 0 as δ tend to 0, and also ignoring the highlighted squared value in the denominator, whose value tends to 1 as δ tend to 0.

Finally, by noticing that for $d \geqslant 2$ it holds that

$$
e \leqslant \left(\frac{d}{d-1}\right)^d \leqslant 4,
$$

it follows that

$$
2\sqrt{2}\cdot 1.4611^d \leqslant \frac{\mathcal{X}_{\text{curr}}}{\mathcal{X}_{\text{LB}}} \leqslant 4\sqrt{\frac{2}{e}}\cdot 1.4611^d.
$$

This shows that our upper bound is at most a factor of roughly 1.4611^d away from the lower bound. These results may have further repercussions in the context of economical covering of the cube. We discuss this issue below in Section 6.

To conclude this section, we provide in Table 3.1 a comparison of sufficient and necessary samples sizes for (ε, δ)-completeness for various values of the clearance δ and the stretch ε in dimensions $d \in \{2, \ldots, 6\}$. In particular, for each combination of parameter values we report the following three quantities: (I) A lower bound (Eq. 3.3) on the *necessary* number of samples for (ε, δ)-completeness, as was derived in [55, Theorem 1], which applies only to the case $\varepsilon = \infty$. (II) The size of our sampling distribution, i.e., $\mathcal{X}_{\text{curr}}$, which guarantees (ε, δ)-completeness, according to Theorem 3.3. (III) An upper bound on the *sufficient* number of samples for (ε, δ)-completeness, denoted by $\mathcal{X}_{\text{prev}}$ above, as was derived in [55, Theorem 2].

		$\varepsilon = \infty$			$\varepsilon = 1$		$\varepsilon = 0.25$		$\varepsilon = 0.1$	
δ	d	lower bound ($\mathcal{X}_{\text{LB}}$)	$\mathcal{X}_{\text{curr}}$	$\mathcal{X}_{\text{prev}}$	$\mathcal{X}_{\text{curr}}$	$\mathcal{X}_{\text{prev}}$	$\mathcal{X}_{\text{curr}}$	$\mathcal{X}_{\text{prev}}$	$\mathcal{X}_{\text{curr}}$	$\mathcal{X}_{\text{prev}}$
0.25	2	0	5	12	13	19	61	109	265	567
	3	0	35	52	35	108	559	1510	4941	$1.79 \cdot 10^4$
	4	0	97	263	97	697	3697	$2.37 \cdot 10^4$	$1.16 \cdot 10^5$	$6.43 \cdot 10^5$
	5	0	275	1478	1267	5000	$4.96 \cdot 10^4$	$4.11 \cdot 10^5$	$2.47 \cdot 10^6$	$2.54 \cdot 10^7$
	6	0	793	9029	4825	$3.90 \cdot 10^4$	$7.94 \cdot 10^5$	$7.74 \cdot 10^6$	$8.11 \cdot 10^7$	$1.09 \cdot 10^9$
0.1	2	3	61	104	85	194	613	1471	3445	8437
	3	15	341	1393	855	3566	$1.99 \cdot 10^4$	$7.50 \cdot 10^4$	$2.58 \cdot 10^5$	$1.03 \cdot 10^6$
	4	88	3697	$2.13 \cdot 10^4$	$1.07 \cdot 10^4$	$7.45 \cdot 10^4$	$7.22 \cdot 10^5$	$4.33 \cdot 10^6$	$2.19 \cdot 10^7$	$1.42 \cdot 10^8$
	5	595	$4.96 \cdot 10^4$	$3.59 \cdot 10^5$	$1.59 \cdot 10^5$	$1.72 \cdot 10^6$	$3.16 \cdot 10^7$	$2.76 \cdot 10^8$	$2.23 \cdot 10^9$	$2.17 \cdot 10^{10}$
	6	4459	$3.80 \cdot 10^5$	$6.58 \cdot 10^6$	$2.77 \cdot 10^6$	$4.32 \cdot 10^7$	$1.32 \cdot 10^9$	$1.91 \cdot 10^{10}$	$2.46 \cdot 10^{11}$	$3.60 \cdot 10^{12}$
0.05	2	21	221	460	365	892	2965	7204	$1.67 \cdot 10^4$	$4.21 \cdot 10^4$
	3	234	3925	$1.31 \cdot 10^4$	9009	$3.54 \cdot 10^4$	$2.01 \cdot 10^5$	$8.13 \cdot 10^5$	$2.77 \cdot 10^6$	$1.15 \cdot 10^7$
	4	3152	$6.70 \cdot 10^4$	$4.23 \cdot 10^5$	$2.35 \cdot 10^5$	$1.59 \cdot 10^6$	$1.64 \cdot 10^7$	$1.04 \cdot 10^8$	$5.45 \cdot 10^8$	$3.55 \cdot 10^9$
	5	$4.82 \cdot 10^4$	$1.81 \cdot 10^6$	$1.51 \cdot 10^7$	$9.24 \cdot 10^6$	$7.88 \cdot 10^7$	$1.49 \cdot 10^9$	$1.46 \cdot 10^{10}$	$1.26 \cdot 10^{11}$	$1.21 \cdot 10^{12}$
	6	$8.13 \cdot 10^5$	$4.09 \cdot 10^7$	$5.83 \cdot 10^8$	$3.39 \cdot 10^8$	$4.25 \cdot 10^9$	$1.58 \cdot 10^{11}$	$2.24 \cdot 10^{12}$	$3.05 \cdot 10^{13}$	$4.49 \cdot 10^{14}$
0.01	2	734	5101	$1.25 \cdot 10^4$	9941	$2.48 \cdot 10^4$	$8.28 \cdot 10^4$	$2.09 \cdot 10^5$	$4.87 \cdot 10^5$	$1.24 \cdot 10^6$
	3	$4.58 \cdot 10^4$	$4.65 \cdot 10^5$	$1.85 \cdot 10^6$	$1.25 \cdot 10^6$	$5.20 \cdot 10^6$	$3.07 \cdot 10^7$	$1.27 \cdot 10^8$	$4.42 \cdot 10^8$	$1.83 \cdot 10^9$
	4	$3.36 \cdot 10^6$	$4.94 \cdot 10^7$	$3.11 \cdot 10^8$	$1.88 \cdot 10^8$	$1.23 \cdot 10^9$	$1.35 \cdot 10^{10}$	$8.73 \cdot 10^{10}$	$4.73 \cdot 10^{11}$	$3.06 \cdot 10^{12}$
	5	$2.80 \cdot 10^8$	$5.96 \cdot 10^9$	$5.78 \cdot 10^{10}$	$3.30 \cdot 10^{10}$	$3.22 \cdot 10^{11}$	$6.76 \cdot 10^{12}$	$6.63 \cdot 10^{13}$	$5.76 \cdot 10^{14}$	$5.66 \cdot 10^{15}$
	6	$2.57 \cdot 10^{10}$	$7.82 \cdot 10^{11}$	$1.17 \cdot 10^{13}$	$6.44 \cdot 10^{12}$	$9.16 \cdot 10^{13}$	$3.71 \cdot 10^{15}$	$5.47 \cdot 10^{16}$	$7.73 \cdot 10^{17}$	$1.14 \cdot 10^{19}$

Table 3.1: Sufficient and necessary samples sizes for (ε, δ)-completeness. A comparison of the specific sample sizes for various values for the attributes of clearance δ, stretch ε, and dimension d. The column "lower bound" reports the value from [55, Theorem 1]. For every combination of attributes we report the value $\mathcal{X}_{\text{curr}}$ from Theorem 3.3, and $\mathcal{X}_{\text{prev}}$ from [55, Theorem 2].

As reported in Theorem 3.3 and [55, Theorem 2], both quantities $\mathcal{X}_{\text{curr}}$ and $\mathcal{X}_{\text{prev}}$ increase exponentially as the dimension increases, and as the clearance δ or the stretch ε decrease. However, the value $\mathcal{X}_{\text{curr}}$ is consistently smaller than $\mathcal{X}_{\text{prev}}$. As the latter value increases it grows more rapidly than $\mathcal{X}_{\text{curr}}$ by at least one order of magnitude. Although in certain cases the value of $\mathcal{X}_{\text{curr}}$ is still quite large for some practical application (particularly when the dimension is larger than 3), we hope that our work would motivate further study into even smaller sampling distributions than our staggered grid, or a more refined version of Theorem 3.3. In this context, we point out that the value $\mathcal{X}_{\text{curr}}$ is larger than the lower bound (where applicable) by one order of magnitude, which suggests that there is room for improvement.

4

Near-optimal tensor roadmaps for MRMP

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

4.1 Basics of multi-robot motion planning

We provide a definition of the multi-robot motion planning (MRMP) problem. We consider the setting of $R \geqslant 2$ identical robots operating in a shared workspace, and denote by $\mathcal{C}_i \subset [0,1]^d$ the configuration space of robot i, $1 \leqslant i \leqslant R$. 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 $1 \leqslant i \leqslant j \leqslant R$.

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., $\mathbb{C} = \mathcal{C}_1 \times \ldots \times \mathcal{C}_R$. That is, a composite configuration $Q = (q_1, ..., q_R) \in \mathbb{C}$ 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 $(q_1, ..., q_R)$ such that (i) $q_i \in \mathcal{C}_i^f$ for every $1 \leqslant i \leqslant R$, and (ii) $q_i \notin I_i^j(q_j)$ for every $1 \leqslant i \neq j \leqslant R$, 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 $1 \leqslant i \leqslant R$, let $\vec{x}^s = (x_1^s, \ldots, x_R^s)$ and $\vec{x}^g = (x_1^g, \ldots, x_R^g)$. 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 $\Sigma = (\sigma_1, ..., \sigma_R)$ of single-robot trajectories $\sigma_i : [0,1] \to \mathcal{C}_i^f$, such that $\Sigma(0) = \vec{x}^s, \Sigma(1) = \vec{x}^g$, and $\Sigma(\tau) \in \mathbb{C}^f$ for all $0 \leqslant \tau \leqslant 1$.

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 applies more generally to cost functions having the following property. Let $\Sigma^1 = (\sigma_1^1, \ldots, \sigma_R^1), \Sigma^2 = (\sigma_1^2, \ldots, \sigma_R^2)$ be multi-robot trajectories, such that $\|\sigma_i^1\| \leqslant (1+\varepsilon)\|\sigma_i^2\|$ for all i, $1 \leqslant i \leqslant R$. Our analysis will work for any cost function c for which $c(\Sigma^1) \leqslant (1+\varepsilon)c(\Sigma^2)$. For example, our analysis can be easily adapted to the case where the cost criterion is the maximum of trajectory lengths, i.e., $\max_{i=1}^{R} \|\sigma_i\|$.

4.2 Tensor roadmaps

We provide a formal definition of the tensor roadmap (TR), which is implicitly explored by sampling-based planners such as `dRRT` [51], `dRRT*` [44], as well as by search-based methods such as `MC-CBS` [32] and `M*` [59].

For every robot i, $1 \leqslant i \leqslant R$, 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 3.2). The TR, denoted by $\widehat{G}(\vec{\mathcal{X}}, \vec{r}) = (\widehat{V}, \widehat{E})$, is the tensor product of $G_1, \ldots, G_R$. 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) $\widehat{V} = \{(v_1, \ldots, v_R) : \forall i, v_i \in V_i\}$, and

(ii) for two vertices $W = (w_1, \ldots, w_R), U = (u_1, \ldots, u_R) \in \widehat{V}$, 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 [2, 14, 43]. 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 $\widehat{G}_{\mathbb{M}(\vec{\mathcal{X}}, \vec{r})}(x^s, x^g)$ 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 $\widehat{G}_{\mathbb{M}(\vec{\mathcal{X}}, \vec{r})}$.

4.3 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, $\mathcal{C}_i^o(\tau) = \mathcal{C}_i^o \bigcup_{j \neq i} I_i^j(\sigma_j(\tau))$ 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 4.1 ($\vec{\delta}$-clearance)**.** *Given a trajectory Σ, we say that Σ has $\vec{\delta}$-clearance for $\vec{\delta} = (\delta_1, \ldots, \delta_R)$ if for each robot i, $1 \leqslant i \leqslant R$ 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 $1 \leqslant i \leqslant R$, $0 \leqslant \tau \leqslant 1$, 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 4.2 (Multi-robot $(\varepsilon, \vec{\delta})$-completeness)**.** *Given R robots, a stretch parameter $\varepsilon > 0$, a vector of R sample sets $\vec{\mathcal{X}} = (\mathcal{X}_1, \ldots, \mathcal{X}_R)$, and a vector of R connection radii $\vec{r} = (r_1, \ldots, r_R)$, 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*

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

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

We are ready to state our main contribution.

Theorem 4.3 (Sufficient conditions for MRMP $(\varepsilon, \vec{\delta})$-completeness)**.** *Let $\varepsilon > 0$ be a stretch factor, let $\vec{\delta}$ be a clearance vector $(\delta_1, \ldots, \delta_R)$, and denote $\omega = \frac{\varepsilon}{2(\varepsilon+2)}$. Define the sampling distributions $\vec{\mathcal{X}} = (\mathcal{X}_1, \ldots, \mathcal{X}_R)$ and radii vector $\vec{r} = (r_1, \ldots, r_R)$, as*

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

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

Proof. Fix $\varepsilon > 0$ and $\vec{\delta} = (\delta_1, \ldots, \delta_R)$ where $\delta_i > 0$ for all $1 \leqslant i \leqslant R$. Let $\vec{x}^s, \vec{x}^g$ be the start and goal composite configurations, respectively. Let $\mathbb{M} = (\mathbb{C}^f, \vec{x}^s, \vec{x}^g)$ be a $\vec{\delta}$-clear MRMP problem. (This implies in particular that the start and goal configurations fulfill the clearance requirements.) Let $\Sigma^* = (\sigma_1^*, \ldots, \sigma_R^*)$ be a $\vec{\delta}$-clear solution which minimizes the expression $\text{cost}(\Sigma)$ over all $\vec{\delta}$-clear solution trajectories Σ.

The overall structure of this proof is as follows. First, we generate for each robot i a `PRM` graph G_i, which is induced by the sample set $\mathcal{X}_i$ and the radius r_i. We show, using Lemma 3.7, that each such G_i induces a collision-free trajectory $\bar{\sigma}_i$ for robot i, which closely follows σ_i^*. Next we show, using additional properties following from Lemma 3.7, that we can specify the arrival time of the individual robots along their corresponding vertices along $\bar{\sigma}_i$, producing the path $\widehat{\sigma}_i$ such that robot-robot collisions are avoided. This implies that the induced tensor graph contains a collision-free composite trajectory $\widehat{\Sigma} = (\widehat{\sigma}_1, \ldots, \widehat{\sigma}_R)$, whose cost is at most $(1+\varepsilon)\text{cost}(\Sigma^*)$. We fill in the details of the proof below.

For every $1 \leqslant i \leqslant R$, define $\mathcal{X}_i = \mathcal{X}_{\beta_i, \delta_i}$, where $\beta_i = \omega\delta_i$. 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 product of $G_1, \ldots, G_R$. Also, define $\rho_i = \frac{\delta_i}{\varepsilon+2}$ and let $T_i := T^{\rho_i}_{\sigma_i^*} = (\tau_0^i, \ldots, \tau_{\ell_i}^i)$, as in Definition 3.6. First, we show that the conditions for Lemma 3.7 hold for each robot i, i.e., $\beta_i^2 + (\rho_i/2)^2 \leqslant \delta_i^2$. Indeed,

$$\begin{aligned}\beta_i^2 + \left(\frac{\rho_i}{2}\right)^2 &= \omega^2\delta_i^2 + \frac{\delta_i^2}{4(\varepsilon+2)^2} \\ &= \left(\frac{\varepsilon^2+1}{(\varepsilon^2+1)+(3\varepsilon^2+8\varepsilon+15)}\right)\delta_i^2 \\ &\leqslant \delta_i^2.\end{aligned}$$

Thus, we can apply Lemma 3.7 with the time sequence T_i for each robot i individually to show that there exists a point set $Z_i := Z^{\rho_i}_{\sigma_i^*} = (z_0^i, \ldots, z_{\ell_i}^i) \subset \mathcal{X}_i \cup \{x_i^s, x_i^g\}$, such that Properties (i)-(vi) of Lemma 3.7 hold. In particular, we have that for all $1 \leqslant i \leqslant R$,

(i') $\|z_k^i - \sigma_i^*(\tau_k^i)\| \leqslant \beta_i$, for all $0 \leqslant k \leqslant \ell_i$;

(ii') $\|z - \sigma_i^*(\tau_{k-1}^i)\| \leqslant \beta_i + \rho_i$, for all $1 \leqslant k \leqslant \ell_i$, $z \in \mathrm{CH}\left(z_{k-1}^i, z_k^i\right)$;

(iii') $\|z - \sigma_i^*(\tau_k^i)\| \leqslant \beta_i + \rho_i$, for all $1 \leqslant k \leqslant \ell_i$, $z \in \mathrm{CH}\left(z_{k-1}^i, z_k^i\right)$;

(iv') $\|z_k^i - z_{k-1}^i\| \leqslant 2\beta_i + \rho_i$, for all $1 \leqslant k \leqslant \ell_i$;

(v') $\mathrm{CH}\left(z_{k-1}^i, z_k^i\right)$ is collision free, for all $1 \leqslant k \leqslant \ell_i$. That is $\mathrm{CH}\left(z_{k-1}^i, z_k^i\right) \subset \mathcal{C}_i^f$, for all $1 \leqslant k \leqslant \ell_i$;

(vi') $\sum_{1\leqslant k\leqslant \ell_i} \|z_k^i - z_{k-1}^i\| \leqslant \left(1 + \frac{2\beta_i}{\rho_i}\right)\|\sigma_i^*\|$.

By setting $r_i = 2\beta_i + \rho_i = \frac{\delta_i(\varepsilon+1)}{(\varepsilon+2)}$, and using Property (iv'), it follows that the edge connecting z_j^i and z_{j-1}^i is considered in the construction of G_i, and from (v'), it is in G_i. Next, due to the fact that $2\beta_i/\rho_i = \varepsilon$, and Property (vi'), it holds that $\|\bar{\sigma}_i\| \leqslant (1+\varepsilon)\|\sigma_i^*\|$, for $\bar{\sigma}_i$ denoting the trajectory induced by $Z^{\rho_i}_{\sigma_i^*}$. Thus, each robot has a collision free path, which is a $(1+\varepsilon)$-approximation for its δ_i-clear path σ_i^*. Finally, observe that $\mathrm{cost}(\bar{\Sigma}) \leqslant (1+\varepsilon)\mathrm{cost}(\Sigma^*)$, where $\bar{\Sigma} = (\bar{\sigma}_1, \ldots, \bar{\sigma}_R)$. Note that this last step also applies to a cost function that returns the maximum length over the R single-robot trajectories.

As robots may collide with one another along the paths $\bar{\Sigma}$, we leverage the above properties (i')-(iv') to show that $\widehat{G}$, the TR where edges describe coordinated collision-free motion where the robots avoid obstacles as well as each other, contains a high-quality composite trajectory which avoids robot-obstacle *and robot-robot* collisions. To do so, we show that we can adjust the positions of the robots along the trajectories $\bar{\sigma}_1, \ldots, \bar{\sigma}_R$, to induce a collision-free trajectory over $\widehat{G}$, which we denote by $\widehat{\Sigma} = (\widehat{\sigma}_1, \ldots, \widehat{\sigma}_R)$.

First, 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 $1 \leqslant i \leqslant R$, 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,

$$\mathcal{L}^o := \{(i_1, \tau_{i_1}, z_{i_1}), (i_2, \tau_{i_2}, z_{i_2}), \ldots, (i_\ell, \tau_{i_\ell}, z_{i_\ell})\},$$

where $\ell = \sum_{i=1}^{R} \ell_i$, $(i_j, \tau_{i_j}, z_{i_j}) \in \mathcal{L}$ for every $1 \leqslant j \leqslant \ell$, and $\tau_{i_j} \leqslant \tau_{i_{j+1}}$ for every $1 \leqslant j \leqslant \ell - 1$. For simplicity, when several robots have the same timestamp, we order the corresponding triplets in increasing order of the robot indices.

Next, we describe an iterative scheme that uses $\mathcal{L}^o$ for generating a sequence of composite vertices $V_0, V_1, \ldots, V_\ell \in \widehat{V}$, such that $(V_j, V_{j+1}) \in \widehat{E}$. First, define $V_0 = \vec{x}^s$. Next, given that $V_j = (v_{j_1}, \ldots, v_{j_R})$ has already been defined for some $1 \leqslant j \leqslant \ell - 1$, set $V_{j+1} = (v_{(j+1)_1}, \ldots, v_{(j+1)_R})$, where $v_{(j+1)_{i_{j+1}}} := z_{i_{j+1}}$, and $v_{(j+1)_{i'}} = v_{j_{i'}}$ 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 $\mathcal{L}^o$.

To illustrate this, we provide the following example.

Example 1. *We provide an example of $\mathcal{L}, \mathcal{L}^o$ and $V_0, \ldots, V_\ell$ for a two-robot setting with $T_1 = (0, 0.3, 0.5, 0.7, 1)$, $Z_1 = (z_0^1, z_1^1, z_2^1, z_3^1, z_4^1)$, $T_2 = (0, 0.2, 0.4, 0.7, 1)$, $Z_2 = (z_0^2, z_1^2, z_2^2, z_3^2, z_4^2)$. Note that z_0^i is robot i's starting point, and z_4^i is its goal. In particular,*

$$\begin{aligned}\mathcal{L} = \{&(1, 0.3, z_1^1), (1, 0.5, z_2^1), (1, 0.7, z_3^1), (1, 1, z_4^1),\\ &(2, 0.2, z_1^2), (2, 0.4, z_2^2), (2, 0.7, z_3^2), (2, 1, z_4^2)\}.\end{aligned}$$

$$\begin{aligned}\mathcal{L}^o = \{&(2, 0.2, z_1^2), (1, 0.3, z_1^1), (2, 0.4, z_2^2), (1, 0.5, z_2^1),\\ &(1, 0.7, z_3^1), (2, 0.7, z_3^2), (1, 1, z_4^1), (2, 1, z_4^2)\},\end{aligned}$$

yielding the vertices $V_0 = (z_0^1, z_0^2), V_1 = (z_0^1, z_1^2), V_2 = (z_1^1, z_1^2), V_3 = (z_1^1, z_2^2), V_4 = (z_2^1, z_2^2), V_5 = (z_3^1, z_2^2), V_6 = (z_3^1, z_3^2), V_7 = (z_4^1, z_3^2), V_8 = (z_4^1, z_4^2)$.

To complete the proof, we first note that $(V_j, V_{j+1}) \in \widehat{E}$ for every $1 \leqslant j \leqslant \ell - 1$. This follows from the values of the connection radii $\vec{r}$ we assigned, and Properties (iv') and (v'), which ensure that $\{v_j^i, v_{j+1}^i\} \in E_i$ for every robot $1 \leqslant i \leqslant R$, where $(v_1, \ldots, v_R) := V_j, (v'_1, \ldots, v'_R) := V_{j+1}$.

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 jth 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 $v_k = v'_k$.

Notice that given two stationary robots $k_1, k_2 \neq i_j$, and assuming that they did not collide for the previous edge (V_{j-1}, V_j), they do not collide with each other for (V_j, V_{j+1}) as well.

It remains to show 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 $V_0, \ldots, V_\ell$, it must be that robot k resides in

$v_k = z^k_{j'} \in Z_k$ such that $\tau^k_{j'} \leqslant \tau_{i_j} \leqslant \tau^k_{j'+1}$, for some $0 \leqslant j' \leqslant \ell_k - 1$. Next, recall that $\sigma^*_k, \sigma^*_{i_j}$ are δ_k and δ_{i_j}-clear, respectively. Thus, it suffices to prove that $\|\sigma^*_{i_j}(\tau_{i_j}) - p\| \leqslant \delta_{i_j}/2$ for every $p \in \mathrm{CH}\left(v_{i_j}, v'_{i_j}\right)$ and $\|\sigma^*_k(\tau_{i_j}) - z^k_{j'}\| \leqslant \delta_k/2$.

Indeed, Property (iii') implies that for $p \in \mathrm{CH}\left(v_{i_j}, v'_{i_j}\right)$ we have that $\left\|p - \sigma^*_{i_j}(\tau_{i_j})\right\| \leqslant \beta_{i_j} + \rho_{i_j} \leqslant \delta_{i_j}/2$. We also have that

$$\begin{aligned}\left\|z^k_{j'} - \sigma^*_k(\tau_{i_j})\right\| &\leqslant \|z^k_{j'} - \sigma^*_k(\tau^k_{j'})\| + \|\sigma^*_k(\tau^k_{j'}) - \sigma^*_k(\tau_{i_j})\| \\ &\leqslant \beta_k + \rho_k \\ &\leqslant \delta_k/2.\end{aligned}$$

Since at time τ_{i_j} both robots i_j and k obey the clearance rules, they are at distance $\delta_{\max} = \max(\delta_{i_j}, \delta_k)$ from one another. As we have shown that during the specified motion, robot i_j has not drifted more than $\delta_{\max}/2$ from its position at time τ_{i_j}, and the stationary placement of robot k during this motion is at most $\delta_{\max}/2$ from its placement at time τ_{i_j}, we are guaranteed that they do not collide with one another during this motion, which concludes the proof. □

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.

4.4 Discussion

Theorem 4.3 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 $\vec{\delta}$-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 $1 \leqslant i \leqslant R$, 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 4.3 within each `PRM` roadmap for specific parameters. In particular, we report in Table 4.1 the value $|\mathcal{X}_{\omega\delta_i,\delta_i}|$, for varying values of the stretch parameter ε, dimension d,

and clearance vector $(\delta_1, \ldots, \delta_R)$, 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 5. This suggests that smaller sample sets are sufficient for $(\varepsilon, \vec{\delta})$-completeness, which we aim to explore in future research (see Section 7).

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	$1.89 \cdot 10^7$	$1.5 \cdot 10^8$	$1.18 \cdot 10^{10}$	$2.46 \cdot 10^{11}$	$7.82 \cdot 10^{12}$

Table 4.1: Multi-robot sample complexity example for $\delta = 0.1$ with varying values for the dimension of each robot, d, and the required maximal stretch factor, ε.

5 Experimental results

We provide experimental results to support our theoretical findings, focusing on the case of multiple disc robots 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 $\widehat{G}$ in Theorem 4.3, 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 7).

Concerning the design of the experiments, we were faced with two challenges. First, current MRMP algorithms [32, 44, 59] (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.

5.1 Implementation

In this section we describe our software implementation with which the experimental results were obtained. We describe what has been implemented, how it was implemented and where to find it.

We implemented the code for moving discs among polygonal obstacles in the plane. In

our implementation all the robots having the same value δ for the clearance parameter. The code can be easily extended to support different δ_i values for the different robots, or work for other identical two-dimensional shapes. Other extensions are possible but may be less easy to implement using our code base.

Our code was added as a planner to **DiscoPygal**, a software environment for developing and testing multi-robot motion planning algorithm in Python, developed at Tel Aviv University [13]. The reading of input files, testing the output for validity, and visualization of the scenes and the robot motion, are implemented by the DiscoPygal framework. The collision detection was implemented using python bindings to the arrangement package [11] of the CGAL library [54], which are part of DiscoPygal. All the code beside that was implemented by us for the research reported in this thesis.

The planner gets two inputs. The first input is a scene, which comprises: start and goal positions for each robot, the radius of the robots (same radius for all the robots in our case), and a set of polygonal obstacles. The second input is the configuration, namely the values for ε and δ, the sample method to use (staggered grid or uniform random), and a few other properties.

The planner can create two types of outputs. First it can create an input file for `MC-CBS` [32]. The second output is a solution path, if one was found. We note that both outputs can be generated at once.

In any case we first calculate the set of sample points (i.e., milestones), and the connection radius, and then create a `PRM` graph for the robots (same one for all robots, although this is easily extendable).

There is option to run our implementation of the A^* algorithm [16] on the tensor roadmap. During the running of A^*, we create the parts of the tensor roadmap, which are discovered in our search. At each iteration we look at the best candidate we found so far, and only add vertices that represent motion by a single robot—this is the easiest way to create new valid nodes in the tensor roadmap, which still satisfy the optimality guarantees. It is possible to move more than one robot at a time, and the code can easily be extended to support it.

Another option is to compute conflicts for `MC-CBS`; these include vertices conflicts, edge and vertex conflicts, and edge to edge conflicts. Once all conflicts were calculated the output is printed to a YAML file. This file is then the input file for our modified version of the implementation of `MC-CBS` [32].

The code described here is publicly available on Git-hub, https://github.com/drorDayan/eps-net.

A more user friendly version of the code containing only the A^* option is also included in DiscoPygal, which can be found in https://bitbucket.org/taucgl/python_rmp_framework/src/master. This is the recommended version to use. Notice that DiscoPygal is currently undergoing a major upgrade, and a new version is anticipated soon.

5.2 Scenarios

The scenarios are illustrated in Figure 5.1. The first scenario, [Top-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* [26]. The [Top-right] 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, [Bottom-left], 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, [Bottom-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.

5.3 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., $\vec{\delta} := \{\delta, \ldots, \delta\}$). For all the scenarios, we set the value δ to be equal to the *static clearance* μ (see caption of Figure 5.1). 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.

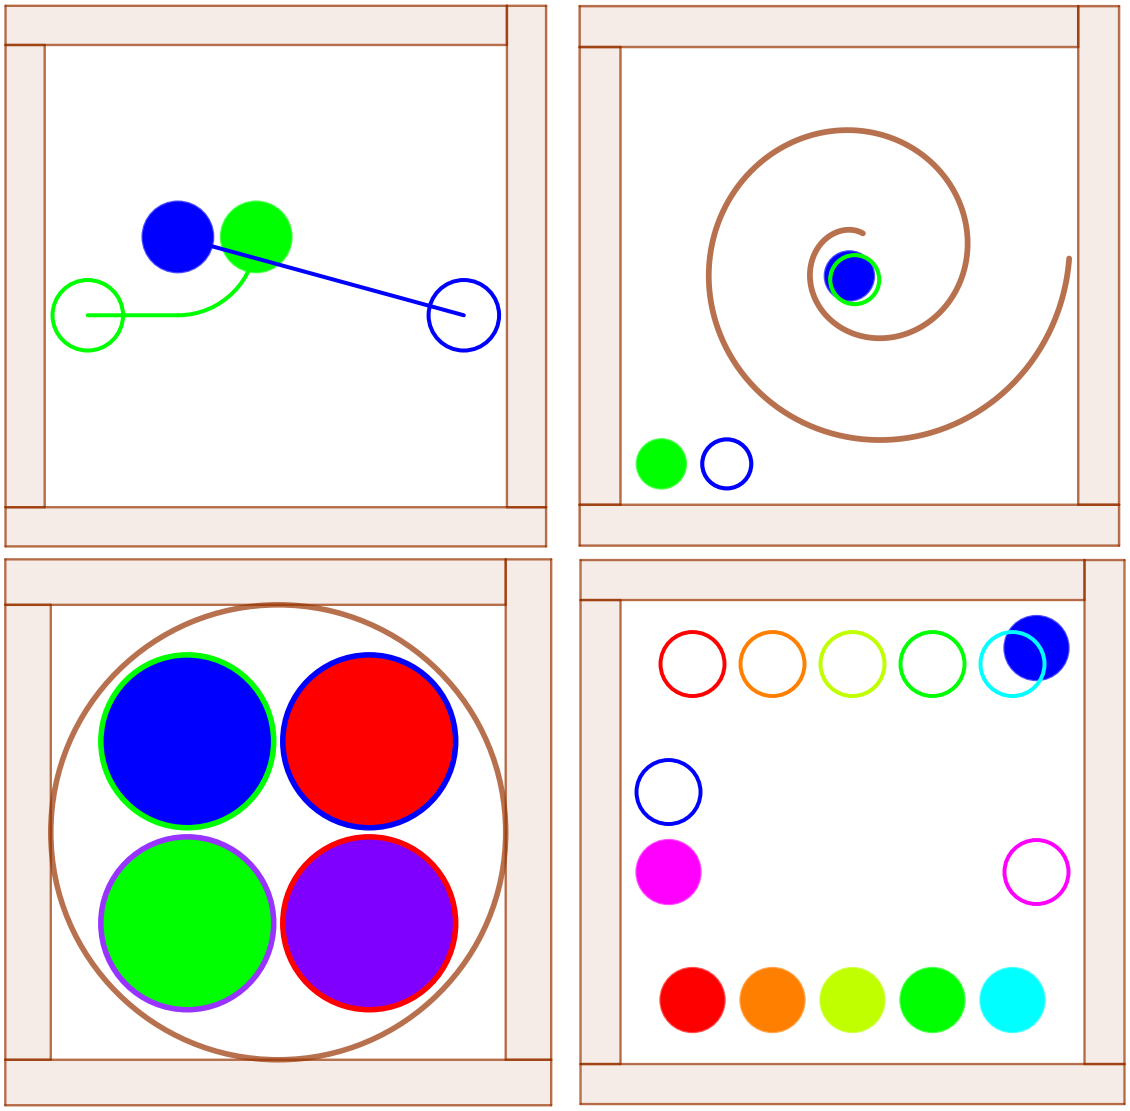

Figure 5.1: Test scenarios for multiple disc robots, where a circle and a disc of the same color represent the start and goal positions, respectively, of a robot. [Top-left] A 2-robot obstacle-free scenario. The optimal trajectories, which were derived in [26], are drawn for each robot. The robots' radius 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. [Top-right] A 2-robot scenario with a spiral obstacle with robots' radius 0.06 and static clearance $\mu = 0.04$. [Bottom-left] Four robots tightly placed within a circular barrier, using robot radius 0.19 and static clearance $\mu = 0.02$. [Bottom-right] A 7-robot scenario with robots' radius 0.08 and the static clearance $\mu = 0.04$.

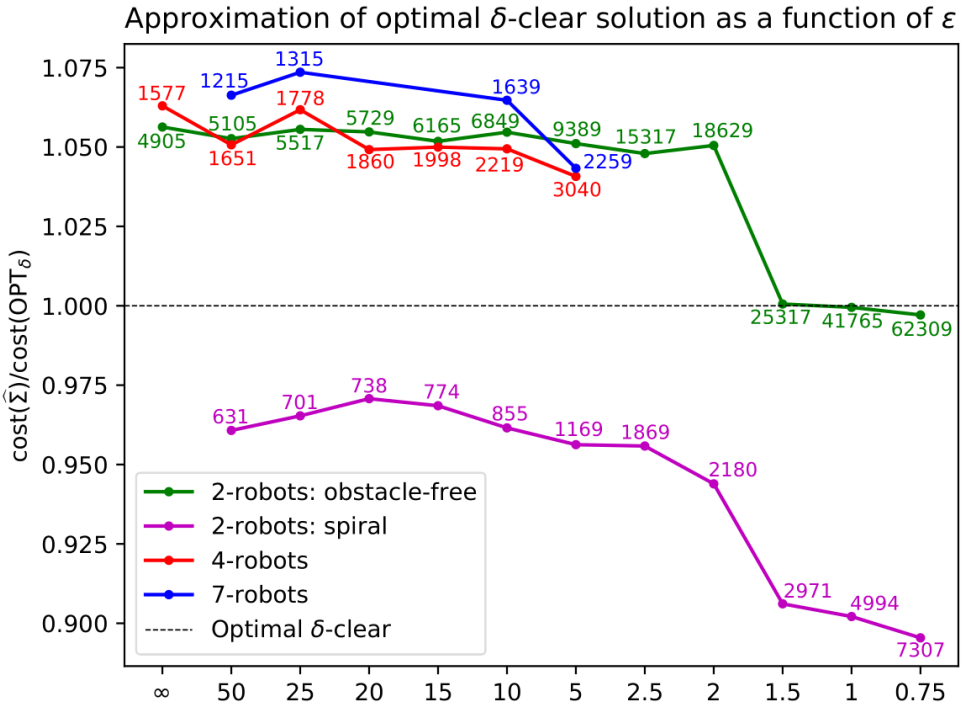

Figure 5.2: We report for each of the four scenarios its approximation factors, which are represented by the ratio between the cost of the solution obtained from the tensor roadmap and OPT_δ. Notice that the tensor-roadmap solution is not necessarily δ-clear and thus can be of lower cost than OPT_δ, 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.

The results are reported in Figure 5.2, 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 Theorem 4.3). The reported approximation factor represents the ratio between the cost of the best solution obtained from the tensor roadmap, denoted by $\widehat{\Sigma}$, and the optimal δ-clear solution whose cost is denoted by OPT_δ. For the first scenario we obtain the value OPT_δ using [26], as we reported earlier. For the [Top-right] 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 [Bottom-left] scenario, OPT_δ is equal to the perimeter of the circle going though the robot centers at the initial positions (as each robot traverses a quarter of the circle). For the [Bottom-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 param-

eter to $\varepsilon = 50$ we obtain an approximation factor of at most 1.075. Moreover, for $\varepsilon \leqslant 1.5$ 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 $\mathcal{X}_{\omega\delta,\delta}$. In some cases using a smaller number of samples may yield better solutions, as in the 2-robot spiral scenario for stretch factors $\varepsilon = 50$ and $\varepsilon = 20$. 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.

5.4 Comparing the staggered grid with random sampling

The staggered grid offers good theoretical guarantees on the necessary sample size for obtaining near-optimal solution in the single robot case and consequently in the multi-robot case as well. An important question is whether the staggered grid also provides an improvement over the somewhat standard random uniform sampling in practice. In this section we demonstrate that the staggered grid provides a practical improvement over random uniform sampling especially for small sample sets.

Next, we show how the staggered grid outperforms random sampling in both of our test scenes with 2 robots, those depicted in the [Top-left] and [Top-right] sub-figures of Figure 5.1. The testing scheme we use is the following: for varying values of ε, we calculate the number of points in the staggered grid, and the connection radius prescribed by Theorem 4.3. We then randomly choose that same number of points uniformly and use the same connection radius for solving the scene. This process is repeated 10 times to account for randomness.

We first look at the success rate. Using the staggered grid we are guaranteed to have a perfect success rate as long as there exists a path with δ-clearance, for a given δ. For a random set of points this is no longer the case—see the success rate reported in Figure 5.3. Note that as ε decreases (and the number of samples increases) the success rate improves. We also observe that the scene with obstacles has a very low success rate for small values of ε.

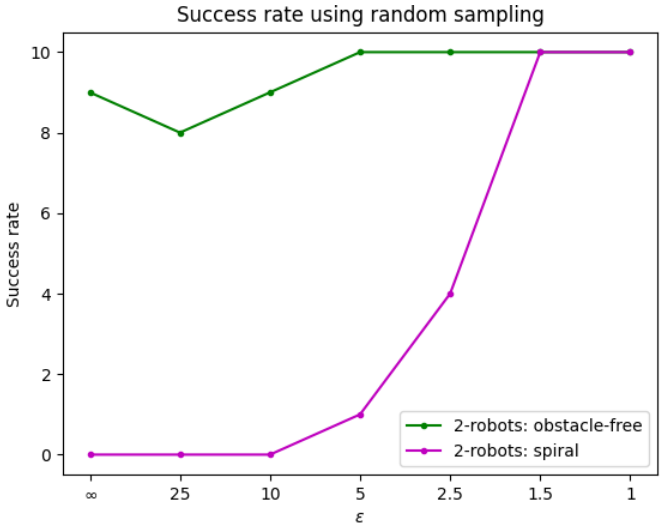

Figure 5.3: The success rate in finding a valid solution for the MRMP problem, using a sample set consisting of random points (uniformly distributed in $[\delta, 1-\delta]^2$). The success rate is measured for 10 runs.

Next we consider the cost of the resulting trajectory. We compare the average cost of the path generated using the randomly sampled points with the cost of the path generated using the staggered grid. As can be observed in Figure 5.4, for small sets of points (large ε values) the staggered grid outperforms an average uniformly chosen random set of sample points. This advantage decreases for smaller ε values. These results are not very surprising, as we know that `PRM` can be near-optimal for infinitely large sets of random points; using a suitable radius, this is proved for `PRM*` in [23].

To summarize, our results emphasize the strength of the staggered grid for large values of ε, when compared with the standard random sampling both in terms of success rate and solution quality.

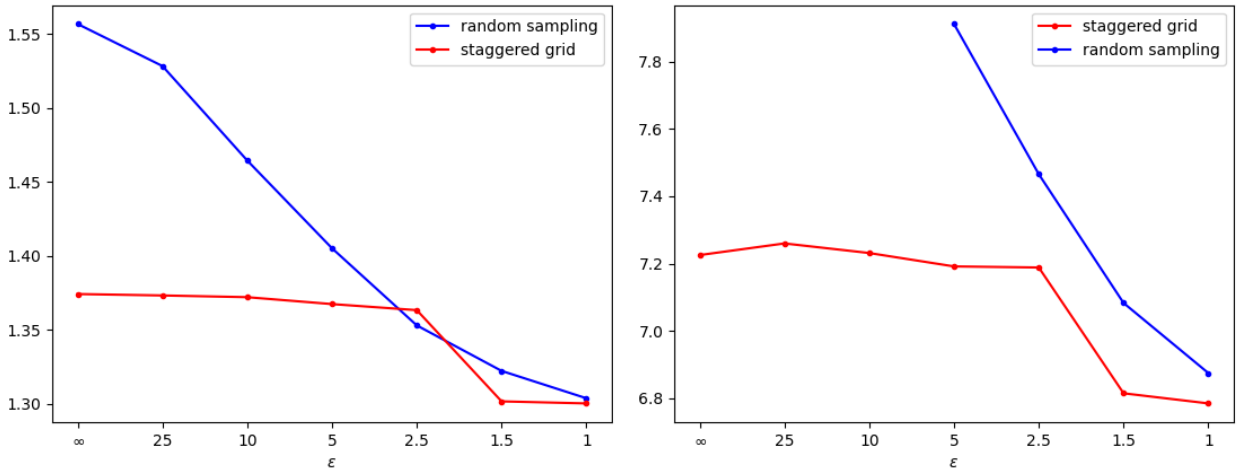

(a) obstacle-free

(b) spiral

Figure 5.4: The average cost of successful runs for the MRMP problem, for the two scenarios of two robots each. We show the average cost using randomly sampled points (blue line) and the cost using the staggered grid (red line).

6

Staggered grid: Comparative size analysis

In this section we present further comparative analysis of the size of the staggered grid. We do that for two reasons. First, for completeness, and secondly because it raises some interesting questions about the lower bound for both the cube cover and the motion-planning problems. Our benchmarks ($\mathcal{Y}_{\text{LB}}$, $\mathcal{Y}_{\text{prev}}$, defined below) will be the previous results from [55, Theorem 3].

Let $\beta, \gamma > 0$ be the parameters as described for Lemma 3.5, and $d \geqslant 2$ the dimension of the γ-cube and the β-spheres. We denote by $\mathcal{Y}_{\text{LB}}$ the lower bound for the cube cover problem. Namely,

$$\mathcal{Y}_{\text{LB}} \approx \sqrt{\pi d}\left(\sqrt{\frac{d}{2\pi e}} \cdot \frac{1-2\gamma}{\beta}\right)^d .$$

We denote by $\mathcal{Y}_{\text{prev}}$ the previous upper bound for the cube cover problem. Namely,

$$\mathcal{Y}_{\text{prev}} \approx \sqrt{\pi d}\left(\sqrt{\frac{2d}{\pi e}} \cdot \frac{1-2\gamma+\beta}{\beta}\right)^d .$$

Lastly, we denote by $\mathcal{Y}_{\text{curr}}$ the number of samples in the staggered grid as in Definition 3.2. Namely,

$$\mathcal{Y}_{\text{curr}} = \left(\left\lceil \frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta} \right\rceil\right)^d + \left(\left\lceil \frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta} \right\rceil + 1\right)^d .$$

To analyze our results we use both the asymptotic regime, where β and γ tend to 0, as well as provide below a table with selected exact values.

We first consider the ratio $\frac{\mathcal{Y}_{\text{curr}}}{\mathcal{Y}_{\text{LB}}}$, where the values β and γ tend to zero.

$$\begin{aligned}\frac{\mathcal{Y}_{\text{curr}}}{\mathcal{Y}_{\text{LB}}} &= \frac{\left(\left\lceil \frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta} \right\rceil\right)^d + \left(\left\lceil \frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta} \right\rceil + \mathbf{1}\right)^d}{\sqrt{\pi d}\left(\sqrt{\frac{d}{2\pi e}} \cdot \frac{1-2\gamma}{\beta}\right)^d} \\ &\overset{(a)}{\approx} \frac{2\left(\frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta}\right)^d}{\sqrt{\pi d}\left(\sqrt{\frac{d}{2\pi e}} \cdot \frac{1-2\gamma}{\beta}\right)^d} \\ &= \frac{2}{\sqrt{\pi d}}\left(\sqrt{\frac{\pi e}{4}}\right)^d \approx \frac{2}{\sqrt{\pi d}} 1.4611^d \,,\end{aligned} \tag{6.1}$$

where the approximation in (a) is due to ignoring the highlighted unit value 1 in the numerator and the roundups (the influence of both values tends to 0 as β or γ tend to 0). This implies that our result yields cover sets that are bigger than the lower bound by an exponential factor in d, but this factor is only 1.4611 compared to an exponential factor of more than 2 in the previous work.

Recall the similar comparison for the motion-planning problem. Namely, the ratio $\frac{\mathcal{X}_{\text{curr}}}{\mathcal{X}_{\text{LB}}}$ between our bound on the number of samples required for motion planning over the lower bound for the motion planning problem, as derived in Eq 3.4; we repeat it here for ease of comparison.

$$2\sqrt{2} \cdot 1.4611^d \leqslant \frac{\mathcal{X}_{\text{curr}}}{\mathcal{X}_{\text{LB}}} \leqslant 4\sqrt{\frac{2}{e}} \cdot 1.4611^d.$$

Note that the base of the exponent is exactly the same, while there is some difference in the constants and the $1/\sqrt{d}$ in the cover ratio. The reason these lower bounds are not in an exact same relation to our bound is that they are both derived from a volume based proof, but these proofs are not identical [55].

Coxeter et al. [8] showed lower bounds on the density of a cover of $\mathbb{R}^d$. While out of scope for this thesis, it should not be too difficult to use their lower bounds to derive better lower bounds for the cube cover problem. More interestingly, these alternative bounds might also be used for deriving improved lower bounds for the motion-planning problem.

Next, we consider the ratio $\frac{\mathcal{Y}_{\text{prev}}}{\mathcal{Y}_{\text{curr}}}$, again for the case where the values β and γ tend to zero.

$$\begin{aligned}\frac{\mathcal{Y}_{\text{prev}}}{\mathcal{Y}_{\text{curr}}} &= \frac{\sqrt{\pi d}\left(\sqrt{\frac{2d}{\pi e}} \cdot \frac{1-2\gamma+\beta}{\beta}\right)^d}{\left(\left\lceil \frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta} \right\rceil\right)^d + \left(\left\lceil \frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta} \right\rceil + 1\right)^d} \\ &\approx \frac{\sqrt{\pi d}}{2} 1.3687^d \,;\end{aligned} \tag{6.2}$$

we omit the calculations as those are similar to Eq 3.2. Note that the cube cover improvement ratio is exactly the same as the motion-planning improvement ratio (Eq 3.2).

Lastly, in Table 6.1 we compare our results with the previous lower and upper bounds for specific values. It shows that for most feasible cases ($d \leqslant 10$) the staggered grid is closer to the lower bound than to the previous upper bound.

		$\beta = 0.1$			$\beta = 0.05$			$\beta = 0.01$		
γ	d	$\mathcal{Y}_{\text{LB}}$	$\mathcal{Y}_{\text{curr}}$	$\mathcal{Y}_{\text{prev}}$	$\mathcal{Y}_{\text{LB}}$	$\mathcal{Y}_{\text{curr}}$	$\mathcal{Y}_{\text{prev}}$	$\mathcal{Y}_{\text{LB}}$	$\mathcal{Y}_{\text{curr}}$	$\mathcal{Y}_{\text{prev}}$
0.05	2	26	61	128	104	181	460	2579	4141	$1.05 \cdot 10^4$
	3	175	559	1910	1393	3925	$1.31 \cdot 10^4$	$1.74 \cdot 10^5$	$3.61 \cdot 10^5$	$1.44 \cdot 10^6$
	4	1330	6497	$3.24 \cdot 10^4$	$2.13 \cdot 10^4$	$6.70 \cdot 10^4$	$4.23 \cdot 10^5$	$1.33 \cdot 10^7$	$3.46 \cdot 10^7$	$2.22 \cdot 10^8$
	5	$1.12 \cdot 10^4$	$9.18 \cdot 10^4$	$6.08 \cdot 10^5$	$3.59 \cdot 10^5$	$1.81 \cdot 10^6$	$1.51 \cdot 10^7$	$1.12 \cdot 10^9$	$4.01 \cdot 10^9$	$3.79 \cdot 10^{10}$
	6	$1.03 \cdot 10^5$	$7.94 \cdot 10^5$	$1.24 \cdot 10^7$	$6.58 \cdot 10^6$	$4.09 \cdot 10^7$	$5.83 \cdot 10^8$	$1.03 \cdot 10^{11}$	$4.68 \cdot 10^{11}$	$7.03 \cdot 10^{12}$
	7	$1.01 \cdot 10^6$	$1.48 \cdot 10^7$	$2.71 \cdot 10^8$	$1.30 \cdot 10^8$	$1.02 \cdot 10^9$	$2.42 \cdot 10^{10}$	$1.01 \cdot 10^{13}$	$6.69 \cdot 10^{13}$	$1.40 \cdot 10^{15}$
	8	$1.06 \cdot 10^7$	$1.43 \cdot 10^8$	$6.31 \cdot 10^9$	$2.72 \cdot 10^9$	$2.80 \cdot 10^{10}$	$1.07 \cdot 10^{12}$	$1.06 \cdot 10^{15}$	$9.01 \cdot 10^{15}$	$2.97 \cdot 10^{17}$
	9	$1.17 \cdot 10^8$	$3.36 \cdot 10^9$	$1.55 \cdot 10^{11}$	$6.01 \cdot 10^{10}$	$1.31 \cdot 10^{12}$	$5.01 \cdot 10^{13}$	$1.17 \cdot 10^{17}$	$1.45 \cdot 10^{18}$	$6.64 \cdot 10^{19}$
	10	$1.37 \cdot 10^9$	$8.79 \cdot 10^{10}$	$4.02 \cdot 10^{12}$	$1.40 \cdot 10^{12}$	$4.32 \cdot 10^{13}$	$2.46 \cdot 10^{15}$	$1.37 \cdot 10^{19}$	$2.32 \cdot 10^{20}$	$1.56 \cdot 10^{22}$
0.01	2	31	61	149	123	221	541	3058	4901	$1.25 \cdot 10^4$
	3	225	855	2406	1798	4941	$1.67 \cdot 10^4$	$2.25 \cdot 10^5$	$4.65 \cdot 10^5$	$1.85 \cdot 10^6$
	4	1870	6497	$4.41 \cdot 10^4$	$2.99 \cdot 10^4$	$8.90 \cdot 10^4$	$5.84 \cdot 10^5$	$1.87 \cdot 10^7$	$4.94 \cdot 10^7$	$3.11 \cdot 10^8$
	5	$1.72 \cdot 10^4$	$9.18 \cdot 10^4$	$8.93 \cdot 10^5$	$5.50 \cdot 10^5$	$2.47 \cdot 10^6$	$2.26 \cdot 10^7$	$1.72 \cdot 10^9$	$5.96 \cdot 10^9$	$5.78 \cdot 10^{10}$
	6	$1.71 \cdot 10^5$	$1.53 \cdot 10^6$	$1.97 \cdot 10^7$	$1.10 \cdot 10^7$	$5.81 \cdot 10^7$	$9.46 \cdot 10^8$	$1.71 \cdot 10^{11}$	$7.82 \cdot 10^{11}$	$1.17 \cdot 10^{13}$
	7	$1.84 \cdot 10^6$	$2.95 \cdot 10^7$	$4.64 \cdot 10^8$	$2.35 \cdot 10^8$	$2.17 \cdot 10^9$	$4.26 \cdot 10^{10}$	$1.84 \cdot 10^{13}$	$1.16 \cdot 10^{14}$	$2.53 \cdot 10^{15}$
	8	$2.10 \cdot 10^7$	$3.14 \cdot 10^8$	$1.17 \cdot 10^{10}$	$5.37 \cdot 10^9$	$6.34 \cdot 10^{10}$	$2.05 \cdot 10^{12}$	$2.10 \cdot 10^{15}$	$1.77 \cdot 10^{16}$	$5.82 \cdot 10^{17}$
	9	$2.53 \cdot 10^8$	$7.52 \cdot 10^9$	$3.10 \cdot 10^{11}$	$1.29 \cdot 10^{11}$	$2.00 \cdot 10^{12}$	$1.04 \cdot 10^{14}$	$2.53 \cdot 10^{17}$	$2.97 \cdot 10^{18}$	$1.42 \cdot 10^{20}$
	10	$3.20 \cdot 10^9$	$8.79 \cdot 10^{10}$	$8.67 \cdot 10^{12}$	$3.28 \cdot 10^{12}$	$6.80 \cdot 10^{13}$	$5.53 \cdot 10^{15}$	$3.20 \cdot 10^{19}$	$5.43 \cdot 10^{20}$	$3.63 \cdot 10^{22}$
0	2	32	61	155	128	221	562	3184	5101	$1.30 \cdot 10^4$
	3	239	855	2543	1910	4941	$1.77 \cdot 10^4$	$2.39 \cdot 10^5$	$4.88 \cdot 10^5$	$1.97 \cdot 10^6$
	4	2027	$1.07 \cdot 10^4$	$4.75 \cdot 10^4$	$3.24 \cdot 10^4$	$1.16 \cdot 10^5$	$6.31 \cdot 10^5$	$2.03 \cdot 10^7$	$5.23 \cdot 10^7$	$3.37 \cdot 10^8$
	5	$1.90 \cdot 10^4$	$9.18 \cdot 10^4$	$9.79 \cdot 10^5$	$6.08 \cdot 10^5$	$2.47 \cdot 10^6$	$2.48 \cdot 10^7$	$1.90 \cdot 10^9$	$6.76 \cdot 10^9$	$6.39 \cdot 10^{10}$
	6	$1.94 \cdot 10^5$	$1.53 \cdot 10^6$	$2.19 \cdot 10^7$	$1.24 \cdot 10^7$	$8.11 \cdot 10^7$	$1.06 \cdot 10^9$	$1.94 \cdot 10^{11}$	$8.98 \cdot 10^{11}$	$1.31 \cdot 10^{13}$
	7	$2.12 \cdot 10^6$	$2.95 \cdot 10^7$	$5.28 \cdot 10^8$	$2.71 \cdot 10^8$	$2.17 \cdot 10^9$	$4.88 \cdot 10^{10}$	$2.12 \cdot 10^{13}$	$1.35 \cdot 10^{14}$	$2.90 \cdot 10^{15}$
	8	$2.46 \cdot 10^7$	$3.14 \cdot 10^8$	$1.35 \cdot 10^{10}$	$6.31 \cdot 10^9$	$6.34 \cdot 10^{10}$	$2.39 \cdot 10^{12}$	$2.46 \cdot 10^{15}$	$2.08 \cdot 10^{16}$	$6.83 \cdot 10^{17}$
	9	$3.03 \cdot 10^8$	$7.52 \cdot 10^9$	$3.66 \cdot 10^{11}$	$1.55 \cdot 10^{11}$	$3.01 \cdot 10^{12}$	$1.23 \cdot 10^{14}$	$3.03 \cdot 10^{17}$	$3.84 \cdot 10^{18}$	$1.70 \cdot 10^{20}$
	10	$3.92 \cdot 10^9$	$2.00 \cdot 10^{11}$	$1.04 \cdot 10^{13}$	$4.02 \cdot 10^{12}$	$1.05 \cdot 10^{14}$	$6.70 \cdot 10^{15}$	$3.92 \cdot 10^{19}$	$6.50 \cdot 10^{20}$	$4.44 \cdot 10^{22}$

Table 6.1: A comparison of the β-cover size for the d-dimensional $(1-2\gamma)$-cube with various values for the attributes of γ, ball radius β, and dimension d. For every combination of attributes we report the following values, $\mathcal{Y}_{\text{LB}}$ is the lower bound for the cube cover problem, $\mathcal{Y}_{\text{prev}}$ is the previous result, both are from [55, Theorem 3]. $\mathcal{Y}_{\text{curr}}$ is the number of points required to cover the cube using the staggered grid. Note that $\gamma = 0$ is the cover of the unit cube, and that for the motion-planning problem $\beta \leqslant \gamma$.

7

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 planing 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 4.3).

A

Better covering of a square

We present a covering for the two-dimensional case suggested by Coxeter et al. [8], which is more economical than our staggered grid. We will compare the *density* of their covering with the density of our staggered grid 3.2 asymptotically, namely when the radius of the covering discs tends to zero.

First we define the density of a covering: It is the ratio between the volume of all objects in the cover and the volume of the object being covered. For a region O covered by a finite set of objects $(o_1, \ldots, o_k)$, the density is defined as $\frac{\sum_{i=0}^{k} \mathrm{vol}(o_i)}{\mathrm{vol}(O)}$. When covering unbounded regions [8], other means are needed in order to calculate that ratio, as we discuss next.

Next we present the methodology used in [8]. In their paper they show a lower bound on the density of covering $\mathbb{R}^d$ using spheres with radius 1. This lower bound is also an exact bound for $d = 2$ which is why we can compare our results to it. In their paper Coxeter *et al.* show that locating such spheres with their centers at each vertex of a regular simplex of edge length $\sqrt{\frac{2(d+1)}{d}}$, will cover the simplex. The density is the volume of the sections of these spheres that intersect the simplex divided by the volume of the simplex. Note that in the plane this lower bound can be attained as it is possible to cover $\mathbb{R}^2$ by interior-disjoint regular 2-simplices (namely, equilateral triangles), but this is not true for higher dimensions. Thus only yielding a lower bound for the density of covering $\mathbb{R}^d$ for $d \geqslant 3$. For $d = 2$ the density proven in [8] is $\frac{2\pi}{3\sqrt{3}} \approx 1.209$.

We now calculate the asymptotic density of the staggered grid when covering the $[\gamma, 1-\gamma]$-square, meaning for the case where the radius of the ball —β, tends to 0. This is the case of interest for us and this will reduce the influence of the β-balls on the edges of the square. As we state in Section 3, the number of β-balls in the staggered grid is $\left(\left\lceil \frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta} \right\rceil\right)^d +$

$\left(\left\lceil \frac{(1-2\gamma)\sqrt{d}}{\sqrt{8}\beta} \right\rceil + 1\right)^d$. The volume of each ball is $\pi\beta^2$. This yields the following expression for the volume of the cover using the staggered grid,

$$\begin{aligned}
&\left(\left(\left\lceil \frac{(1-2\gamma)}{2\beta} \right\rceil\right)^2 + \left(\left\lceil \frac{(1-2\gamma)}{2\beta} \right\rceil + \mathbf{1}\right)^2\right)(\pi\beta^2) \\
&\overset{(a)}{\approx} 2\left(\frac{(1-2\gamma)}{2\beta}\right)^2(\pi\beta^2) \\
&= (1-2\gamma)^2\frac{\pi}{2}\,,
\end{aligned}$$

where the approximation in (a) is due to ignoring the round-ups and the highlighted unit value 1, whose influence tends to 0 as β tend to 0.

The volume of the $[\gamma, 1-\gamma]$-square is $(1-2\gamma)^2$, Thus, the density of the staggered grid when covering the $[\gamma, 1-\gamma]$-square is,

$$\frac{(1-2\gamma)^2\frac{\pi}{2}}{(1-2\gamma)^2} = \frac{\pi}{2} \approx 1.5708.$$

Now we look at the density of the staggered grid divided by the density from [8]. This yields the following ratio,

$$\frac{\frac{\pi}{2}}{\frac{2\pi}{3\sqrt{3}}} = \frac{\pi}{2}\frac{3\sqrt{3}}{2\pi} = \frac{3\sqrt{3}}{4} \approx 1.299\,,$$

which means that asymptotically, for $d = 2$, the covering using [8] is more economical than our staggered grid by a multiplicative factor of 1.299.

We leave the employment of the approach in [8] in higher dimensions ($d > 2$) for future research. In their approach they cover the space with regular simplices and then cover each simplex with balls. This works well in two dimensions, where the equilateral triangles tile the plane. This is no longer the case in higher dimensions. In dimensions three and higher, there is a gap between the lower bound [8] and the best known covering.

Bibliography

[1] 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*, 12(4):1309–1317, 2015.

[2] R. Balakrishnan and P. Paulraja. Hamilton cycles in tensor product of graphs. *Discrete Mathematics*, 186(1):1 – 13, 1998.

[3] D. Bertschinger and P. Schnider. Weighted epsilon-nets. *CoRR*, abs/2002.08693, 2020.

[4] A. Bicchi, A. Fagiolini, and L. Pallottino. Towards a society of robots. *IEEE Robotics Automation Magazine*, 17(4):26–36, 2010.

[5] P. Brass, W. O. Moser, and J. Pach. *Research problems in discrete geometry*. Springer Science & Business Media, 2006.

[6] B. Carl and I. Stephani. *Entropy, Compactness and the Approximation of Operators*. Cambridge Tracts in Mathematics. Cambridge University Press, 1990.

[7] 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)*, pages 4543–4550, 2020.

[8] H. Coxeter, L. Few, and C. Rogers. Covering space with equal spheres. *Mathematika*, 6(2):147–157, 1959.

[9] D. Dayan, K. Solovey, M. Pavone, and D. Halperin. Near-optimal multi-robot motion planning with finite sampling. *CoRR*, abs/2011.08944, 2020.

[10] D. Dayan, K. Solovey, M. Pavone, and D. Halperin. Near-optimal multi-robot motion planning with finite sampling. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

[11] E. Fogel, D. Halperin, and R. Wein. *CGAL Arrangements and Their Applications - A Step-by-Step Guide*, volume 7 of *Geometry and Computing*. Springer, 2012.

[12] M. T. Goodrich and H. Brönnimann. Almost optimal set covers in finite VC-dimension. *Discrete & computational geometry*, 14(4):463–479, 1995.

[13] N. Goren and D. Halperin. DiscoPygal: Multi robot motion planning in Pyhton on CGAL, 2021. Manuscript.

[14] R. Hammack, W. Imrich, and S. Klavzar. *Handbook of Product Graphs, Second Edition*. CRC Press, Inc., USA, 2nd edition, 2011.

[15] 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)*, pages 10–12. IEEE, 2019.

[16] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.

[17] K. Hauser and Y. Zhou. Asymptotically optimal planning by feasible kinodynamic planning in a state-cost space. *IEEE Trans. Robotics*, 32(6):1431–1443, 2016.

[18] R. Herbrich. *Learning kernel classifiers: theory and algorithms*. MIT press, 2001.

[19] 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*, pages 4869–4873, 2017.

[20] 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*, 3(4):76–88, 1984.

[21] 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*, 34(7):883–921, 2015.

[22] J. K. Johnson. On the relationship between dynamics and complexity in multi-agent collision avoidance. *Autonomous Robots*, 42(7):1389–1404, 2018.

[23] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *International Journal of Robotics Research*, 30(7):846–894, 2011.

[24] L. E. Kavraki, M. N. Kolountzakis, and J. . Latombe. Analysis of probabilistic roadmaps for path planning. *IEEE Transactions on Robotics and Automation*, 14(1):166–171, 1998.

[25] L. E. Kavraki, P. Svestka, J. . Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.

[26] D. G. Kirkpatrick and P. Liu. Characterizing minimum-length coordinated motions for two discs. *CoRR*, abs/1607.04005, 2016.

[27] 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)*, pages 6344–6350, 2020.

[28] 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.*, 4(2):277–283, 2019.

[29] 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*, pages 1081–1087, 2014.

[30] 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)*, pages 995–1001, 2000.

[31] S. M. LaValle. *Planning algorithms*. Cambridge university press, 2006.

[32] 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*, pages 7627–7634, 2019.

[33] P. Long, T. Fan, X. Liao, W. Liu, H. Zhang, and J. Pan. Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6252–6259. IEEE, 2018.

[34] 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*, pages 7651–7658, 2019.

[35] C. Rogers. A note on coverings. *Mathematika*, 4(1):1–6, 1957.

[36] C. A. Rogers. A Note on Coverings and Packings. *Journal of the London Mathematical Society*, s1-25(4):327–331, 10 1950.

[37] F. Rossi, S. Bandyopadhyay, M. T. Wolf, and M. Pavone. Multi-agent algorithms for collective behavior: A structural and application-focused atlas. *CoRR*, abs/2103.11067, 2021.

[38] N. Rubin. Stronger bounds for weak epsilon-nets in higher dimensions. *CoRR*, abs/2104.12654, 2021.

[39] M. Rufli, J. Alonso-Mora, and R. Siegwart. Reciprocal collision avoidance with motion continuity constraints. *IEEE Trans. Robotics*, 29(4):899–912, 2013.

[40] C. Sammut and G. I. Webb. *Encyclopedia of machine learning*. Springer Science & Business Media, 2011.

[41] 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)*, pages 2112–2119, 2002.

[42] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence*, 219:40–66, 2015.

[43] Shitov. Counterexamples to Hedetniemi's conjecture. *Annals of Mathematics*, 190(2):663, 2019.

[44] R. Shome, K. Solovey, A. Dobson, D. Halperin, and K. E. Bekris. dRRT*: Scalable and informed asymptotically-optimal multi-robot motion planning. *Auton. Robots*, 44(3-4):443–467, 2020.

[45] I. Solis, J. Motes, R. Sandström, and N. M. Amato. Representation-optimal multi-robot motion planning using conflict-based search. *IEEE Robotics and Automation Letters*, 6(3):4608–4615, 2021.

[46] 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*, pages 799–816, 2018.

[47] K. Solovey and D. Halperin. k-Color multi-robot motion planning. *International Journal of Robotic Research*, 33(1):82–97, 2014.

[48] K. Solovey and D. Halperin. On the hardness of unlabeled multi-robot motion planning. *International Journal of Robotic Research*, 35(14):1750–1759, 2016.

[49] 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)*, pages 2189–2195, 2020.

[50] K. Solovey and M. Kleinbort. The critical radius in sampling-based motion planning. *International Journal of Robotics Research*, 39(2-3):266–285, 2020.

[51] 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*, 35(5):501–513, 2016.

[52] K. Solovey, J. Yu, O. Zamir, and D. Halperin. Motion planning for unlabeled discs with optimality guarantees. In *Robotics: Science and Systems*, 2015.

[53] P. G. Spirakis and C.-K. Yap. Strong NP-hardness of moving many discs. *Information Processing Letters*, 19(1):55–59, 1984.

[54] The CGAL Project. *CGAL User and Reference Manual*. CGAL Editorial Board, 5.3 edition, 2021.

[55] 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.

[56] M. Turpin, N. Michael, and V. Kumar. Trajectory planning and assignment in multirobot systems. In *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, pages 175–190, 2012.

[57] M. Vidyasagar. *A Theory of Learning and Generalization*. Springer-Verlag, Berlin, Heidelberg, 2nd edition, 2002.

[58] P. Švestka and M. H. Overmars. Coordinated path planning for multiple robots. *Robotics and Autonomous Systems*, 23(3):125–152, 1998.

[59] G. Wagner and H. Choset. Subdimensional expansion for multirobot path planning. *Artificial Intelligence*, 219:1–24, 2015.

[60] Z. Yan, N. Jouandeau, and A. A. Cherif. A survey and analysis of multi-robot coordination. *International Journal of Advanced Robotic Systems*, 10(12):399, 2013.

תקציר

התיזה עוסקת בבעיית *תכנון תנועה* (motion planning) לרובוטים תוך שימוש בכמות סופית של נקודות דגימה.

הגישה אשר עומדת בבסיס התיזה הינה תכנון תנועה מבוסס דגימות. בשיטה זו ממירים את הבעיה מהעולם הרציף לעולם הבדיד בעזרת דגימות של נקודות וחיבור בין נקודות שקרובות זו לזו. שיטה זו הינה השיטה הרווחת כיום לפתרון בעיות בהן מעורבים מספר רובוטים או בעיות תנועה עבור רובוט בודד בעל מספר רב של דרגות חופש של תנועה.

הבעיה הראשונה בה עוסקת התיזה היא בעיית כיסוי קוביה (במימד כללי) ע״י כדורים ברדיוס קטן כלשהו. עבור בעיה זו אנו מציעים שיטה בשם ה-staggered grid ומוכיחים שהיא משפרת באופן משמעותי תוצאות קודמות (אך אינה אופטימלית). בעבודה אנו מנתחים את השיטה שלנו אל מול חסמים עליונים ותחתונים מעבודות קודמות.

תוך שימוש בשיטת הדגימה המצוינת לעיל, אנו מראים שיפור משמעותי למול תוצאות קודמות עבור בעיית תכנון תנועה לרובוט יחיד בעזרת מספר סופי של נקודות דגימה.

לאחר מכן אנו מרחיבים את התוצאה למספר כלשהו של רובוטים. בחלק זה אנו מציגים תוצאה ראשונה מסוגה: תכנון תנועה למספר רב של רובוטים כלליים עם ערבויות על איכות המסלולים *תוך שימוש במספר סופי של נקודות דגימה*. חלק זה הינו החלק המרכזי של התיזה והוא מסתמך רבות על הפרקים הקודמים. התוצאה התיאורטית מגובה בניסויים בתוכנה בהם אנו מראים כי התוצאות הפרקטיות אף טובות ממה שהתיאוריה מבטיחה.

לסיום אנו סוקרים מספר נושאים מעניינים נוספים אשר עלו מהמחקר. ראשית אנו משווים בין דגימה בשיטתנו למול דגימה רנדומלית אחידה ומראים את השיפור בסיכויי ההצלחה ובאיכות התוצאות. בנוסף אנו סוקרים יותר לעומק את בעיית הכיסוי, מעלים רעיונות כיצד ניתן לשפר את שיטת הדגימה שלנו ומשערים איך ניתן לנצל חסמים משופרים על כיסוי קוביה כדי לשפר את החסמים עבור בעיית תכנון התנועה.

מאמר המסכם את תוצאות התזה התפרסם בכנס ICRA 2021:

D. Dayan, K. Solovey, M. Pavone, and D. Halperin. Near-Optimal Multi-Robot Motion Planning with Finite Sampling. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

תכנון תנועה כמעט אופטימלי תוך שימוש בכמות דגימות סופית עבור מספר רובוטים

חיבור זה מוגש כחלק ממילוי הדרישות לקבלת
התואר "מוסמך למדעים" (M.Sc.) בבית הספר למדעי המחשב
באוניברסיטת תל-אביב
ע"י

דרור דיין

העבודה הוכנה באוניברסיטת תל-אביב
בהדרכתו של פרופסור דן הלפרין
חשוון התשפ"ב