

The Critical Radius in Sampling-based Motion Planning

Kiril Solovey and Michal Kleinbort

Blavatnik School of Computer Science, Tel Aviv University, Israel

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 Karaman 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 $(\frac{\log n}{n})^{1/d}$. 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.

I. 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 [12], nowadays it is applied in various fields. Notably, motion planning arises in coordination of multiple autonomous vehicles [8], steering surgical needles [3], and planning trajectories of spacecrafts in orbit [32], 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.

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 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

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. K.S. is supported by the Clore Israel Foundation. M.K. is supported in part by Yandex. Contact information: kirilsol@post.tau.ac.il; michal.kleinbort@gmail.com

the probability of finding a solution tends to 1 as the number of samples n tends to infinity. Moreover, some recently-introduced sampling-based techniques are also guaranteed to return high-quality¹ solutions that tend to the optimum as n diverges—a property called *asymptotic optimality* (AO).

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.

A. Contribution

We develop a new analysis of PRM [16] for uniform random sampling in Euclidean space, which relies on a novel connection between sampling-based planners and *continuum percolation* (see, e.g., [7]). Our analysis is tight and proves the existence of a *critical connection radius* $r_n^* = \gamma^* n^{-1/d}$, where $\gamma^* > 0$ is a constant², 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*³ (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 [33] (see Section II) 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* [14] is AnO with respect to the path-length cost, while BTT [29] is AO with respect to the bottleneck cost. RRG behaves similarly for the two cost functions. Our results are also applicable to *multi-robot* motion planners such as the recently introduced dRRT* [6], and M* [35] when applied

¹Quality can be measured in terms of energy, length of the plan, clearance from obstacles, etc.

² $0.4 \leq \gamma^* \leq 0.6$ for all $d \geq 2$.

³AnO 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.

to a continuous domain. See Figure 1 for additional PRM-based planners to which our analysis is applicable (see more information in the extended version of the paper [30]).




Fig. 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.

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., [14, 15, 28, 31], which provided a rough estimate of this number, that is proportional to $O(\log n)$.

B. Organization

In Section II we discuss related work. Then we proceed to basic definitions and problem statement in Section III. 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 IV. This section also contains an outline of the proof. In Section V 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 VI. In Section VII we present experimental work, which validates our theory. We conclude this paper with directions for future work (Section VIII).

An extended version of the paper [30] provides supplementary material, and includes proofs that were omitted from the main document, additional experimental results, and a table with values of γ^* . It also presents an extension of the analysis to FMT*, BTT, RRG and discusses the implications of our results to the multi-robot planners dRRT* and M*. Additionally, it presents a new 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.

II. RELATED WORK

This section is devoted to a literature review of sampling-based planners with emphasis on their theoretical analysis.

The influential work [15] 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 PRM⁴ with a specific value of the connection radius r_n ; an AO variant of RRT [18] termed RRT*; RRG, which can be viewed as a combination between RRT and PRM. The analysis in [15] 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 in [15], other AO planners have emerged (see e.g., [1, 2, 10]). The paper [14] 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 [15], 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 in [33]. 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 Dobson et al. [5], which bounds the success rate with an expression that depends on the amount of deviation from the optimum.

A recent work [31] 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., [21]). Their work shows that one can slightly reduce the PRM and FMT* radius obtained in [14]. Furthermore, the connection with RGGs yields additional analyses of different extensions of PRM, which have not been analyzed before in a mathematically-rigorous setting.

We also mention that a number of methods have been developed to reduce the running time or space requirements of existing planners by relaxing AO to AnO, see, e.g., LBT-RRT [24], MPLB [23], SPARS2 [4], and RSEC [25].

A. Extensions

The aforementioned papers deal mainly with the cost function of *path length*. Two recent works [28, 29] 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 [29] and Fréchet matching between curves [13], 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

⁴Throughout this work we will refer to the more general algorithm PRM rather than PRM*.

([28, 29]) that BTT is AO, with respect to bottleneck cost, for the reduced connection radius that was obtained in [31].

The results reported until this point have dealt exclusively with holonomic robotic systems. Two recent papers [26, 27] develop the theoretical foundations of PRM and FMT*-flavored planners when applied to robots having differential constraints. Li et al. [19] develop an AO algorithm that does not require a steering function, as PRM for instance does.

III. 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 .

A. 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 VIII), 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 $s, t \in \mathcal{F}$.

B. Cost function

It is usually desirable to obtain paths that minimize a given criterion. In this paper we consider the following two cost functions. Given a path π , its *length* is denoted by $c_\ell(\pi)$, and its *bottleneck cost* is defined to be $c_b(\sigma, \mathcal{M}) = \max_{\tau \in [0, 1]} \mathcal{M}(\pi(\tau))$, where $\mathcal{M} : \mathcal{C} \rightarrow \mathbb{R}$ is a *cost map*.

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 1. 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 2. 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 3. 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 4. The *robust optimum* with respect to c_b is defined as $c_b^* = \inf \{c_b(\pi, \mathcal{M}) \mid \pi \in \Pi_{s,t}^\mathcal{F} \text{ is } \mathcal{M}\text{-robust}\}$.

C. 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 5 ([7]). 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 $k \geq 0$, $\Pr[|\mathcal{X} \cap D| = k] = e^{-\lambda|D|} \frac{(\lambda|D|)^k}{k!}$.

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 6. ([21]) 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 $\|x - y\| \leq r$.

IV. 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 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 7. The PRM graph \mathcal{P}_n is the union between $\mathcal{G}(\mathcal{X}_n \cap \mathcal{F}; r_n)$ and the supplementary edges

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

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 constant⁵, 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.⁶*
- 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 $c_b(\pi'_n, \mathcal{M}) \leq (1 + \varepsilon)c_b^*$.

A. 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 VI. 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 V-A). 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 V-B). 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 VI): 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 [22] 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 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 V we provide background on continuum percolation in unbounded and bounded domains, and prove two key lemmata (Lemma 2 and Lemma 3). In Section VI we return to the setting of motion planning and utilize the aforementioned results in the proof of Theorem 1.

V. 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.

A. The basics

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

Definition 8. 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. ([11, 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. See proof in [30].

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 $\theta(n, r) > 0$.*

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

Theorem 3. ([20, Theorem 2.3]) *With probability 1, $\mathcal{G}(\mathcal{X}_n; r)$ contains at most one infinite connected component.*

B. 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.

⁵In particular, for any $d \geq 2$, $\gamma^* \in (0.4, 0.6)$ (see [30]).

⁶Let 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 $\lim_{n \rightarrow \infty} \Pr[B(A_n)] = 1$.

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*

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

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 [22]⁷. 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, $\text{diam}(C'_i) \leq \log^{-1} n - r_n$ is required for C'_i 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 for $i > 1$.

Theorem 2 in [22] 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. ■




Fig. 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.

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

⁷We mention that [22] 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.

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 $\Pr[\mathfrak{E}_n^2 | \mathfrak{E}_n^1] \geq 1 - n^{-1}$.*

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

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

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 [22], which bounds the number of vertices from the unbounded component. In particular, there exists $\beta_0 > 0$ such that

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

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 $1 - n^{-1}$. ■

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. ([9, 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 $\text{dist}(\mathcal{G}_n, x, x') \leq \xi \|x - x'\|$.*

VI. PROOF OF THEOREM 1

We provide a full proof for the positive setting with length cost. The proof for the bottleneck case is very similar to the length case, and it is therefore deferred to the extended version [30]. The incompleteness proof for $r_n < r_n^*$ is presented there as well.

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 component C_∞ . Recall that C_n denotes the largest connected component of $C_\infty \cap [0, 1]^d$. Also note that $r_n^{st} = h'_n$, where h'_n is defined in Lemma 3.

Recall that c_ℓ^* denotes the robust optimum, with respect to path length (Definition 2). Fix $\varepsilon > 0$. By definition, there exists a robust path $\pi_\varepsilon \in \Pi_{s,t}^{\mathcal{F}}$ and $\delta > 0$ such that $c_\ell(\pi_\varepsilon) \leq (1+\varepsilon)c_\ell^*$ and $\mathcal{B}_\delta(\pi_\varepsilon) \subset \mathcal{F}$. See illustration in Figure 3.




Fig. 3. 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 $C \setminus \mathcal{F}$. $p_1 = s, p_2, \dots, p_{k-1}, p_k = t$ 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 .

We now define a sequence of k points p_1, \dots, p_k along π_ε that are separated by exactly $\delta/2\xi$ units, where ξ is as defined in Theorem 4. In particular, define $k = \lceil c_\ell(\pi_\varepsilon) \cdot 2\xi/\delta \rceil$, set $p_1 = s, p_k = t$, and assign p_i along π_ε , such that $c_\ell(\pi_\varepsilon^{i-1,i}) = \delta/2\xi$, where $\pi_\varepsilon^{i-1,i}$ represents the subpath of π_ε starting at p_{i-1} and ending at p_i . Notice that k is finite.

Claim 1. Denote by \mathfrak{E}_n^4 the event that for all $i \in [k]$ there exists $q_i \in C_n$ such that $q_i \in \mathcal{B}_{r_n^{st}}(p_i)$ and $q_i \in C_n$. Then $\Pr[\mathfrak{E}_n^4 | \mathfrak{E}_n^1] \geq 1 - k \Pr[\overline{\mathfrak{E}_n^2} | \mathfrak{E}_n^1]$ (see definition of $\mathfrak{E}_n^1, \mathfrak{E}_n^2$ in Lemma 2 and Lemma 3, respectively).

Proof: Define $H'_n(x) \subset \mathbb{R}^d$ to represent a d -dimensional (axis-aligned) hypercube of side length h'_n that is centered in $x \in \mathbb{R}^d$. Formally, $H'_n(x) = x + h'_n \cdot [-\frac{1}{2}, \frac{1}{2}]^d$. Observe that $H'_n(p_i) \subset H_n$ for n large enough. Also note that $H'_n(p_i) \subset \mathcal{B}_{r_n^{st}}(p_i)$. Thus, the result follows from the union bound. ■

Suppose that $\mathfrak{E}_n^1, \mathfrak{E}_n^4$ are satisfied. Let $q_1, \dots, q_k \in C_n$ be the points obtained from Claim 1. These points reside in a single connected component of \mathcal{G}_n . Define the path

$$\pi_n := s \rightarrow q_1 \rightsquigarrow q_2 \rightsquigarrow \dots \rightsquigarrow q_k \rightarrow t,$$

where $s \rightarrow q_1$ represents a straight-line path from s to q_1 , and $q_i \rightsquigarrow q_{i+1}$ represents the shortest path from q_i to q_{i+1} in \mathcal{G}_n . The next claim states that if in addition \mathfrak{E}_n^3 is satisfied (Theorem 4), then π_n is also collision free.

Claim 2. Suppose that $\mathfrak{E}_n^1, \mathfrak{E}_n^3, \mathfrak{E}_n^4$ are satisfied. Then $\pi_n \in \Pi_{s,t}^{\mathcal{F}}$ is a path in \mathcal{P}_n (with probability 1).

Proof: First, observe that the straight-line paths $s \rightarrow q_1, q_k \rightarrow t$ are contained in \mathcal{P}_n due to the definition of PRM and the fact that $r_n^{st} < \delta$. Let us consider a specific subpath $q_i \rightsquigarrow q_{i+1}$. Recall from Claim 1 and by definition of p_i, p_{i+1} that $\text{dist}(\mathcal{G}_n, q_i, q_{i+1}) \leq 2r_n^{st} + \xi \|p_{i+1} - p_i\| = o(1) + \delta/2$, which is bounded by δ . Thus, for any point q'_i along $q_i \rightsquigarrow q_{i+1}$ it holds that $\|q'_i - p_i\| < \delta, \|q'_i - p_{i+1}\| < \delta$. Thus, $\text{Im}(\pi_n) \subset \mathcal{B}_\delta(\pi_\varepsilon) \subset \mathcal{F}$. ■

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

Claim 3. 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: By Theorem 4 and the triangle inequality, it follows that

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

To conclude, Claim 2 and Claim 3 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:

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

By Claim 1 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,

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

VII. 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 [14, 15].


We validate our theory for PRM (Theorem1) and FMT* (See [30]) for the path-length cost c_ℓ . 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 rate when compared to the results obtained using the original FMT* connection radius.

A. Implementation details

In our experiments, we used the Open Motion Planning Library (OMPL 1.3) [34] on a 2.6GHz×2 Intel Core i5 processor with 16GB of memory. Result 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. 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 [15]. 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 [14]. 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 [15]). The figure to the right depicts the colors and labeling that will be used throughout this section.



B. Results

Euclidean space. The scenario we consider 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




Fig. 4. Two of the scenarios used in the experiments. (a) dD Hypercube with 2^d hypercubical obstacles, and (b) 3D Cubicles. Start and target configurations for a robot are depicted in green and red, respectively. Scenario (b) is provided with the OMPL distribution.

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 . The results for dimensions 4 and 8 are depicted in Figure 5a and Figure 5b, respectively. See [30] for plots for $d = 12$.

The plots demonstrate the following trend: for each radius r the cost obtained by PRM converges to some constant⁸ times the optimal cost, which is marked with the dashed red curve in the “Cost vs. n ” plot. Clearly, r_{PRM^*} yields the best cost but at the price of increased running time. $r_{10} = r_{FMT^*}$ obtains the next best cost, with improved running time, and so on. Note that for $d = 4$, already for $n = 1K$ a solution is found for all radii except r_0 , whereas for $d = 8$ and $n = 5K$ 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_{FMT^*}, r_{PRM^*} , with a slight penalty (a factor of roughly 2) in the resulting costs. In [30] we also study how r_i affects the size of C_n .




Fig. 5. PRM in the (a) 4D and (b) 8D Hypercube.

General Euclidean space. We consider the following d -

⁸We note that it is possible that for larger values of n the cost values for different radii will eventually converge to the same value.

dimensional scenario (based on a scenario from [31]), for $d \in \{4, 8\}$, depicted in Figure 4a 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 $(1, \dots, 1)$.

Here we use FMT* with radii ranging from r_0 to r_{10} . We also ran PRM which exhibited a similar behavior. Figure 6a 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 $d = 8$ we obtain similar results. However, the success rates deteriorate as the dimension increases (see extended version [30]).

General non-Euclidean space. We use the Cubicles scenario (see Figure 4b) provided with the OMPL distribution [34]. Here the goal is to find a collision-free path for two L-shaped robots that need to exchange their positions (in green and red). Since the robots are allowed to translate and rotate, then $d = 12$ and \mathcal{C} is non-Euclidean. Although our theory does directly apply here, we chose to test it empirically for such a setting. A similar experiment involving only one L-shaped robot is discussed in [30].

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 [17], 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 6b) we increased all radii by a multiplicative factor of 10 in order to increase the success rates. We mention that this yielded similar behavior to that when using the connection scheme suggested in [17].

Observe in Figure 6b, 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.




Fig. 6. (a) FMT* for a point robot in the 4D hypercube with obstacles scenario (Figure 4a) (b) FMT* for two rigid-body robots moving in OMPL’s Cubicles scenario (Figure 4b). 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).

VIII. 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 [26, 27].

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 [19, 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.

REFERENCES

- [1] 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.
- [2] 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.
- [3] Cenk Baykal and Ron Alterovitz. Asymptotically optimal design of piecewise cylindrical robots using motion planning. In *Robotics: Science and Systems*, 2017.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] Emilio Frazzoli and Marco Pavone. Multi-vehicle routing. In *Encyclopedia of Systems and Control*. 2015.
- [9] Tobias Friedrich, Thomas Sauerwald, and Alexandre Stauffer. Diameter and broadcast time of random geometric graphs in arbitrary dimensions. *Algorithmica*, 67(1):65–88, 2013.
- [10] 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.
- [11] Geoffrey R. Grimmett. *Percolation*, volume 321 of *Comprehensive Studies in Mathematics*. Springer, 1999.
- [12] 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>.
- [13] 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>.
- [14] 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.
- [15] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *International Journal of Robotics Research*, 30(7):846–894, 2011.
- [16] 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.
- [17] 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.
- [18] Steven M. LaValle and James J. Kuffner Jr. Randomized kinodynamic planning. *I. J. Robotics Res.*, 20(5):378–400, 2001.
- [19] Yanbo Li, Zakary Littlefield, and Kostas E. Bekris. Asymptotically optimal sampling-based kinodynamic planning. *I. J. Robotics Res.*, 35(5):528–564, 2016.
- [20] 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.
- [21] Mathew D. Penrose. *Random geometric graphs*. Oxford University Press, 2003.
- [22] Mathew D. Penrose and Agoston Pisztora. Large deviations for discrete and continuous percolation. *Advances in Applied Probability*, 28(1):2952, 1996.
- [23] 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.
- [24] Oren Salzman and Dan Halperin. Asymptotically near-optimal RRT for fast, high-quality motion planning. *IEEE Trans. Robotics*, 32(3):473–483, 2016.
- [25] 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.
- [26] 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.
- [27] 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.

- [28] 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.
- [29] 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.
- [30] Kiril Solovey and Michal Kleinbort. The critical radius in sampling-based motion planning. *CoRR*, abs/1709.06290, 2017. URL <http://arxiv.org/abs/1709.06290>.
- [31] Kiril Solovey, Oren Salzman, and Dan Halperin. New perspective on sampling-based motion planning via random geometric graphs. In *Robotics: Science and Systems*, 2016.
- [32] 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.
- [33] 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.
- [34] 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>.
- [35] Glenn Wagner and Howie Choset. Subdimensional expansion for multirobot path planning. *Artif. Intell.*, 219: 1–24, 2015.