

The Critical Radius in Sampling-based Motion Planning

Kiril Solovey* and Michal Kleinbort†

December 27, 2018

Abstract

We develop a new analysis of sampling-based motion planning in Euclidean space with uniform random sampling, which significantly improves upon the celebrated result of Karman and Frazzoli (2011) and subsequent work. Particularly, we prove the existence of a critical connection radius proportional to $\Theta(n^{-1/d})$ for n samples and d dimensions: Below this value the planner is guaranteed to fail (similarly shown by the aforementioned work, *ibid.*). More importantly, for larger radius values the planner is asymptotically (near-)optimal. Furthermore, our analysis yields an explicit lower bound of $1 - O(n^{-1})$ on the probability of success. A practical implication of our work is that asymptotic (near-)optimality is achieved when each sample is connected to only $\Theta(1)$ neighbors. This is in stark contrast to previous work which requires $\Theta(\log n)$ connections, that are induced by a radius of order $\left(\frac{\log n}{n}\right)^{1/d}$. Our analysis is not restricted to PRM and applies to a variety of “PRM-based” planners, including RRG, FMT* and BTT. Continuum percolation plays an important role in our proofs. Lastly, we develop similar theory for all the aforementioned planners when constructed with deterministic samples, which are then sparsified in a randomized fashion. We believe that this new model, and its analysis, is interesting in its own right.

1 Introduction

Motion planning is a fundamental problem in robotics concerned with allowing autonomous robots to efficiently navigate in environments cluttered with obstacles. Although motion planning has originated as a strictly theoretical problem in computer science [18], nowadays it is applied in various fields. Notably, motion planning arises in coordination of multiple autonomous vehicles [12], steering surgical needles [5], and planning trajectories of spacecrafts in orbit [47], to name just a few examples.

Motion planning is notoriously challenging from a computational perspective due to the continuous and high-dimensional search space it induces, which accounts for the structure of the robot, the physical constraints that it needs to satisfy, and the environment in which it operates. Low-dimensional instances of the problem can typically be solved by efficient and complete algorithms [17], and the same applies to several high-dimensional instances (see, e.g., [1, 44]). However, in general, motion planning is computationally intractable (see, e.g., [7, 35, 42]).

Nowadays the majority of practical approaches for motion planning capture the connectivity of the free space by sampling (typically in a randomized fashion) configurations and connecting nearby samples, to form a graph data structure. Although such *sampling-based planners* are inherently incomplete, i.e., cannot detect situations in which a solution (collision-free path) does

*Department of Aeronautics and Astronautics, Stanford University, CA 94305, US. email: kirilsole@stanford.edu

†Blavatnik School of Computer Science, Tel Aviv University, Israel. email: balasmic@post.tau.ac.il

not exist, most have the desired property of being able to find a solution *eventually*, if one exists. That is, a planner is *probabilistically complete* (PC) if the probability of finding a solution tends to 1 as the number of samples n tends to infinity. Moreover, some recent sampling-based techniques are also guaranteed to return high-quality solutions that tend to the optimum as n diverges—a property called *asymptotic optimality* (AO). Quality can be measured in terms of energy, length of the plan, clearance from obstacles, etc.

An important attribute of sampling-based planners, which dictates both the running time and the quality of the returned solution, is the number of neighbors considered for connection for each added sample. In many techniques this number is directly affected by a connection radius r_n : Decreasing r_n reduces the number of neighbors. This in turn reduces the running time of the planner for a given number of samples n , but may also reduce the quality of the solution or its availability altogether. Thus, it is desirable to come up with a radius r_n that is small, but not to the extent that the planner loses its favorable properties of PC and AO.

1.1 Contribution

We develop a new analysis of PRM [24] for uniform random sampling in Euclidean space, which relies on a novel connection between sampling-based planners and *continuum percolation* (see, e.g., [11]). Our analysis is tight and proves the existence of a *critical connection* radius $r_n^* = \gamma^* n^{-1/d}$, where $\gamma^* > 0$ is a constant (in particular, $0.4 \leq \gamma^* \leq 0.6$ for all $d \geq 2$), and $d \geq 2$ is the dimension: If $r_n < r_n^*$ then PRM is guaranteed to fail, where d is the dimension. Above the threshold, i.e., when $r_n > r_n^*$, PRM is AO for the bottleneck cost, and *asymptotically near optimal*¹ (AnO) with respect to the path-length cost. Furthermore, our analysis yields concrete bounds on the probability of success, which is lower-bounded by $1 - O(n^{-1})$. Notice that this bound is comparable to the one obtained in [48] (see Section 2) although we show this for a much smaller radius.

Our analysis is not restricted to PRM and applies to a variety of planners that maintain PRM-like roadmaps, explicitly or implicitly. For instance, when r_n is above the threshold, FMT* [21] is AnO with respect to the path-length cost, while BTT [43] is AO with respect to the bottleneck cost. RRG [23] behaves similarly for the two cost functions. Our results are also applicable to *multi-robot* motion planners such as the recently introduced dRRT* [10], and M* [51] when applied to a continuous domain. See Figure 1 for additional PRM-based planners to which our analysis is applicable, and which are mentioned further on.

A practical implication of our work is that AO (or AnO), under the regime of uniform random sampling, can be achieved even when every sample is connected to $\Theta(1)$ neighbors. This is in stark contrast with previous work, e.g., [21, 23, 41, 45], which provided a rough estimate of this number, that is proportional to $O(\log n)$. Interestingly, our $\Theta(1)$ bound for uniform samples is comparable to the best known bound of for deterministic samples [22].

Lastly, we also study the asymptotic properties of the aforementioned planners when constructed using a deterministic point process, rather than uniform random sampling which we have mentioned so far. We introduce an analysis which shows that the graphs constructed using such deterministic samples can be sparsified in a randomized fashion while maintaining AO or AnO of the planners. We term this regime as *semi-deterministic sampling*.

¹AnO means that the cost of the solution tends to at most a constant factor times the optimum, compared with AO in which this constant is equal to one.

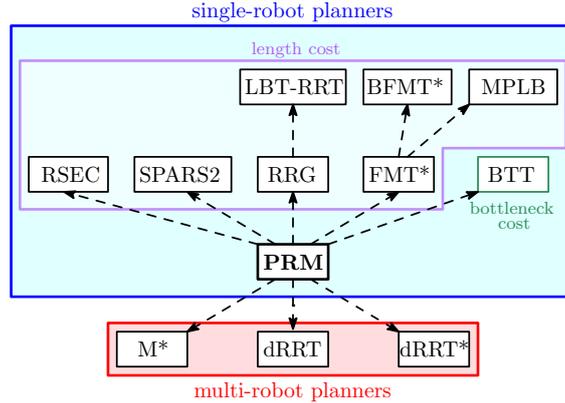


Figure 1: The PRM dynasty. Single robot and multi-robot planners, are bounded into the blue and red frames, respectively. Planners inside the magenta frame aim to minimize the path-length cost, whereas BTT is designed for bottleneck cost, and RRG works for both costs (see more information below). Roughly speaking, arrow from “A” to “B” indicates that theoretical properties of “A” extend to “B”, or that the latter maintains “A” as a substructure.

1.2 Organization

In Section 2 we discuss related work. Then we proceed to basic definitions and problem statement in Section 3. In that section we also include a precise description of the robotic system our theory is developed for. Our central contribution (Theorem 1), which states the existence of a critical radius r_n^* with respect to PRM, is presented in Section 4. This section also contains an outline of the proof. In Section 5 we lay the foundations of the proof and discuss important aspects of continuum percolation, that would later be employed in the main proof, which appears in Section 6. In Section 7 we extend the analysis to FMT*, BTT, RRG and discuss the implications of our results to the multi-robot planners dRRT* and M*. In Section 8 we present experimental work, which validates our theory. We conclude this paper with directions for future work (Section 9).

In the appendix we include proofs that were omitted from the main document and a table with values of γ^* . We also present in the appendix a new analysis for the aforementioned planners when constructed with deterministic samples, which are then sparsified in a randomized fashion. We believe that this new model, and its analysis, is interesting in its own right.

2 Related work

This section is devoted to a literature review of sampling-based planners with emphasis on their theoretical analysis. Sampling-based techniques were first described in the mid 90’s and several prominent examples of that time, that are still used extensively today, include PRM [24], RRT [29], and EST [20]. Those planners are also known to be PC (see [25, 28, 27, 20]).

More recent work has been concerned with the quality of the returned solution. [32] were one of the first to address this matter in a mathematically-rigorous manner. This work proved that RRT can produce arbitrarily-bad (long) paths with non-negligible probability.

The influential work of [23] laid the theoretical foundations for analyzing quality in sampling-based planning. The authors introduced a set of techniques to prove AO. Using their framework, they showed that the following algorithms are AO: PRM*, which is a special case of PRM (throughout this work we will refer to the more general algorithm PRM rather than PRM*) with a specific value of the connection radius r_n ; an AO variant of RRT [29] termed RRT*; RRG, which can be viewed as a combination between RRT and PRM. The analysis in [23] establishes that

$r_n = \Theta((\log n/n)^{1/d})$ guarantees AO, where the configuration space of the robot is assumed to be $[0, 1]^d$. This indicates that the expected number of neighbors used per vertex should be $O(\log n)$. The authors also proved that for sufficiently-small radii of order $O(n^{-1/d})$ the planner is guaranteed to fail (asymptotically) in finding any solution.

Following the breakthrough of [23], other AO planners have emerged (see e.g., [2, 4, 14]). [21] introduced FMT^* , which is a single-query planner that traverses an implicitly-represented PRM graph, and is comparable in performance to RRT^* . The authors refined the proof technique of [23], which allowed them to slightly reduce the connection radius r_n necessary to FMT^* and PRM to achieve AO. We do mention that here again $r_n = \Theta((\log n/n)^{1/d})$. BFMT^* , which is a bidirectional version of FMT^* , was introduced by [48]. In this paper the authors also proved that the success rate of PRM , FMT^* , BFMT^* can be lower bounded by $1 - O(n^{-\eta/d} \log^{-1/d} n)$, where $\eta > 0$ is a tuning parameter. In this context, we also mention the work of [9], which bounds the success rate with an expression that depends on the amount of deviation from the optimum.

A recent work by [45] developed a different method for analyzing sampling-based planners. It exploits a connection with *random geometric graphs* (RGGs), which have been extensively studied (see, e.g., [33]). Their work shows that one can slightly reduce the PRM and FMT^* radius obtained in [21]. Furthermore, the connection with RGGs yields additional analyses of different extensions of PRM , which have not been analyzed before in a mathematically-rigorous setting.

A number of methods have been developed to reduce the running time or space requirements of existing planners by relaxing AO constraints to AnO. For instance, LBT-RRT [37] interpolates between the quick RRT and the AO yet slower RRG , according to a predefined parameter. MPLB [36] can be viewed as a relaxation of FMT^* . SPARS2 [8] and RSEC [38] perform a sparsification of PRM in an online or offline fashion, respectively, to reduce the space footprint of the produced roadmap.

2.1 Extensions

The aforementioned papers deal mainly with the cost function of *path length*. Two recent works [41, 43] considered the *bottleneck-pathfinding* problem in a sampling-based setting and introduced the BTT algorithm, which traverses an implicitly-represented PRM graph. The bottleneck-cost function, which arises for instance in high-clearance multi-robot motion [43] and Fréchet matching between curves [19], is defined as follows: Every robot configuration x is paired with a value $\mathcal{M}(x)$, and the cost of a path is the maximum value of \mathcal{M} along any configuration on the path. It was shown [41, 43] that BTT is AO, with respect to bottleneck cost, for the reduced connection radius that was obtained by [45].

The results reported until this point have dealt exclusively with holonomic robotic systems. Two recent papers by [39, 40] develop the theoretical foundations of PRM and FMT^* -flavored planners when applied to robots having differential constraints. [30] develop an AO algorithm that does not require a steering function, as PRM for instance does. Interestingly, the authors of the last paper also obtain bounds on the rate of convergence of their algorithm.

3 Preliminaries

We provide several basic definitions that will be used throughout the paper. Given two points $x, y \in \mathbb{R}^d$, denote by $\|x-y\|$ the standard Euclidean distance. Denote by $\mathcal{B}_r(x)$ the d -dimensional ball of radius $r > 0$ centered at $x \in \mathbb{R}^d$ and $\mathcal{B}_r(\Gamma) = \bigcup_{x \in \Gamma} \mathcal{B}_r(x)$ for any $\Gamma \subseteq \mathbb{R}^d$. Similarly, given a curve $\pi : [0, 1] \rightarrow \mathbb{R}^d$ define $\mathcal{B}_r(\pi) = \bigcup_{\tau \in [0, 1]} \mathcal{B}_r(\pi(\tau))$. Given a subset $D \subset \mathbb{R}^d$ we denote by $|D|$ its Lebesgue measure. All logarithms are at base e .

3.1 Motion planning

Denote by \mathcal{C} the configuration space of the robot, and by $\mathcal{F} \subseteq \mathcal{C}$ the free space, i.e., the set of all collision free configurations. Though our proofs may be extended to more complex robotic systems (see discussion in Section 9), in this work we investigate the geometric (holonomic) setting of the problem in which no constraints are imposed on the motion of the robot. Additionally, we assume that \mathcal{C} is some subset of the Euclidean space. In particular, $\mathcal{C} = [0, 1]^d \subset \mathbb{R}^d$ for some fixed $d \geq 2$. We also assume that for any two configurations $x, x' \in \mathcal{C}$ the robot is capable of following precisely the straight-line path from x to x' .

Given start and target configurations $s, t \in \mathcal{F}$, the problem consists of finding a continuous path (curve) $\pi : [0, 1] \rightarrow \mathcal{F}$ such that $\pi(0) = s, \pi(1) = t$. That is, the robot starts its motion along π on s , and ends in t , while remaining collision free. An instance of the problem is defined for a given (\mathcal{F}, s, t) , where $s, t \in \mathcal{F}$.

3.2 Cost function

It is usually desirable to obtain paths that minimize a given criterion. In this paper we consider the following two cost functions.

Definition 1. Given a path σ , its *length* is

$$c_\ell(\sigma) = \sup_{n \in \mathbb{N}_+, 0 = \tau_1 \leq \dots \leq \tau_n = 1} \sum_{i=2}^n \|\pi(\tau_i) - \pi(\tau_{i-1})\|.$$

Definition 2. Given a path σ , and a *cost map* $\mathcal{M} : \mathcal{C} \rightarrow \mathbb{R}$, its *bottleneck cost* is

$$c_b(\sigma, \mathcal{M}) = \max_{\tau \in [0, 1]} \mathcal{M}(\pi(\tau)).$$

We proceed to describe the notion of *robustness*, which is essential when discussing properties of sampling-based planners. Given a subset $\Gamma \subset \mathcal{C}$ and two configurations $x, y \in \Gamma$, denote by $\Pi_{x,y}^\Gamma$ the set of all continuous paths, whose image is in Γ , that start in x and end in y , i.e., if $\pi \in \Pi_{x,y}^\Gamma$ then $\pi : [0, 1] \rightarrow \Gamma$ and $\pi(0) = x, \pi(1) = y$.

Definition 3. Let (\mathcal{F}, s, t) be a motion-planning problem. A path $\pi \in \Pi_{s,t}^\mathcal{F}$ is *robust* if there exists $\delta > 0$ such that $\mathcal{B}_\delta(\pi) \subset \mathcal{F}$. We also say that (\mathcal{F}, s, t) is *robustly feasible* if there exists such a robust path.

Definition 4. The *robust optimum* with respect to c_ℓ is defined as

$$c_\ell^* = \inf \{c_\ell(\pi) \mid \pi \in \Pi_{s,t}^\mathcal{F} \text{ is robust}\}.$$

The corresponding definition for the bottleneck cost is slightly more involved.

Definition 5. Let \mathcal{M} be a cost map. A path $\pi \in \Pi_{s,t}^\mathcal{F}$ is *\mathcal{M} -robust* if it is robust and for every $\varepsilon > 0$ there exists $\delta > 0$ such that for every $x \in \mathcal{B}_\delta(\pi)$, $\mathcal{M}(x) \leq (1 + \varepsilon)c_b(\pi, \mathcal{M})$. We also say that \mathcal{M} is *well behaved* if there exists at least one \mathcal{M} -robust path.

Definition 6. The *robust optimum* with respect to c_b is defined as

$$c_b^* = \inf \{c_b(\pi, \mathcal{M}) \mid \pi \in \Pi_{s,t}^\mathcal{F} \text{ is } \mathcal{M}\text{-robust}\}.$$

3.3 Poisson point processes

We draw our main analysis techniques from the literature of continuum percolation, where point samples are generated with the following distribution. Thus we will use this point distribution in PRM, which would be formally defined in the following section.

Definition 7 [11]. A random set of points $\mathcal{X} \subset \mathbb{R}^d$ is a *Poisson point process* (PPP) of density $\lambda > 0$ if it satisfies the conditions:

1. For mutually disjoint domains $D_1, \dots, D_\ell \subset \mathbb{R}^d$, the random variables $|D_1 \cap \mathcal{X}|, \dots, |D_\ell \cap \mathcal{X}|$ are mutually independent.
2. For any bounded domain $D \subset \mathbb{R}^d$ we have that for every $k \geq 0$,

$$\Pr[|\mathcal{X} \cap D| = k] = e^{-\lambda|D|} \frac{(\lambda|D|)^k}{k!}.$$

Informally, property (1) states that the number of points from \mathcal{X} in any two disjoint sets $D_i, D_j \subset \mathbb{R}^d$ is independent. Another useful property that follows from (2) is that the expected number of points in a certain region is known precisely and corresponds to the volume of this region. That is, for any $D \subset \mathbb{R}^d$ it holds that $E(|\mathcal{X} \cap D|) = \lambda|D|$.

The following provides a simple recipe for generating PPP.

Claim 1 [11]. Let N be a Poisson random variable with mean λ . For a given $z \in \mathbb{Z}^d$ draw a sample $N_z \in N$ and define $\mathcal{X}_z = \{X_1, \dots, X_{N_z}\}$ to be N_z points chosen independently and uniformly at random from $z + [0, 1]^d$. Then $\mathcal{X} = \bigcup_{z \in \mathbb{Z}^d} \mathcal{X}_z$ is a PPP of density λ .

It will be convenient to think about PRM as a subset of the following *random geometric graph* (RGG). We will describe various properties of this graph in later sections.

Definition 8. [33] Let $\mathcal{X} \subset \mathbb{R}^d$ be a PPP. Given $r > 0$, the random geometric graph $\mathcal{G}(\mathcal{X}; r)$ is an undirected graph with the vertex set \mathcal{X} . Given two vertices $x, y \in \mathcal{X}$, $(x, y) \in \mathcal{G}(\mathcal{X}; r)$ if $\|x - y\| \leq r$.

4 Analysis of PRM

In this section we provide a mathematical description of PRM, which essentially maintains an underlying RGG with PPP samples. We proceed to describe our main contribution (Theorem 1) which is concerned with the conditions for which PRM converges to the (robust) optimum. We then provide an outline of the proof, in preparation for the following sections.

Recall that the configuration space of the robot is represented by $\mathcal{C} = [0, 1]^d$, and the free space is denoted by $\mathcal{F} \subseteq \mathcal{C}$. The motion-planning problem (\mathcal{F}, s, t) will remain fixed throughout this section.

Recall that PRM accepts as parameters the number of samples $n \in \mathbb{N}_+$ and a connection radius r_n . Denote by \mathcal{X}_n a PPP with mean density n . In relation to the definitions of the previous section, the graph data structure obtained by the *preprocessing stage* of PRM can be viewed as an RGG. For instance, when $\mathcal{F} = \mathcal{C}$, the graph obtained by PRM is precisely $\mathcal{G}(\mathcal{X}_n \cap [0, 1]^d; r_n)$. In the more general case, when $\mathcal{F} \subset \mathcal{C}$, PRM produces the graph $\mathcal{G}(\mathcal{X}_n \cap \mathcal{F}; r_n)$. As \mathcal{F} can be non-convex, we emphasize that the latter notation describes the maximal subgraph of $\mathcal{G}(\mathcal{X}_n; r_n)$ such that vertices and edges are contained in \mathcal{F} .

In the *query stage*, recall that PRM accepts two configurations $s, t \in \mathcal{F}$, which are then connected to the preprocessed graph. Here we slightly diverge from the standard definition of PRM in the literature. In particular, instead of using the same radius r_n when connecting s, t , we use the (possibly larger) radius r_n^{st} . The graph obtained after query is formally defined below:

Definition 9. The PRM graph \mathcal{P}_n is the union between $\mathcal{G}(\mathcal{X}_n \cap \mathcal{F}; r_n)$ and the supplementary edges

$$\bigcup_{y \in \{s, t\}} \{(x, y) \mid x \in \mathcal{X}_n \cap \mathcal{B}_{r_n^{st}}(y) \text{ and } xy \subset \mathcal{F}\}.$$

Remark. We emphasize that the larger radius r_n^{st} is only used when s, t are connected to the preprocessed RGG.

We reach our main contribution.

Theorem 1. *Suppose that (\mathcal{F}, s, t) is robustly feasible. Then there exists a critical radius $r_n^* = \gamma^* n^{-1/d}$, where γ^* is a constant (see supplementary material), such that the following holds:*

- i. *If $r_n < r_n^*$ and $r_n^{st} = \infty$ then PRM fails (to find a solution) a.a.s.²*
- ii. *Suppose that $r_n > r_n^*$. There exists $\beta_0 > 0$ such that for $r_n^{st} = \frac{\beta \log^{1/(d-1)} n}{n^{1/d}}$, where $\beta \geq \beta_0$, and any $\varepsilon > 0$ the following holds with probability $1 - O(n^{-1})$:*
 1. *\mathcal{P}_n contains a path $\pi_n \in \Pi_{s,t}^{\mathcal{F}}$ with $c_\ell(\pi_n) \leq (1 + \varepsilon)\xi c_\ell^*$, where ξ is independent of n ;*
 2. *If \mathcal{M} is well behaved then \mathcal{P}_n contains a path $\pi'_n \in \Pi_{s,t}^{\mathcal{F}}$ with $c_b(\pi'_n, \mathcal{M}) \leq (1 + \varepsilon)c_b^*$.*

Remark: Theorem 1 also holds, with a slight modification, when the PPP is replaced with the standard binomial point process (BPP), in which the number of sampled points is fixed a priori. The only difference is that the probability of success should be reduced from $1 - O(n^{-1})$ to $1 - O(n^{-1/2})$. This follows from Lemma 1 in [13], which states that if a property holds for PPP then it also holds for BPP, albeit with slightly smaller probability.

4.1 Outline of proof

For the remainder of this section we briefly describe our technique for proving this theorem, in preparation for the full proof, which is given in Section 6. The critical radius r_n^* defined above, coincides with the *percolation threshold*, which determines the emergence of a connected component of \mathcal{G} that is of infinite size. In particular, if $r_n < r_n^*$ then $\mathcal{G}(\mathcal{X}_n; r_n)$ breaks into tiny connected components of size $O(\log n)$ each. Thus, unless s, t are infinitesimally close, no connected component can have both s and t simultaneously.

More interestingly, the radius of $r_n > r_n^*$ leads to the emergence of a unique infinite component of $\mathcal{G}(\mathcal{X}_n; r_n)$. That is, in such a case one of the components of $\mathcal{G}(\mathcal{X}_n; r_n)$ must contain an infinite number of vertices (Section 5.1). Denote this component by C_∞ .

In contrast to $\mathcal{G}(\mathcal{X}_n; r_n)$, which is defined for the unbounded space \mathbb{R}^d , our motion-planning problem is bounded to $\mathcal{C} = [0, 1]^d$. Thus, the next step is to investigate the properties of $\mathcal{G}(\mathcal{X}_n; r_n)$ when restricted to $[0, 1]^d$ (Section 5.2). Denote by C_n the largest connected component of $C_\infty \cap [0, 1]^d$. This structure plays a key role in our proof (see Section 6): With high probability (to be defined), there exist vertices of C_n that are sufficiently close to s and t , respectively, so that a connection between the two vertices can be made through C_n .

Of course, this overlooks the fact that some portions of C_n lie in forbidden regions of \mathcal{C} . Thus, we also have to take the structure of \mathcal{F} into consideration. To do so, we rely on [34] to prove that any small subset of $[0, 1]^d$ must contain at least one point of C_n (see lemmata 2 and 3). This allows us to trace the robust optimum (and collision-free) path with points from C_n .

The final ingredient, which allows to bound the path length along $\mathcal{G}(\mathcal{X}_n; r_n) \cap [0, 1]^d$, is Theorem 4. It states that the distance over this graph is proportional to the Euclidean distance between

²Let A_1, A_2, \dots be random variables in some probability space and let B be an event depending on A_n . We say that B occurs *asymptotically almost surely* (a.a.s., in short) if $\lim_{n \rightarrow \infty} \Pr[B(A_n)] = 1$.

the end points. This also ensures that the trace points from C_n can be connected with collision-free paths over the graph.

To conclude, in Section 5 we provide background on continuum percolation in unbounded and bounded domains, and prove two key lemmata (Lemma 2 and Lemma 3). In Section 6 we return to the setting of motion planning and utilize the aforementioned results in the proof of Theorem 1.

5 Elements of continuum percolation

In this section we describe some of the properties of the unbounded graph $\mathcal{G}(\mathcal{X}_n; r_n)$ that will be employed in our analysis in the following section.

5.1 The basics

A fundamental question is when \mathcal{G} contains an infinite connected component around the origin.

Definition 10. The *percolation probability* $\theta(n, r)$ is the probability that the origin $o \in \mathbb{R}^d$ is contained in a connected component of $\mathcal{G}(\mathcal{X}_n \cup \{o\}; r)$ of an infinite number of vertices. That is, if C_o denotes the set of vertices connected to o in the graph, then $\theta(n, r) = \Pr(|C_o| = \infty)$.

We say that a graph *percolates* iff $\theta(n, r) > 0$. Note that the selection of the origin is arbitrary, and the following result can be obtained for any $x \in \mathbb{R}^d$ alternative to o .

Theorem 2. [16, Theorem 12.35] *There exists a critical radius $r_n^* = \gamma^* n^{-1/d}$, where γ^* is a constant, such that $\theta(n, r_n) = 0$ when $r_n < r_n^*$, and $\theta(n, r_n) > 0$ when $r_n > r_n^*$.*

The following lemma states that the infinite connected component exists with probability strictly 0 or 1.

Lemma 1. *Let $\psi(n, r)$ be the probability that $\mathcal{G}(\mathcal{X}_n; r)$ contains an infinite connected component, i.e., without conditioning on any specific additional vertex. Then $\psi(n, r) = 0$ when $\theta(n, r) = 0$ and $\psi(n, r) = 1$ when $\theta(n, r) > 0$.*

Proof. Suppose that $\theta(n, r) = 0$, and for any $x \in \mathbb{R}^d$, denote by $\theta_x(n, r)$ the percolation probability of $\mathcal{G}(\mathcal{X}_n \cup \{x\}; r)$, and note that $\theta_x(n, r) = 0$. Define $Z_r = \{r \cdot z \mid z \in \mathbb{Z}^d\}$, and observe that for any $y \in \mathbb{R}^d$ there exists $z \in Z_r$ such that $y \in \mathcal{B}_r(z)$. By definition, $\mathcal{G}(\mathcal{X}_n \cup \{x\}; r)$ percolates, i.e., the infinite component touches $\mathcal{B}_r(x)$, iff there exists $y \in C_\infty$ such that $\|x - y\| \leq r$. By definition of Z_r , if the latter event occurs then there exists $z \in Z_r$ such that $\|z - y\| \leq r$. Thus, by applying the union bound, and noting that a sum of countable number of zeros is still zero, we establish that

$$\begin{aligned} \psi(n, r) &= \Pr \left[\exists x \in \mathbb{R}^d, \mathcal{G}(\mathcal{X}_n \cup \{x\}; r) \text{ percolates} \right] \\ &= \Pr \left[\exists z \in Z_r, \mathcal{G}(\mathcal{X}_n \cup \{z\}; r) \text{ percolates} \right] \\ &\leq \sum_{z \in Z_r} \theta_z(n, r) = 0. \end{aligned}$$

For the other direction we employ Kolmogorov's zero-one law (see, e.g., [6, Theorem 1, p.36]). Informally, it states that an event, e.g., existence of an infinite connected component in \mathcal{G} , that occurs independently of any finite subset of independent random variables, e.g., points from \mathcal{X}_n , has a ability of either 0 or 1. Thus, as $\psi(n, r) \geq \theta(n, r)$ it immediately follows that $\psi(n, r) = 1$ when $\theta(n, r) > 0$. \square

The following theorem establishes that the infinite connected component is unique.

Theorem 3. [31, Theorem 2.3] *With probability 1, $\mathcal{G}(\mathcal{X}_n; r)$ contains at most one infinite connected component.*

5.2 Bounded domains

We study different properties of $\mathcal{G}(\mathcal{X}_n; r)$ when it is restricted to the domain $[0, 1]^d$. In case that $\theta(n, r) > 0$, we use C_∞ to refer to the infinite connected component of the unbounded graph $\mathcal{G}(\mathcal{X}_n; r)$. Note that C_∞ exists (Lemma 1) and is unique (Theorem 3) with probability 1.

Denote by C_n the largest connected component of $C_\infty \cap [0, 1]^d$. By definition, C_n is also a subgraph of $\mathcal{G}(\mathcal{X}_n \cap [0, 1]^d; r_n)$. The following lemma shows that with high probability all the points from C_∞ , that are sufficiently close to the center of $[0, 1]^d$, are members of C_n .

Lemma 2. *Let $r_n > r_n^*$. Define*

$$H_n = [0, 1]^d \setminus \mathcal{B}_{1/\log n} \left(\mathbb{R}^d \setminus [0, 1]^d \right).$$

Denote by \mathfrak{E}_n^1 the event that $C_\infty \cap H_n \subset C_n$. Then there exist $n_0 \in \mathbb{N}$ and $\alpha > 0$ such that for any $n > n_0$ it holds that

$$\Pr[\mathfrak{E}_n^1] \geq 1 - \exp\left(-\alpha n^{1/d} \log^{-1} n\right).$$

Proof. This statement is an adaptation of Lemma 8 and Theorem 2 of [34]. We mention that [34] uses a slightly different, but nevertheless equivalent model. While we consider an RGG that is bounded to $[0, 1]^d$ and PPP of density n , they consider the domain $[0, n]^d$ with density 1. It is only a matter of rescaling and variable substitution to import their results to our domain. Let C'_1, \dots, C'_k denote the connected components of $C_\infty \cap [0, 1]^d$, and set C_n to be the component C'_i with the largest number of vertices. Without loss of generality, $C_n = C'_1$. See illustration in Figure 2.

Observe that if $C'_i \subset C_\infty$ then it must have at least one vertex $x'_i \in C'_i$ that lies closely to the boundary of $[0, 1]^d$, i.e., $\|x - \partial([0, 1]^d)\| \leq r_n$, as otherwise C'_i will not be able to connect to the rest of C_∞ . Thus, it must be the case that $\text{diam}(C'_i) \leq \log^{-1} n - r_n$ in order to be able to reach H_n , where $\text{diam}(D) = \sup_{x, x' \in D} \|x - x'\|$ defines the diameter of a given $D \subset \mathbb{R}^d$. We shall show that this does not hold.

Theorem 2 in [34] states that for any ϕ_n large enough there exist α', n_0 such that for any $n > n_0$ it holds with probability at least $1 - \exp(-\alpha' n^{1/d} \phi_n)$ that $\text{diam}(C'_i) < \phi_n$ for any $1 < i \leq k$. By setting $\alpha = \alpha'/2$, $\phi_n = 1/(2 \log n)$ the conclusion immediately follows. \square

Observe that for any fixed point x internal to $[0, 1]^d$ there exists $n_0 \in (0, \infty)$ such that for any $n > n_0$ it holds that $x \in H_n$. The following lemma bounds the probability of having at least one point from C_n in a (small) subregion of H_n . Note that the value β corresponds to the value used in Theorem 1.

Lemma 3. *Let $r_n > r_n^*$. Define $H'_n \subset H_n$ to be a hypercube of side length*

$$h'_n = \frac{\beta \log^{1/(d-1)} n}{n^{1/d}}.$$

Denote by \mathfrak{E}_n^2 the event that $H'_n \cap C_n \neq \emptyset$. Then there exists $n_0 \in \mathbb{N}$ and $\beta_0 > 0$ such that for any $n > n_0$, $\beta > \beta_0$ it holds that $\Pr[\mathfrak{E}_n^2 | \mathfrak{E}_n^1] \geq 1 - n^{-1}$.

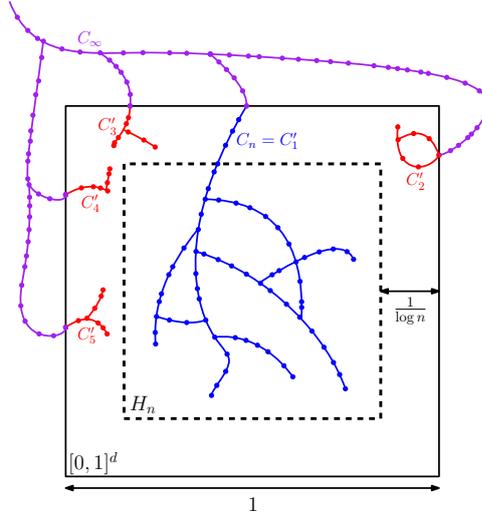


Figure 2: Illustration for Lemma 2. The outer cube represents $[0, 1]^d$, whereas the inner cube depicted with dashed boundary is H_n . Observe that H_n has side length of $1 - 2/\log n$. The purple, blue and red graphs combined describe C_∞ , whereas C_n is depicted in blue, and C'_2, C'_3, C'_4, C'_5 are in red.

Proof. Define $G_n = \mathcal{G}(\mathcal{X}_n, r_n) \cap H'_n$, $n' = E[|\mathcal{X}_n \cap H'_n|]$ and observe that $n' = n \cdot |H'_n| = \beta^d \log^{d/(d-1)} n$.

We treat G_n as a subset of \mathbb{R}^d in order to apply a rescaling argument. Observe that (scalar) multiplication of every point of H'_n with $1/h'_n$ yields a translation of $[0, 1]^d$. We will use the superscript $1/h'_n$ to describe this rescaling to a given object. For instance, applying the same transformation on G_n yields the graph G_n^{1/h'_n} , which has the same topology as G_n . Denote by $x_\ell \in H'_n$ the (lexicographically) smallest point of H'_n , i.e., $x_\ell = (1/\log n, \dots, 1/\log n)$. Notice that

$$\begin{aligned} G_n^{1/h'_n} - x_\ell^{1/h'_n} &= \mathcal{G}(\mathcal{X}_n^{1/h'_n} \cap [0, 1]^d, r_n/h'_n) \\ &= \mathcal{G}(\mathcal{X}_{n'} \cap [0, 1]^d, r_{n'}), \end{aligned}$$

where the minus sign in the left-hand side represents a translation by a vector. This implies that G_n , which is defined over H'_n behaves as $\mathcal{G}(\mathcal{X}_{n'} \cap [0, 1]^d; r_{n'})$. This allows to leverage Theorem 1 from [34], which bounds the number of vertices from the unbounded component. In particular, there exists $\beta_0 > 0$ such that

$$\begin{aligned} \Pr[C_\infty \cap H'_n = \Theta(n')] &\geq 1 - \exp\left(-\beta_0^{-(d-1)} n'^{\frac{d-1}{d}}\right) \\ &= 1 - \exp\left(-\beta_0^{-(d-1)} \beta^{d-1} \log n\right) \\ &\geq 1 - \exp(-\log n) = 1 - n^{-1}. \end{aligned}$$

While this is an overkill for our purpose, it does the job in proving that $\Pr[C_\infty \cap H'_n \neq \emptyset] \geq 1 - n^{-1}$. As we assume that \mathfrak{E}_n^1 holds (Lemma 2), it follows that $C_n \cap H'_n \neq \emptyset$ holds with probability at least $1 - n^{-1}$. \square

The following statement allows to bound the graph distance between two connected vertices. We endow every edge of the graph with a length attribute that represents the Euclidean distance between the edges' endpoints. For every two vertices x, x' of \mathcal{G} , $\text{dist}(\mathcal{G}, x, x')$ denotes the length of the shortest (weighted) path on \mathcal{G} between the two vertices.

Theorem 4. [13, Theorem 3] Let $r_n r_n^*$. There exists a constant $\xi \geq 1$, independent of n , such that $\Pr[\mathfrak{E}_n^3] = 1 - O(n^{-1})$, where the event \mathfrak{E}_n^3 is defined as follows: For any two vertices x, x' in the same connected component of $\mathcal{G}(\mathcal{X}_n \cap [0, 1]^d; r_n)$, with $\|x - x'\| = \omega(r_n)$, it holds that $\text{dist}(\mathcal{G}_n, x, x') \leq \xi \|x - x'\|$.

6 Proof of Theorem 1

Proofs for all the three settings of the main theorem are given individually in the subsections below.

6.1 Case i

Here we provide proof for the first (and easy) part of the theorem, which states that PRM fails when $r_n < \gamma^* n^{-1/d}$, even with $r_n^{st} = \infty$. We mention that our proof is a simpler and shorter version of a similar proof that was given in [23] for a slightly different setting.

We show below that in the subcritical regime, i.e., $r_n < r_n^*$, the graph breaks into many small connected components. In particular, the probability of having the largest connected component of size m decays exponentially in m . Denote by $L(\mathcal{G})$ the size of the largest connected component in \mathcal{G} .

Proposition 1. Let $r_n < r_n^*$. Then $L(\mathcal{G}_n) = O(\log n)$ a.a.s.

Proof. First, we start by stating that there exist a constant $\zeta > 0$ and an integer m_0 such that for all $m \geq m_0$ it holds that $\Pr[L(\mathcal{G}_n) \geq m] \leq ne^{-m\zeta}$. This is a simplified version of Proposition 11.2 in [33], which is given for n uniformly sampled points, to the case of PPP, which we have here. In particular, the original proof relies on a relation between these two distributions, which transforms uniform sampling into a PPP and induces an additional factor that is not necessary in our setting. In particular, the factor “ $\exp(-\mu n)$ ” which appears in Equation 11.1 in [33] should be eliminated.

To conclude the proof, we set $m = \alpha \cdot \log n$, where $\alpha > 1/\zeta$, similarly to the proof of Theorem 11.1, Equation 11.4 in [33]. Observe that $ne^{-m\zeta} \rightarrow 0$ as n tends to ∞ . \square

Now, consider the configuration space $C = [0, 1]^2$ depicted in Figure 3 (similar examples can be devised for any $d \geq 2$). The white and blue regions represent the free space. Observe that any path connecting s to t must go through the blue region, whose width is greater than $1/2$. Also note that any point contained in the blue region cannot be connected by an edge from s or t , which deems the large value of r_n^{st} as irrelevant.

Suppose that $\mathcal{G}_n = \mathcal{G}(\mathcal{X}_n, r_n)$ contains a path π_n (black dashed curve) connecting s and t , and denote by $\pi'_n \subset \pi_n$ the subpath that is contained in the blue region. Also, denote by s' and t' the first and last points along π'_n , respectively. Obviously $1/2 \leq \|s' - t'\| \leq c_\ell(\pi_n)$. Then, the number of edges of \mathcal{G}_n , which induce π'_n , is at least $\|s' - t'\|/r_n = \Omega(n^{1/d})$. However, this is in contradiction with the fact that every connected component of \mathcal{G}_n is of logarithmic size in n (Proposition 1).

6.2 Case ii.1

We provide a full proof for the positive setting with length cost. The proof for the bottleneck case, which appears later on, is very similar to the length case.

Suppose that $r_n > r_n^*$, $r_n^{st} = \frac{\beta \log^{1/(d-1)} n}{n^{1/d}}$, $\beta > \beta_0$. For simplicity, we set $r_n = \gamma n^{-1/d}$, where $\gamma > \gamma^*$. By Lemma 1 and Theorem 3, $\mathcal{G}(\mathcal{X}_n; r_n)$ contains a unique infinite connected

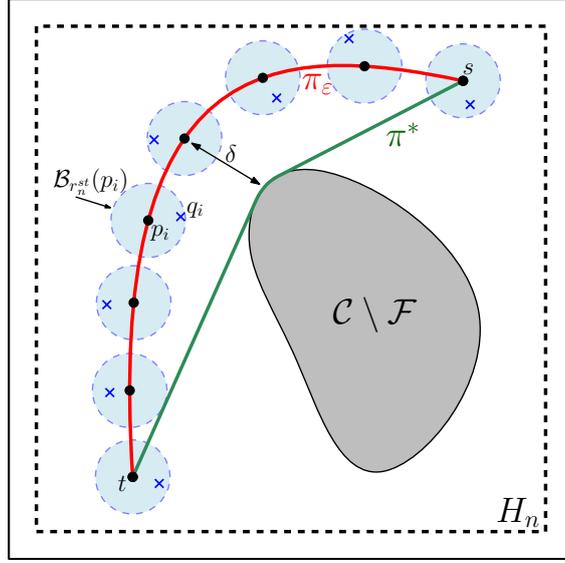


Figure 4: Illustration for the proof of Theorem 1.ii. The outer cube represents the configuration space, whereas the dashed cube is H_n . The gray area represents the forbidden regions. The robust feasible path π^* and π_ε are depicted in green and red, respectively. Observe that every point along π_ε is at least δ away from $\mathcal{C} \setminus \mathcal{F}$. $p_1 = s, p_2, \dots, p_{k-1}, p_k = t$ are depicted as black bullets, where $k = 8$. $\mathcal{B}_{r_n^{st}}(p_i)$ are depicted as blue circles, while the blue cross in each such circle represents q_i .

The next claim states that the length of π_n is a constant factor from the optimum.

Claim 4. *Suppose that $\mathfrak{E}_n^1, \mathfrak{E}_n^3, \mathfrak{E}_n^4$ are satisfied. We have that $c_\ell(\pi_n) \leq (1 + \varepsilon)\xi c_\ell^*$ (with probability 1).*

Proof. First, observe that by the triangle inequality, it holds that

$$\begin{aligned}
& \|q_{i-1} - q_i\| \\
& \leq \|q_{i-1} - p_{i-1}\| + \|p_i - p_{i-1}\| + \|q_i - p_i\| \\
& \leq \|q_{i-1} - p_{i-1}\| + c_\ell(\pi_\varepsilon^{i-1,i}) + \|q_i - p_i\| \\
& \leq 2r_n + c_\ell(\pi_\varepsilon^{i-1,i}) = o(1) + c_\ell(\pi_\varepsilon^{i-1,i}).
\end{aligned}$$

By Theorem 4 and the triangle inequality, it follows that

$$\begin{aligned}
c_\ell(\pi_n) &= \|s - q_1\| + \text{dist}(\mathcal{G}_n, q_1, q_k) + \|t - q_k\| \\
&\leq 2r_n^{st} + \sum_{i=2}^k \text{dist}(\mathcal{G}_n, q_{i-1}, q_i) \\
&\leq o(1) + \sum_{i=2}^k \xi \|q_{i-1} - q_i\| \\
&\leq o(1) + \xi \sum_{i=2}^k c_\ell(\pi_\varepsilon^{i-1,i}) \\
&= o(1) + \xi c_\ell(\pi_\varepsilon) \leq o(1) + (1 + \varepsilon)\xi c_\ell^*.
\end{aligned}$$

We note that in order to eliminate the $o(1)$ summand one can fix ε but work with $\pi_{\varepsilon'}$ instead of π_ε , where $\varepsilon' < \varepsilon$. For simplicity, we chose to leave the proof as is. \square

To conclude, Claim 3 and Claim 4 show the existence of a (collision free) path π_n in \mathcal{P}_n , whose length is at most $(1 + \varepsilon)\xi c_\ell^*$. It remains to bound the probability that $\mathfrak{E}_n^1, \mathfrak{E}_n^3, \mathfrak{E}_n^4$ hold simultaneously:

$$\begin{aligned} \Pr[\mathfrak{E}_n^1 \wedge \mathfrak{E}_n^3 \wedge \mathfrak{E}_n^4] &= \Pr[\mathfrak{E}_n^3 \wedge \mathfrak{E}_n^4 | \mathfrak{E}_n^1] \cdot \Pr[\mathfrak{E}_n^1] \\ &= \left(1 - \Pr[\overline{\mathfrak{E}_n^3 \wedge \mathfrak{E}_n^4} | \mathfrak{E}_n^1]\right) \cdot \Pr[\mathfrak{E}_n^1] \\ &= \left(1 - \Pr[\overline{\mathfrak{E}_n^3} \vee \overline{\mathfrak{E}_n^4} | \mathfrak{E}_n^1]\right) \cdot \Pr[\mathfrak{E}_n^1] \\ &\geq \left(1 - \Pr[\overline{\mathfrak{E}_n^3} | \mathfrak{E}_n^1] - \Pr[\overline{\mathfrak{E}_n^4} | \mathfrak{E}_n^1]\right) \cdot \Pr[\mathfrak{E}_n^1] \end{aligned}$$

By Claim 2 and Lemma 3, $\Pr[\overline{\mathfrak{E}_n^4} | \mathfrak{E}_n^1] \leq k \Pr[\overline{\mathfrak{E}_n^2} | \mathfrak{E}_n^1] \leq kn^{-1}$. Also note that $\Pr[\mathfrak{E}_n^3 | \mathfrak{E}_n^1] \geq \Pr[\mathfrak{E}_n^3]$, since $\Pr[\mathfrak{E}_n^3 | \overline{\mathfrak{E}_n^1}] = 0$. Therefore,

$$\begin{aligned} \Pr[\mathfrak{E}_n^1 \wedge \mathfrak{E}_n^3 \wedge \mathfrak{E}_n^4] &\geq \left(1 - \Pr[\overline{\mathfrak{E}_n^3} | \mathfrak{E}_n^1] - \Pr[\overline{\mathfrak{E}_n^4} | \mathfrak{E}_n^1]\right) \cdot \Pr[\mathfrak{E}_n^1] \\ &\geq (1 - O(n^{-1}) - kn^{-1}) \left(1 - \exp(-\alpha n^{1/d} \log^{-1} n)\right) \\ &= 1 - O(n^{-1}). \end{aligned}$$

6.3 Case ii.2

We prove the asymptotic optimality of PRM for bottleneck cost, with appropriate values of r_n, r_n^{st} . The proof follows very similar lines to that of the length cost. Fix $\varepsilon > 0$. For simplicity we assume that $\varepsilon < 1$ (the proof can be adapted to larger values of ε). By definition, there exists an \mathcal{M} -robust path $\pi' \in \Pi_{s,t}^{\mathcal{F}}$ and $\delta, \delta' > 0$ such that

- (a) $\mathcal{B}_\delta(\pi') \subset \mathcal{F}$,
- (b) $c_b(\pi', \mathcal{M}) \leq (1 + \varepsilon/3)c_b^*$,
- (c) and $\forall x \in \mathcal{B}_{\delta'}(\pi') : \mathcal{M}(x) \leq (1 + \varepsilon/3)c_b(\pi', \mathcal{M})$.

Now, define $\delta^* = \min\{\delta, \delta'\}$. By substituting δ with δ^* , and π_ε with π' in the proof of Theorem 1.ii, it follows that if $\mathfrak{E}_n^1, \mathfrak{E}_n^3, \mathfrak{E}_n^4$ are satisfied, \mathcal{P}_n contains a path $\pi'_n \in \Pi_{s,t}^{\mathcal{F}}$ such that $\text{Im}(\pi'_n) \subset \mathcal{B}_{\delta'}(\pi')$. This implies that

$$\begin{aligned} c_b(\pi'_n, \mathcal{M}) &\leq (1 + \varepsilon/3)c_b(\pi', \mathcal{M}) \\ &\leq (1 + \varepsilon/3)^2 c_b^* \leq (1 + \varepsilon)c_b^*, \end{aligned}$$

which concludes this proof.

7 Analysis of other planners

In this section we describe the implications of Theorem 1 to additional planners, which are closely related to PRM. In all the results below, the values $r_n^*, \gamma^*, \beta, \xi$ are identical to those in Theorem 1. Furthermore, for simplicity of presentation, we assume that (\mathcal{F}, s, t) is robustly feasible, and \mathcal{M} is well behaved, where relevant.

We first consider FMT* and BTT, and then proceed to RRG, whose analysis is more involved. In all the following algorithms we assume that the sampling method is a PPP.

7.1 The FMT* and the BTT planners

The following two corollaries are direct consequences of Theorem 1 as the two algorithms FMT* and BTT traverse an implicitly-represented PRM graph while minimizing the cost functions c_ℓ and c_b , respectively.

Corollary 1. *If $r_n < r_n^*$ and $r_n^{st} = \infty$ then FMT* fails a.a.s. If $r_n > r_n^*$ and $r_n^{st} = \frac{\beta \log^{1/(d-1)} n}{n^{1/d}}$, then for any $\varepsilon > 0$ FMT* returns a path $\pi_n \in \Pi_{s,t}^{\mathcal{F}}$ with $c_\ell(\pi_n) \leq (1 + \varepsilon)\xi c_\ell^*$ with probability $1 - O(n^{-1})$.*

We remark that FMT* can ignore certain edges of the underlying PRM graph, which may result in a path of lower quality than that obtained by PRM. However, this situation only occurs for vertices which lie within distance of r_n from $C \setminus \mathcal{F}$ (see [21, Remark 3.3]). As our analysis of PRM in the previous section utilizes only vertices and edges that are far from obstacles our proofs extend to FMT* straightforwardly.

Corollary 2. *If $r_n < r_n^*$ and $r_n^{st} = \infty$ then BTT fails a.a.s. If $r_n > r_n^*$ and $r_n^{st} = \frac{\beta \log^{1/(d-1)} n}{n^{1/d}}$, then for any $\varepsilon > 0$ BTT returns a path $\pi'_n \in \Pi_{s,t}^{\mathcal{F}}$ with $c_b(\pi'_n, \mathcal{M}) \leq (1 + \varepsilon)c_b^*$ with probability $1 - O(n^{-1})$.*

7.2 The RRG planner

We consider the incremental planner RRG, which can be viewed as a cross between RRT and PRM. Due to this relation we can extend the analysis of PRM to RRG. (The term ‘‘incremental’’ refers to planners that generate samples one after the other, and connect the current samples to previous ones.)

We introduce an incremental version of a PPP to extend our theory to RRG.

Claim 5. *Let N be a Poisson random variable with mean 1. At each iteration i draw a sample $N_i \in N$ and define $\mathcal{X}_i = \{X_1, \dots, X_{N_i}\}$ to be N_i points chosen independently and uniformly at random from $[0, 1]^d$. The set $\mathcal{X} = \bigcup_{1 \leq i \leq n} \mathcal{X}_i$, obtained after n iterations, is a PPP of density $\lambda = n$.*

Note that the sum of n i.i.d. Poisson random variables with means $\lambda_1, \lambda_2, \dots, \lambda_n$ is a Poisson random variable with mean $\lambda = \sum_{i=1}^n \lambda_i$. Claim 5 follows directly from this property.

We now describe an adaptation of RRG for an incremental PPP. Given a start configuration $s \in \mathcal{F}$ and a goal region $\mathcal{X}_{\text{goal}} \subseteq \mathcal{F}$, RRG initializes a roadmap with a single node s . Let η denote the constant used by RRG for local steering (see [23]). Similarly to PRM, RRG employs the connection radii $r_n^{\text{RRG}}, r_n^{s \text{RRG}}$.

Let N be a Poisson random variable with mean 1 and let n_{i-1} denote the number of nodes of the constructed roadmap after $i - 1$ iterations. At the i th iteration we draw a sample $N_i \in N$ and define $\mathcal{X}_i = \{X_1, \dots, X_{N_i}\}$ to be N_i points chosen independently and uniformly at random from $[0, 1]^d$. RRG will iteratively process the samples X_1, \dots, X_{N_i} , as follows: For a sample $X_k \in \mathcal{X}_i$ RRG will first locate the nearest node x_{near} to X_k . Let x_{new} be a node at distance at most η from x_{near} at the direction of X_k . RRG will attempt to connect x_{new} to x_{near} . If the connection attempt is successful, the new node x_{new} will be added to the roadmap. Then, RRG will attempt to connect x_{new} to all the existing nodes within a ball of radius $\min\{r_n^{\text{RRG}}, \eta\}$. It will additionally attempt to connect x_{new} to s if $\|x_{\text{new}} - s\| \leq \min\{r_n^{s \text{RRG}}, \mu\}$. Note that for the i th iteration of RRG we set $n = n_{i-1}$.

The connection radii $r_n^{\text{RRG}}, r_n^{s \text{RRG}}$ that will be used during the i th iteration are fixed, as both are functions of the number of nodes in the roadmap at the beginning of the iteration. Both radii will

be set to decrease with the number of roadmap nodes, as in PRM. However, since we fix the radius at each iteration of RRG we use a slightly larger radius than the one obtained had we considered the current roadmap size. This will clearly keep all relevant connections and perhaps even add more.

Theorem 5. *Suppose that $(\mathcal{F}, s, \mathcal{X}_{goal})$ is robustly feasible (we extend Definition 3 to describe a robustly feasible path $(\mathcal{F}, s, \mathcal{X}_{goal})$ for RRG). Let \mathcal{R}_n denote the roadmap constructed by RRG after n iterations with*

$$r_n^{\text{RRG}} = (1 + \mu)r_n, \quad r_n^{\text{sRRG}} = (1 + \mu)r_n^{\text{st}},$$

where r_n, r_n^{st} are as in Theorem 1.ii, and μ is any positive constant. Then for any $\varepsilon > 0$ the following holds with probability $1 - O(n^{-1})$:

1. \mathcal{R}_n contains a path $\pi_n \in \Pi_{s, \mathcal{X}_{goal}}^{\mathcal{F}}$ with $c_\ell(\pi_n) \leq (1 + \varepsilon)\xi c_\ell^*$;
2. If \mathcal{M} is well behaved then, \mathcal{R}_n contains a path $\pi'_n \in \Pi_{s, \mathcal{X}_{goal}}^{\mathcal{F}}$ with $c_b(\pi'_n, \mathcal{M}) \leq (1 + \varepsilon)c_b^*$.

7.2.1 Proof of Theorem 5

We provide proof only for c_ℓ , as the case for c_b is almost identical. In what follows we consider a robust path $\pi_\varepsilon \in \Pi_{s, \mathcal{X}_{goal}}^{\mathcal{F}}$ for which $c_\ell(\pi_\varepsilon) \leq (1 + \varepsilon)c_\ell^*$. Let $\delta > 0$ denote the clearance of π_ε , and L denote its length. Finally, set $\kappa = \min(\delta, \eta)/5$.

Our proof uses Theorem 1 as a central ingredient, due to the following observation, which is formalized below: For n large enough, with probability $1 - O(e^{-n})$, \mathcal{R}_n contains a PRM roadmap $\mathcal{P}_{n'}$ in the vicinity of π_ε , where n' is slightly smaller than n . This in turn follows from the unique structure of RRG, which can be viewed as a combination of PRM and RRT, as RRG supplements the RRT tree with additional edges.

In particular, the following corollary can be deduced from the probabilistic completeness of RRT with samples from a PPP (see Appendix B), which does not rely on the parameters $r_n^{\text{RRG}}, r_n^{\text{sRRG}}$.

Claim 6. *It is possible to tile π_ε with $m = L/\kappa$ balls of radius κ , such that after n iterations of RRG, with probability at least $1 - a \cdot e^{-bn}$ for some constants $a, b \in \mathbb{R}_{>0}$, every ball will contain at least one RRG vertex.*

The following lemma formalizes the connection between RRG and PRM.

Lemma 4. *Fix n and define $n' = (1 + \mu)^{-d}n$. Let \mathcal{R}_n be an RRG graph constructed after n iterations with connection radii $r_n^{\text{RRG}}, r_n^{\text{sRRG}}$. With probability at least $1 - a \cdot e^{-b'n}$, for some constants $a, b' \in \mathbb{R}_{>0}$, the graph $\mathcal{R}_n \cap \mathcal{B}_{\kappa/2}(\pi)$ contains a PRM graph $\mathcal{P}_{n'} \cap \mathcal{B}_{\kappa/2}(\pi)$, constructed with the radius $r_{n'}$.*

Proof. Consider the $n - n' = n(1 - (1 + \mu)^{-d})$ first iterations of RRG. From Claim 6 we obtain that given a robust path π_ε , it can be tiled using a constant number of balls of radius κ such that with probability of at least $1 - ae^{-b(n-n')} = 1 - ae^{-b'n}$ every ball will contain an RRG vertex after n iterations of RRG, for some constants $a, b \in \mathbb{R}_{>0}$, and $b' = b(1 - (1 + \mu)^{-d})$. Now, let us consider a new sample $x \in \mathcal{B}_{\kappa/2}(\pi)$ that is added in iteration $n'' > n - n'$. By Claim 6, $\mathcal{R}_{n''}$ includes (with certain probability) a node v that can be connected to x ; that is, both $\|x - v\| \leq \eta$ and the straight-line path from x to v is collision free. Thus, v is added as a node to RRG and then connected in a PRM-fashion to all its neighbors within a radius of $r_{n''}^{\text{RRG}}$.

Note that the smallest value of the latter radius is when $n'' = n$. Observe that

$$\begin{aligned} r_n^{\text{RRG}} &= (1 + \mu)r_n = \gamma \left((1 + \mu)^{-d}n \right)^{-1/d} \\ &= \gamma(n')^{-1/d} = r_{n'}. \end{aligned}$$

Additionally,

$$\begin{aligned} r_n^{s\text{RRG}} &= (1 + \mu)r_n^{st} \\ &= (1 + \mu) \frac{\beta \log^{1/(d-1)} n}{n^{1/d}} \\ &= \frac{\beta \log^{1/(d-1)} n}{((1 + \mu)^{-d}n)^{1/d}} \\ &\geq \frac{\beta \log^{1/(d-1)} ((1 + \mu)^{-d}n)}{((1 + \mu)^{-d}n)^{1/d}} = r_{n'}^{s\text{RRG}}. \end{aligned}$$

This implies that all samples added after iteration $n - n'$ are connected in a PRM-fashion with a radius $r_{n'}$. Thus, $\mathcal{R}_n \cap \mathcal{B}_{\kappa/2}(\pi)$ contains a PRM graph $\mathcal{P}_{n'} \cap \mathcal{B}_{\kappa/2}(\pi)$, constructed with the radius $r_{n'}$ with probability $1 - a \cdot e^{-b'n}$, as required. \square

We are now ready to finish the proof of Theorem 5. Let us denote by A the event that the graph \mathcal{R}_n contains a PRM graph $\mathcal{P}_{n'} \cap \mathcal{B}_{\kappa/2}(\pi_\varepsilon)$. Given that A occurs, we denote by B the event that the constructed roadmap $\mathcal{R}_n \cap \mathcal{B}_{\kappa/2}(\pi_\varepsilon)$ maintains the good properties of PRM, as defined in Theorem 1.ii.1. That is, $\mathcal{P}_{n'} \cap \mathcal{B}_{\kappa/2}(\pi_\varepsilon)$ contains a path $\pi_{n'} \in \Pi_{s, \mathcal{X}_{\text{goal}}}^{\mathcal{F}}$ with $c_\ell(\pi_{n'}) \leq (1 + \varepsilon)\xi \cdot c_\ell^*$;

Given that A occurs, from Theorem 1 we obtain that B occurs on the PRM-like graph inside $\mathcal{B}_{\kappa/2}(\pi)$ with probability of at least

$$\begin{aligned} 1 - O((n')^{-1}) &= 1 - O\left((1 + \mu)^{-d}n^{-1}\right) \\ &= 1 - O(n^{-1}). \end{aligned}$$

That is, $\Pr[B|A] \geq 1 - O(n^{-1})$. Additionally, Lemma 4 states that $\Pr[A] \geq 1 - ae^{-b'n}$ for some b' . Note that we would like to bound the probability that both A and B occur simultaneously. From the definition of conditional probability we have that $\Pr[A \cap B] = \Pr[B|A] \cdot \Pr[A]$. Therefore,

$$\begin{aligned} \Pr[A \cap B] &\geq (1 - O(n^{-1})) \cdot (1 - ae^{-b'n}) \\ &\geq 1 - O(e^{-b'n}) - O(n^{-1}) \\ &= 1 - O(n^{-1}). \end{aligned}$$

7.3 Multi-robot planners

In this section we briefly state the implications of our analysis of PRM for sampling-based multi-robot motion planning. In particular, we consider the planners M^* [51], dRRT [46] and dRRT* [10].

The multi-robot setting involves $m \geq 2$ robots operating in a shared workspace. In this more challenging setting collisions between different robots must be avoided, in addition to standard robot-obstacle collisions. Single-robot sampling-based planners can be applied directly to the problem by considering the robot fleet as one highly-complex robot. However, recently-introduced planners that are tailored for the multi-robot case have proved to be much more effective than the aforementioned naive approach.

Such recent approaches include M^* , $dRRT$, and $dRRT^*$. A common ground to these techniques is an implicit construction of a composite roadmap \mathbb{G} [46], which results from a tensor product between m single-robot roadmaps G_1, \dots, G_m —for every $1 \leq i \leq m$ the graph G_i is a PRM embedded in the configuration space of robot i . It was recently proved that if G_1, \dots, G_m are AO, with respect to the individual robots for which they are defined, then \mathbb{G} is AO with respect to the multi-robot problem [10].

Thus, by Theorem 1 it follows that the underlying single-robot PRM graphs G_1, \dots, G_m can be constructed using a smaller connection radius, while still guaranteeing the AO (or AnO) of \mathbb{G} . This in turn, guarantees AO (or AnO) of M^* , $dRRT^*$, and PC of $dRRT$.

8 Experimental results

We present experiments demonstrating the effect of using different values of the connection radius r_n on the running time, cost of solution for the length cost c_ℓ , and success rate, when running the algorithms PRM and FMT* on problems of dimensions up to 12. In particular, r_n ranges between the critical radius (Theorem 1) and previously obtained upper bounds from [21, 23].

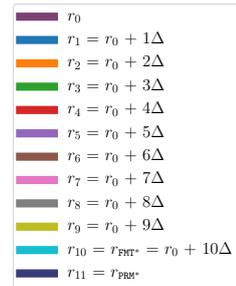
We validate our theory for PRM (Theorem 1) and FMT* (Corollary 1). We observe that smaller connection radii, than previously-obtained bounds, still allow the planners to converge to high-quality, near-optimal paths. Furthermore, we identify situations in which using a radius that is close to r_n^* allows to obtain a high-quality solution more quickly. Moreover, although the resulting cost for the smaller radii can be slightly worse, we observe that postprocessing the paths using standard simplification methods yields solutions that are only marginally inferior to the best (postprocessed) solution. Specifically, in harder scenarios the advantage in using a smaller connection radius is more prominent; in some cases we obtain a reduction of 50% in running time, with an improved cost, and similar success rates when compared to the results obtained using the original FMT* connection radius.

8.1 Implementation details

In our experiments, we used the Open Motion Planning Library (OMPL 1.3, [49]) on a 2.6GHz \times 2 Intel Core i5 processor with 16GB of memory. Results were averaged over 50 runs and computed for dimensions up to 12.

The planners that we used are PRM and the batch variant of FMT*, which were adapted for samples from a PPP, where n is the expected number of samples. Specifically, given the expected number of samples n , these variants generate a set of samples according to the recipe in Claim 1. The two planners use the connection radii r_n, r_n^{st} . Note that r_n^{st} should be at least $\frac{\beta \log^{1/(d-1)} n}{n^{1/d}}$, but the exact value of β is unknown. For simplicity, we set r_n^{st} to be identical to r_{PRM^*} , defined in [23]. We emphasize that although we use an asymptotically smaller value, it still yields (empirically) convergence in cost. This suggests that the bound on r_n^{st} can be further reduced.

Given a scenario and a value n , we define a set of $k + 1$ increasing connection radii, $\{r_0, \dots, r_k\}$, as follows. We set the minimal connection radius to be $r_0 = \gamma n^{-1/d}$, where $\gamma = 1$. Note that γ is larger than γ^* by a factor of roughly 2. The maximal connection radius, denoted by $r_k = r_{FMT^*}$, is as defined in [21]. For each $1 \leq i \leq k - 1$ we define $r_i = r_0 + i \cdot \Delta$, where $\Delta = (r_k - r_0)/k$. Now, for every scenario and number of samples n we run our planning algorithm with r_n^{st} , and $r_n \in \{r_0, \dots, r_k\}$. Note that all our plots are for $k = 10$, and that in some experiments an additional radius $r_{k+1} = r_{PRM^*} > r_{FMT^*}$ appears as well (see [23]). The figure to the right depicts the colors and labeling that will be used throughout this section.



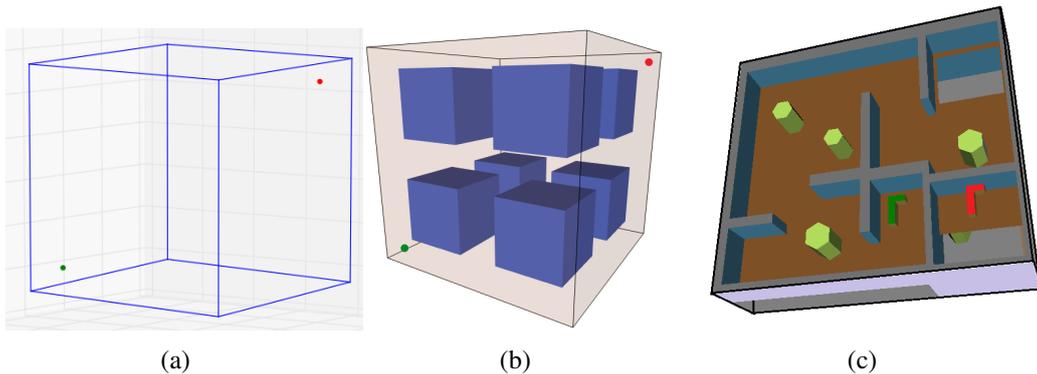


Figure 5: Scenarios used in experiments. (a) d D Hypercube, (b) d D Hypercube with 2^d hypercubical obstacles, and (c) 3D Cubicles. Start and target configurations for a robot are depicted in green and red, respectively. Scenario (c) is provided with the OMPL distribution.

8.2 Results

Euclidean space. The scenario we consider (see Figure 5a) consists of a point robot moving in the obstacle-free unit d -dimensional hypercube. Therefore, $\mathcal{F} = \mathcal{C} = [0, 1]^d$. We set the start and target positions of the robot to be $s = (0.1, \dots, 0.1)$ and $t = (0.9, \dots, 0.9)$, respectively.

We run PRM and plot (i) the overall running time, (ii) the normalized cost (c_ℓ) of the obtained solution, where a value of 1 represents the best possible cost, and (iii) the portion of successful runs—all as a function of the expected number of samples n . Results are depicted in Figure 6.

The plots demonstrate the following trend: for each radius r the cost obtained by PRM converges to some constant times the optimal cost, which is marked with the dashed red curve in the “Cost vs. n ” plot. (We note that it is possible that for larger values of n the cost values for different radii will eventually converge to the same value.) Clearly, r_{PRM^*} yields the best cost but at the price of increased running time. $r_{10} = r_{\text{FMT}^*}$ obtains the next best cost, with improved running time, and so on. Note that for $d = 4$, already for $n = 1\text{K}$ a solution is found for all radii except r_0 , whereas for $d = 8$ and $n = 5\text{K}$ a solution is found for all radii above r_4 . It is important to note that there is a clear speedup in the running times of PRM when using r_i for $i \leq 9$, over $r_{\text{FMT}^*}, r_{\text{PRM}^*}$, with a slight penalty (a factor of roughly 2) in the resulting costs.

For the first set of experiments of PRM in an obstacle free d -dimensional hypercube, we also study how the value of r_i affects the size of the connected components. Note that Theorem 1 in [34] states that $|C_n| = \Theta(n)$. In particular, we measure the size of the two largest connected components, denoted by C_n and C'_n , as a function of both the radius r_i and the expected number of samples n . The results for $d = 2, d = 12$ are summarized in Table 1. Already for r_2 , C_n is significantly larger than C'_n . However, for r_0 there is no clear difference. That is, we do not see in practice the expected emergence of the “huge” component. As n increases, the proportion of C_n increases as well, whereas that of C'_n shrinks. Also note that for specific r_i, n the maximal component is smaller for $d = 12$ than for $d = 2$.

General Euclidean space. We consider the following d -dimensional scenario (based on a scenario from [45]), for $d \in \{4, 8\}$, depicted in Figure 5b in 3D, and a point robot: $\mathcal{C} = [0, 1]^d$ is subdivided into 2^d sub-cubes by halving it along each axis. Each sub-cube contains a centered d -dimensional axis-aligned hypercubical obstacle that covers 25% of the sub-cube. The start position s of the robot is placed on the diagonal between $(0, \dots, 0)$ and $(1, \dots, 1)$, such that it is equidistant from the origin and from the closest hypercubical obstacle. t is selected similarly, with respect to $(1, \dots, 1)$.

d	n	r_0		r_2		r_{10}	
		$ C_n /n$	$ C'_n /n$	$ C_n /n$	$ C'_n /n$	$ C_n /n$	$ C'_n /n$
2	1K	0.17	0.1	0.88	0.05	1.0	0.0
2	5K	0.08	0.05	0.97	0.001	0.999	0.0
2	10K	0.05	0.04	0.98	0.003	1.0	0.0
2	50K	0.02	0.01	0.99	0.001	1.0	0.0
12	1K	0.04	0.02	0.84	0.003	1.0	0.0
12	5K	0.12	0.01	0.91	0.001	1.0	0.0
12	10K	0.21	0.005	0.94	0.001	1.0	0.0

Table 1: The average proportion of the two largest connected components C_n, C'_n obtained using PRM as a function of r_i, n, d .

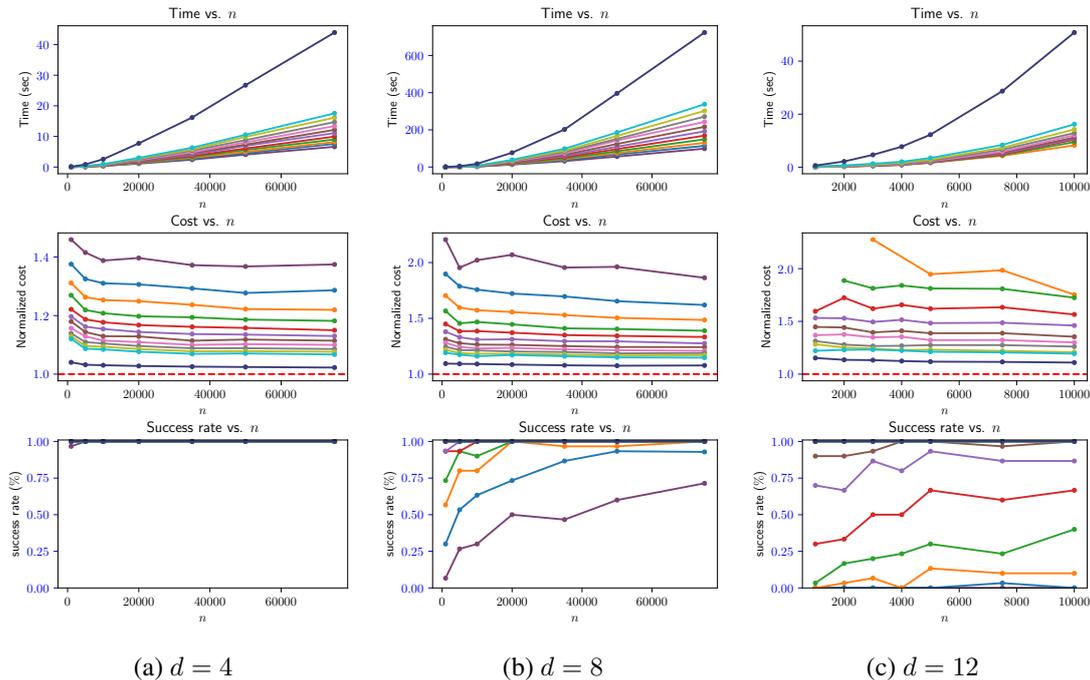


Figure 6: PRM in the (a) 4D, (b) 8D, and 12D obstacle free hypercube.

Here we use FMT^* with radii ranging from r_0 to r_{10} . We also ran PRM which exhibited a similar behavior. Figure 7a presents the results for $d = 4$. We plot the average cost after simplifying the resulting paths in addition to the average original cost. We mention that there is a difference in cost between the various radii and that costs obtained using larger radii are often better. However, after applying OMPL’s default path-simplification procedure we obtain paths with negligible differences in cost. This suggests that paths obtained using the smaller connection radii are of the same homotopy class as the path obtained using r_{FMT^*} .

For $d = 8$ (Figure 7b), as in the case of $d = 4$, costs obtained using larger radii are often better. However, applying OMPL’s default path-simplification procedure yields paths with negligible differences in cost. We do mention, though, that the success rates of all radii deteriorate, when compared to $d = 4$.

General non-Euclidean space. We use the Cubicles scenario (see Figure 5c) provided with the OMPL distribution [49]. Here the goal is to find a collision-free path for one or two L-shaped robots that need to exchange their positions. In the single-robot case, the robot needs to

move between the two configurations depicted in red and green. In the two-robot case, these two configurations represent the start positions of the robots, where the goal is for each robot to finish at the start of the other robot. Since the robots are allowed to translate and rotate, then $d = 6$ or $d = 12$, respectively, and \mathcal{C} is non-Euclidean. Although our theory does not directly apply here, we chose to test it empirically for such a setting.

We ran our experiments with FMT^* using a connection radii ranging from r_0 to $r_{10} = r_{\text{FMT}^*}$. Initially, no solution was found in all runs. Indeed, as was mentioned in [26], this is not surprising, since the theory from which the radius values are derived assumes that the configuration space is Euclidean. In the same paper the authors propose a heuristic for effectively using radius-based connections in motion planning. In our experiments (Figure 8b) we increased all radii by a multiplicative factor of 3 and 10, respectively, in order to increase the success rates. We mention that this yielded similar behavior to that when using the connection scheme suggested in [26].

Results for the single-robot case are depicted in 8a. In this case, even with 25K samples already with r_3 (green) we were able to obtain high success rates. The running time was slightly better than that of r_{10} , whereas the cost was higher than the one obtained using r_{10} . However, after applying the smoothing procedure the difference in cost decreased significantly.

Let us proceed to the two-robot setting. Observe in Figure 8b, that the variance in cost after path simplification is again significantly smaller than the variance in the original cost. Clearly, smaller radii exhibit shorter running times, but with smaller success rates. However, since the success rate improves as the number of samples n increases, one could use the smaller radii with larger n and still obtain a solution with comparable cost and success rate in shorter running time. Indeed, using r_3 (green) with 75K samples we obtain a solution whose cost (after simplification) is slightly better than the cost obtained using r_{10} with 42K samples (after simplification). Moreover, the running time of the former is roughly 50% of the time taken for the latter, and both obtain similar success rates. This indicates that in such settings, one could benefit from using smaller radii in terms of favorable running times and obtain a comparable or even better cost with similar success rates.

9 Future work

In this work we leveraged techniques from percolation theory to develop stronger analysis of PRM-based sampling-based motion planners. In the hope of providing mathematically rigorous presentation, while still being accessible, we chose to focus the discussion on simplified, possibly unrealistic, robotic systems. In particular, we assumed a holonomic system having a Euclidean configuration space.

Our immediate future goal is to extend our theory to non-Euclidean spaces, such as those arising from rigid bodies translating and rotating in space. The next challenge will be to extend the model to robots with differential constraints. While the latter task seems daunting, we should keep in mind that the state space of such systems can be modeled as a differential manifold. This may allow to locally apply our Euclidean-space techniques to analyze manifold spaces. Indeed, a similar approach has already been considered in previous work [39, 40].

The next algorithmic challenge is to consider robotic systems for which precise *steering*, i.e., solving the *two-point boundary value problem*, cannot be performed, at least not efficiently (see discussion in [30, Section 1.3]). In such cases, PRM-based planners (as we considered here), or RRT*-based techniques, cannot be applied. We pose the following question: Is it possible to extend existing planners to work in the absence of a steering function, while maintaining their theoretical properties? We plan to tackle this question in the near future.

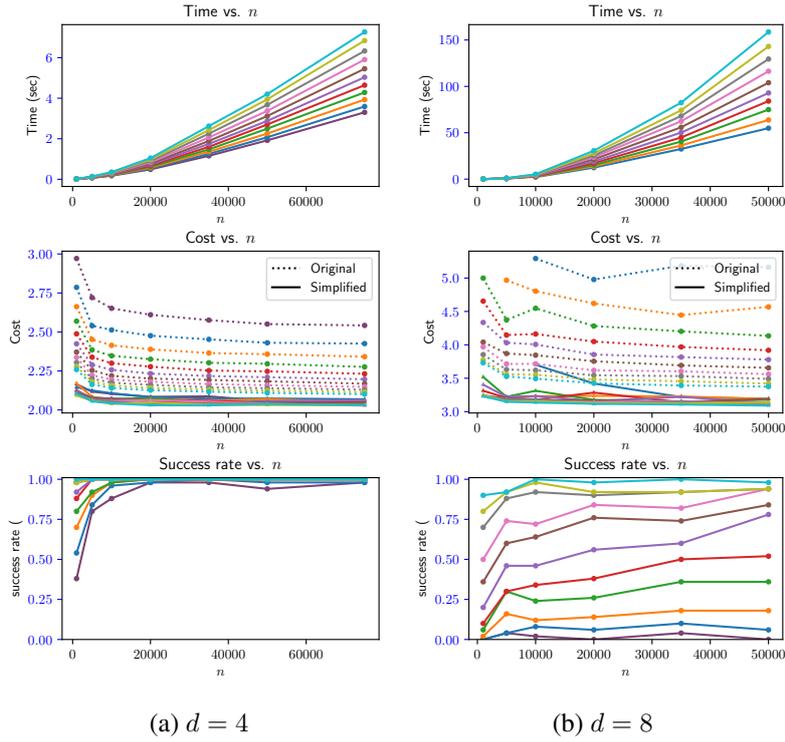


Figure 7: FMT* for a point robot in the (a) 4D and (b) 8D hypercube with obstacles scenario (Figure 5b). Both the average original cost and the average cost after simplification are presented for each radius. There is a difference in cost between the various radii (dashed lines), that diminishes after simplifying the resulting paths (solid lines).

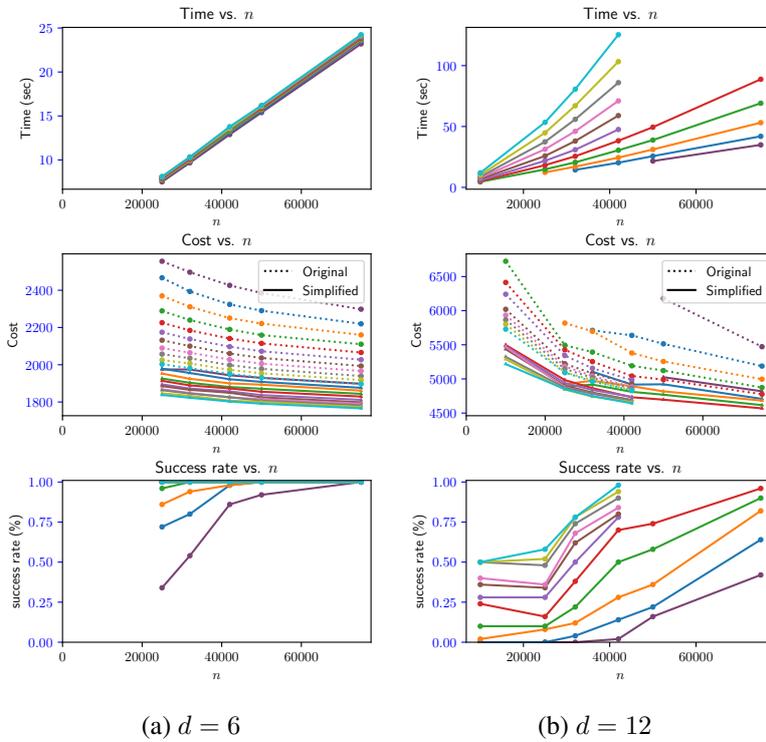


Figure 8: FMT* for (a) one and (b) two rigid-body robot(s) moving in OMPL's Cubicles scenario (Figure 5c). Both the average original cost and the average cost after simplification are presented for each radius.

10 Acknowledgments

We thank the participants of the *Workshop on Random Geometric Graphs and Their Applications to Complex Networks*, 2015, and in particular Carl P. Dettmann, Mathew Penrose, Paul Balister, and Yuval Peres, for fruitful discussions. We also thank Dan Halperin for commenting on early drafts of the paper.

This work has been supported in part by the Israel Science Foundation (grant no. 825/15), by the Blavatnik Computer Science Research Fund, and by the Blavatnik Interdisciplinary Cyber Research Center at Tel Aviv University. M.K. is supported in part by the Yandex Foundation. This work was prepared as part of K.S.'s Ph.D. thesis in Tel Aviv University, with the support of the Clore Israel Foundation.

Appendix

A Theory for (semi-)deterministic sampling

So far we have studied the theoretical properties of various sampling-based planners constructed with the PPP sample set \mathcal{X}_n . In this section we study similar aspects but for a different sampling regime. Particularly, we replace \mathcal{X}_n with \mathcal{L}_n which represents a rectilinear lattice, also known as a Sukharev sequence. A recent paper by [22] studies the theoretical properties of PRM and FMT^* when constructed with \mathcal{L}_n (and other deterministic low-dispersion sequences). Here we introduce a new analysis which shows that PRM constructed with \mathcal{L}_n can be further sparsified while maintaining AO. This result can also be extended to FMT^* , BTT , RRG and other planners, as was done in Section 7 for the setting of PPP.

A.1 The basics

We start with some basic definitions. For a given n define $\tilde{r}_n^* = n^{-1/d}$ and $\mathcal{L}_n = \tilde{r}_n^* \cdot \mathbb{Z}_d$. That is, \mathcal{L}_n contains all the points of a rectangular lattice, with edge length \tilde{r}_n^* . Observe that $\mathcal{G}(\mathcal{L}_n; \tilde{r}_n^*)$ is a deterministic graph which connects by an edge every two points $x, x' \in \mathcal{L}_n$ that differ in exactly one coordinate and the distance between them is exactly \tilde{r}_n^* . We proceed to define a randomized subset of $\mathcal{G}(\mathcal{L}_n; \tilde{r}_n^*)$:

Definition 11. Given $0 \leq p \leq 1$, the graph $\mathcal{G}_n^p = \mathcal{G}(\mathcal{L}_n; \tilde{r}_n^*; p)$ is defined in the following manner. Every point of \mathcal{L}_n is added independently with probability p to \mathcal{G}_n^p as a vertex. Two vertices $x, x' \in \mathcal{G}_n^p$ are connected by an edge if $\|x - x'\| \leq \tilde{r}_n^*$.

This corresponds to a well-studied model of *site percolation*. Similarly to the continuous model studied in the previous section, we say that \mathcal{G}_n^p *percolates* if it contains an infinite connected component C_∞^p that includes the origin o . Also denote by C_n^p the largest connected component of $[0, 1]^d \cap C_\infty^p$.

The following theorem corresponds to a combination of Theorem 2, Lemma 1, and Theorem 3 for the continuous case.

Theorem 6. [11] *There exists a critical value $0 < p^* < 1$ such that for $p < p^*$ the graph \mathcal{G}_n^p percolates with probability zero. If $p > p^*$ then \mathcal{G}_n^p percolates with non-zero probability. Furthermore, if $p > p^*$ then C_∞^p is unique with probability 1.*

The next two lemmata serve as discrete versions of Lemma 2 and Lemma 3, respectively. Their proofs are very similar to the original proofs and are therefore omitted. We do mention that they rely on Theorem 5 and Theorem 4 of [34], respectively.

Lemma 5. Suppose that $p > p^*$. Denote by \mathfrak{E}_n^1 the event that $C_\infty^p \cap H_n \subset C_n^p$, where H_n is as defined in Lemma 2. Then there exists $n_0 \in \mathbb{N}$ and $\alpha > 0$ such that for any $n > n_0$ it holds that

$$\Pr[\mathfrak{E}_n^1] \geq 1 - \exp\left(-\alpha n^{1/d} \log^{-1} n\right).$$

Lemma 6. Suppose that $p > p^*$. Define $H_n'' \subset H_n$ to be a hypercube of side length

$$h_n'' = \frac{\beta' \log^{1/(d-1)} n}{n^{1/d}}.$$

Denote by \mathfrak{E}_n^2 the event that $H_n'' \cap C_n^p \neq \emptyset$. Then there exist $n_0 \in \mathbb{N}$ and $\beta'_0 > 0$ such that for any $n > n_0$, $\beta' > \beta'_0$ it holds that $\Pr[\mathfrak{E}_n^2 | \mathfrak{E}_n^1] \geq 1 - n^{-1}$.

Now we proceed to the discrete counterpart of Theorem 4. For any $v \in \mathbb{R}^d$ define $x(v) = \operatorname{argmin}_{x \in \mathcal{L}_n} \|x - v\|$. The following is a simplified version of [3, Theorem 1.1].

Theorem 7. Suppose that $p > p^*$. There exists a positive constant $\xi' \geq 1$ such that for any two fixed points $v, v' \in \mathbb{R}^d$ it holds that

$$\begin{aligned} \Pr[x(v), x(v') \in \mathcal{C}_\infty, \operatorname{dist}(\mathcal{G}_n^p, x(v), x(v')) \\ > \xi' \|x(v) - x(v')\|] = o(1). \end{aligned}$$

For the subcritical regime ($p < p^*$), we have the following theorem. Define

$$S_k = \left\{ x \in \mathcal{L}_n \mid \frac{\|x\|_1}{\tilde{r}_n^*} \leq k \right\}.$$

That is, S_k represents all the points in \mathcal{L}_n that can be reached from o by at most k edges of length \tilde{r}_n^* over the corresponding rectilinear lattice. Denote by ∂S_k the boundary of S_k , i.e., all the points that are exactly k edges from o . Let A_k be the event that \mathcal{G}_n^p contains a path from o to some vertex in S_k .

Theorem 8. ([16, Theorem 5.4]) If $p < p^*$, there exists a constant $c > 0$ such that $\Pr[A_k] < e^{-ck}$ for all k .

This theorem implies that in the subcritical regime it is improbable that \mathcal{G}_n^p has a path connecting two lattice points $x(v), x(v')$ as defined above, as such a path must have $\omega(1)$ edges.

A.2 Back to motion planning

To conclude this section, we define a sparsified version of PRM defined on the deterministic samples \mathcal{L}_n , and state its theoretical properties. The proof is omitted as it is almost identical to the one obtained in Section 4, with the difference that it now uses theorems and lemmata presented earlier in this section, rather than the ones from Section 5.

Recall that PRM uses the two connection radii r_n, r_n^{st} . For simplicity we fix $r_n = \tilde{r}_n^*$.

Definition 12. The sparsified PRM graph \mathcal{P}_n^p is the union between $\mathcal{G}_n^p \cap \mathcal{F}$ and the supplementary edges

$$\bigcup_{y \in \{s, t\}} \{(x, y) \mid x \in \mathcal{G}_n^p \cap \mathcal{B}_{r_n^{st}}(y) \text{ and } xy \subset \mathcal{F}\}.$$

Theorem 9. Suppose that (\mathcal{F}, s, t) is robustly feasible. Then there exists $p^* > 0$ such that the following holds:

- i. If $p < p^*$ and $r_n^{st} = \infty$ then PRM fails (to find a solution) a.a.s.
- ii. Suppose that $p > p^*$. There exists $\beta'_0 > 0$ such that for $r_n^{st} = \frac{\beta' \log^{1/(d-1)} n}{n^{1/d}}$, where $\beta' > \beta'_0$, and any $\varepsilon > 0$ the following holds a.a.s.:
 1. \mathcal{P}_n^p contains a path $\pi_n \in \Pi_{s,t}^{\mathcal{F}}$ with $c_\ell(\pi_n) \leq (1 + \varepsilon)\xi' \cdot c_\ell^*$;
 2. If \mathcal{M} is well behaved then, \mathcal{P}_n^p contains a path $\pi'_n \in \Pi_{s,t}^{\mathcal{F}}$ with $c_b(\pi'_n, \mathcal{M}) \leq (1 + \varepsilon)c_b^*$.

B Probabilistic completeness of RRT with samples from a PPP

This section is devoted to the proof of Claim 6. A probabilistic completeness proof of RRT, for the setting of uniform random sampling, is given in [27, Section 3]. We describe here how to modify this proof for samples from a PPP, which are of interest in the present paper.

Note that RRT is an incremental planning algorithm, which generates samples one after the other, and connects the current sample to previous ones. Therefore, it should employ the incremental PPP sampling defined in Claim 5.

We wish to apply the line of arguments as for the uniform-sampling case described in [27, Theorem 1]. This theorem assumes that a path π from start to target exists, whose clearance is $\delta > 0$ and length is L . The path π is covered with a set of m balls of radius r , where $m = L/r$ and r is a constant depending on both the steering parameter η of RRT and the clearance δ . The proof describes the process as a Markov chain with transition probability p , where p is the probability to produce an RRT vertex in the $(i + 1)$ st ball $\mathcal{B}_r(x_{i+1})$ given that there exists an RRT vertex in the i th ball $\mathcal{B}_r(x_i)$. It then shows that the probability that RRT will not produce samples within all m balls decays to zero exponentially as the number of samples tends to infinity. To adapt the proof for samples from a PPP, it suffices to outline how to compute the transition probability p ,

Recall that at each iteration i we draw a sample $N_i \in N$, where N is a Poisson random variable with mean 1, and define $\mathcal{X}_i = \{X_1, \dots, X_{N_i}\}$ to be N_i points chosen independently and uniformly at random from $[0, 1]^d$. Let A denote the event that $N_i \geq 1$, and let B denote the event that at least one of X_1, \dots, X_{N_i} falls inside $\mathcal{B}_r(x_{i+1})$. Since $p = \Pr[A \cap B] = \Pr[A] \cdot \Pr[B|A]$, we will bound $\Pr[A]$, $\Pr[B|A]$. Observe that

$$\Pr[A] = \Pr[N_i \geq 1] = 1 - \Pr[N_i = 0] = 1 - e^{-1},$$

and note that $\Pr[B|A] \geq |\mathcal{B}_r|/|\mathcal{F}|$. Therefore,

$$p = \Pr[A \cap B] \geq (1 - e^{-1}) \cdot |\mathcal{B}_r|/|\mathcal{F}|.$$

The rest of the proof of Claim 6 follows the same lines as [27, Theorem 1].

C Critical values

In this section we provide for reference a list of (estimates) of the values γ^* and p^* .

C.1 Continuum percolation

Table 2 presents an estimate of γ^* such that \mathcal{G}_n percolates only if $r_n > \gamma^* n^{-1/d}$, for $2 \leq d \leq 11$. It is directly derived from [50, Table I], where an estimate of the critical node degree $\Delta(n, r_n)$ is derived. We mention that for d large enough $\gamma^* \sim \frac{1}{2b_d^{1/d}}$ [50], where b_d is volume of the Lebesgue measure of the unit ball in \mathbb{R}^d .

d	γ^*
2	0.5992373341(3)
3	0.4341989179(2)
4	0.4031827664(5)
5	0.4007817822(7)
6	0.4067135508(5)
7	0.4178609367(3)
8	0.4317877097(6)
9	0.447061366(4)
10	0.462335684(3)
11	0.4773913785(3)

Table 2: Values of γ^* in continuum percolation.

d	p^*
2	0.5
3	0.24881182(10)
4	0.1601314(13)
5	0.118172(1)
6	0.0942019(6)
7	0.0786752(3)
8	0.06770839(7)
9	0.05949601(5)
10	0.05309258(4)
11	0.04794969(1)
12	0.04372386(1)
13	0.04018762(1)

Table 3: Values of p^* in lattice percolation. The case for $d = 2$ is due to [11]; $d = 3$ is due to [52]; $d \in \{4, \dots, 13\}$ is due to [15].

C.2 Lattice percolation

Table 3 presents the value of p^* for $d = 2$, the best known estimates for $4 \leq d \leq 13$. Numbers in round brackets are single standard deviations.

References

- [1] Aviv Adler, Mark de Berg, Dan Halperin, and Kiril Solovey. Efficient multi-robot motion planning for unlabeled discs in simple polygons. *IEEE Trans. Automation Science and Engineering*, 12(4):1309–1317, 2015.
- [2] Ron Alterovitz, Sachin Patil, and Anna Derbakova. Rapidly-exploring roadmaps: Weighing exploration vs. refinement in optimal motion planning. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3706–3712, 2011.
- [3] Peter Antal and Agoston Pisztor. On the chemical distance for supercritical bernoulli percolation. *The Annals of Probability*, 24(2):1036–1048, 1996.

- [4] Oktay Arslan and Panagiotis Tsiotras. Use of relaxation methods in sampling-based algorithms for optimal motion planning. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2413–2420, 2013.
- [5] Cenk Baykal and Ron Alterovitz. Asymptotically optimal design of piecewise cylindrical robots using motion planning. In *Robotics: Science and Systems*, 2017.
- [6] Bela Bollobás and Oliver Riordan. *Percolation*. Cambridge University Press, 2006.
- [7] John Canny. *The complexity of robot motion planning*. MIT press, 1988.
- [8] Andrew Dobson and Kostas E. Bekris. Sparse roadmap spanners for asymptotically near-optimal motion planning. *International Journal of Robotics Research*, 33(1):18–47, 2014.
- [9] Andrew Dobson, George V. Moustakides, and Kostas E. Bekris. Geometric probability results for bounding path quality in sampling-based roadmaps after finite computation. In *IEEE International Conference on Robotics and Automation*, pages 4180–4186, 2015.
- [10] Andrew Dobson, Kiril Solovey, Rahul Shome, Dan Halperin, and Kostas E. Bekris. Scalable asymptotically-optimal multi-robot motion planning. In *International Symposium on Multi-Robot and Multi-Agent Systems*, 2017.
- [11] Massimo Franceschetti and Ronald Meester. *Random Networks for Communication: From Statistical Physics to Information Systems*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2008.
- [12] Emilio Frazzoli and Marco Pavone. Multi-vehicle routing. In John Baillieul and Tariq Samad, editors, *Encyclopedia of Systems and Control*, pages 1–11. Springer London, 2013.
- [13] Tobias Friedrich, Thomas Sauerwald, and Alexandre Stauffer. Diameter and broadcast time of random geometric graphs in arbitrary dimensions. *Algorithmica*, 67(1):65–88, 2013.
- [14] Jonathan D. Gammell, Siddhartha S. Srinivasa, and Timothy D. Barfoot. Batch informed trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3067–3074, 2015.
- [15] Peter Grassberger. Critical percolation in high dimensions. *Phys. Rev. E*, 67:036101, Mar 2003.
- [16] Geoffrey R. Grimmett. *Percolation*, volume 321 of *Comprehensive Studies in Mathematics*. Springer, 1999.
- [17] Dan Halperin, Lydia Kavraki, and Kiril Solovey. Robotics. In Jacob E. Goodman, Joseph O’Rourke, and Csaba D. Toth, editors, *Handbook of Discrete and Computational Geometry*, chapter 51. CRC Press LLC, 3rd edition, 2016. URL <http://www.csun.edu/~ctoth/Handbook/HDCG3.html>.
- [18] Dan Halperin, Oren Salzman, and Micha Sharir. Algorithmic motion planning. In Jacob E. Goodman, Joseph O’Rourke, and Csaba D. Toth, editors, *Handbook of Discrete and Computational Geometry*, chapter 50. CRC Press LLC, 3rd edition, 2016. URL <http://www.csun.edu/~ctoth/Handbook/HDCG3.html>.

- [19] Rachel M. Holladay, Oren Salzman, and Siddhartha Srinivasa. Minimizing task space frechet error via efficient incremental graph search. *CoRR*, abs/1710.06738, 2017. URL <http://arxiv.org/abs/1710.06738>.
- [20] David Hsu, Jean-Claude Latombe, and Rajeev Motwani. Path planning in expansive configuration spaces. *International Journal of Computational Geometry and Applications*, 9(4/5), 1999.
- [21] Lucas Janson, Edward Schmerling, Ashley A. Clark, and Marco Pavone. Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions. *International Journal of Robotics Research*, 34(7):883–921, 2015.
- [22] Lucas Janson, Brian Ichter, and Marco Pavone. Deterministic sampling-based motion planning: Optimality, complexity, and performance. *International Journal of Robotics Research*, 2017.
- [23] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *International Journal of Robotics Research*, 30(7):846–894, 2011.
- [24] Lydia E. Kavraki, Petr Svestka, Jean-Claude Latombe, and Mark Overmars. Probabilistic roadmaps for path planning in high dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.
- [25] Lydia E. Kavraki, Mihail N. Kolountzakis, and Jean-Claude Latombe. Analysis of probabilistic roadmaps for path planning. *IEEE Trans. Robotics and Automation*, 14(1):166–171, 1998.
- [26] Michal Kleinbort, Oren Salzman, and Dan Halperin. Collision detection or nearest-neighbor search? On the computational bottleneck in sampling-based motion planning. In *International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2016.
- [27] Michal Kleinbort, Kiril Solovey, Zakary Littlefield, Kostas E. Bekris, and Dan Halperin. Probabilistic completeness of RRT for geometric and kinodynamic planning with forward propagation. *IEEE Robotics and Automation Letters*, 2018. To appear.
- [28] Andrew M. Ladd and Lydia E. Kavraki. Measure theoretic analysis of probabilistic path planning. *IEEE Trans. Robotics and Automation*, 20(2):229–242, 2004.
- [29] Steven M. LaValle and James J. Kuffner Jr. Randomized kinodynamic planning. *I. J. Robotics Res.*, 20(5):378–400, 2001.
- [30] Yanbo Li, Zakary Littlefield, and Kostas E. Bekris. Asymptotically optimal sampling-based kinodynamic planning. *I. J. Robotics Res.*, 35(5):528–564, 2016.
- [31] Ronald Meester and Rahul Roy. Uniqueness of unbounded occupied and vacant components in boolean models. *The Annals of Applied Probability*, 4(3):933–951, 1994.
- [32] Oren Nechushtan, Barak Raveh, and Dan Halperin. Sampling-diagram automata: A tool for analyzing path quality in tree planners. In *International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, pages 285–301, 2010.
- [33] Mathew D. Penrose. *Random geometric graphs*. Oxford University Press, 2003.
- [34] Mathew D. Penrose and Agoston Pisztor. Large deviations for discrete and continuous percolation. *Advances in Applied Probability*, 28(1):2952, 1996.

- [35] John H Reif. Complexity of the movers problem and generalization. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 421–427, 1979.
- [36] Oren Salzman and Dan Halperin. Asymptotically-optimal motion planning using lower bounds on cost. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4167–4172, 2015.
- [37] Oren Salzman and Dan Halperin. Asymptotically near-optimal RRT for fast, high-quality motion planning. *IEEE Trans. Robotics*, 32(3):473–483, 2016.
- [38] Oren Salzman, Doron Shaharabani, Pankaj K. Agarwal, and Dan Halperin. Sparsification of motion-planning roadmaps by edge contraction. *I. J. Robotics Res.*, 33(14):1711–1725, 2014.
- [39] Edward Schmerling, Lucas Janson, and Marco Pavone. Optimal sampling-based motion planning under differential constraints: The drift case with linear affine dynamics. In *IEEE Conference on Decision and Control*, pages 2574–2581, 2015.
- [40] Edward Schmerling, Lucas Janson, and Marco Pavone. Optimal sampling-based motion planning under differential constraints: The driftless case. In *IEEE International Conference on Robotics and Automation*, pages 2368–2375, 2015.
- [41] Kiril Solovey and Dan Halperin. Sampling-based bottleneck pathfinding with applications to Fréchet matching. In *European Symposium on Algorithms*, pages 76:1–76:16, 2016.
- [42] Kiril Solovey and Dan Halperin. On the hardness of unlabeled multi-robot motion planning. *I. J. Robotics Res.*, 35(14):1750–1759, 2016.
- [43] Kiril Solovey and Dan Halperin. Efficient sampling-based bottleneck pathfinding over cost maps. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017. to appear.
- [44] Kiril Solovey, Jingjin Yu, Or Zamir, and Dan Halperin. Motion planning for unlabeled discs with optimality guarantees. In *Robotics: Science and Systems*, 2015.
- [45] Kiril Solovey, Oren Salzman, and Dan Halperin. New perspective on sampling-based motion planning via random geometric graphs. In *Robotics: Science and Systems*, 2016.
- [46] Kiril Solovey, Oren Salzman, and Dan Halperin. Finding a needle in an exponential haystack: Discrete RRT for exploration of implicit roadmaps in multi-robot motion planning. *I. J. Robotics Res.*, 35(5):501–513, 2016.
- [47] J. A. Starek, E. Schmerling, G. D. Maher, B. W. Barbee, and M. Pavone. Real-time, propellant-optimized spacecraft motion planning under Clohessy-Wiltshire-Hill dynamics. In *IEEE Aerospace Conference*, Big Sky, Montana, 2016.
- [48] Joseph A. Starek, Javier V. Gómez, Edward Schmerling, Lucas Janson, Luis Moreno, and Marco Pavone. An asymptotically-optimal sampling-based algorithm for Bi-directional motion planning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2072–2078, 2015.
- [49] Ioan A. Şucan, Mark Moll, and Lydia E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 19(4):72–82, 2012. <http://ompl.kavrakilab.org>.

- [50] S. Torquato and Y. Jiao. Effect of dimensionality on the continuum percolation of overlapping hyperspheres and hypercubes. ii. simulation results and analyses. *The Journal of Chemical Physics*, 137(7):074106, 2012.
- [51] Glenn Wagner and Howie Choset. Subdimensional expansion for multirobot path planning. *Artif. Intell.*, 219:1–24, 2015.
- [52] Junfeng Wang, Zongzheng Zhou, Wei Zhang, Timothy M. Garoni, and Youjin Deng. Bond and site percolation in three dimensions. *Phys. Rev. E*, 87:052107, May 2013.