

New perspective on sampling-based motion planning via random geometric graphs*

Kiril Solovey, Oren Salzman and Dan Halperin

Blavatnic School of Computer Science
Tel Aviv University, Israel
email: {kirisol,orensal,danha}@post.tau.ac.il

Abstract—Roadmaps constructed by many sampling-based motion planners coincide, in the absence of obstacles, with standard models of random geometric graphs (RGGs). Those models have been studied for several decades and by now a rich body of literature exists analyzing various properties and types of RGGs. In their seminal work on optimal motion planning Karaman and Frazzoli [31] conjectured that a sampling-based planner has a certain property if the underlying RGG has this property as well. In this paper we settle this conjecture and leverage it for the development of a general framework for the analysis of sampling-based planners. Our framework, which we call *localization-tessellation*, allows for easy transfer of arguments on RGGs from the free unit-hypercube to spaces punctured by obstacles, which are geometrically and topologically much more complex. We demonstrate its power by providing alternative and (arguably) simple proofs for probabilistic completeness and asymptotic (near-)optimality of probabilistic roadmaps (PRMs). Furthermore, we introduce two variants of PRMs, analyze them using our framework, and discuss the implications of the analysis.

I. INTRODUCTION

Motion planning is a fundamental research area in robotics with applications in diverse domains such as graphical animation, surgical planning, computational biology and computer games. For a general overview of the subject and its applications, see, e.g., [14, 37, 38].

The basic problem of motion planning is concerned with finding a collision-free path for a robot in a *workspace* cluttered with static obstacles. The spatial pose of the robot, or its *configuration*, is uniquely defined by its degrees of freedom (DOFs). The set of all configurations \mathcal{C} is termed the *configuration space* of the robot, and decomposes into the disjoint sets of free and forbidden configurations, namely \mathcal{F} and $\mathcal{C} \setminus \mathcal{F}$, respectively. Thus, given start and target configurations, the problem can be restated as the task of finding a continuous curve in \mathcal{F} connecting the two configurations. This can be very challenging, as \mathcal{F} can be exponentially complex (see, e.g., [13, 57, 63]) in the number of DOFs.

The high computational complexity of exact solutions to motion planning have led to the development of sampling-based planners. These algorithms, which trade completeness with applicability in practical settings, aim to capture the connectivity of \mathcal{F} in a graph data structure, called a roadmap, by randomly sampling \mathcal{C} . Most of the theoretical properties of these algorithms are stated in terms of their *asymptotic* behavior, i.e., assuming that the number of samples is sufficiently large: The property of *probabilistic completeness* indicates that a given algorithm will eventually find a solution (if one exists); algorithms that are known to be *asymptotically optimal* also return a solution whose cost converges to the optimum.

Interestingly, roadmaps constructed by many sampling-based planners coincide, in the absence of obstacles, with standard models of random geometric graphs (RGGs). These models have been studied for several decades and by now a rich body of literature exists analyzing various properties and types of RGGs. Indeed, in their seminal work on optimal motion planning, Karaman and Frazzoli [31] observed this relation. They employed techniques that were initially developed for the analysis of RGGs to the study of sampling-based planners. Subsequent proofs regarding completeness and optimality of new planners (see, e.g., [22, 29, 58]) rely, to some extent, on the proofs in [31]. Karaman and Frazzoli conjectured that a sampling-based planner possesses a certain property if the underlying RGG has this property as well (see [31, Section 6]). The validity of this conjecture, which is settled in this paper, allows to import existing results on RGGs directly to the corresponding sampling-based planners.

Contribution. We introduce the *localization-tessellation* framework for the analysis of sampling-based algorithms in motion planning. Our framework facilitates the extension of properties of RGGs to sampling-based techniques in motion planning. This is done using conceptually simple ideas and elementary tools in probability theory. The underlying result of the framework is that RGGs demonstrate similar behavior in the absence as well as in the presence of obstacles. The framework consists of two main components. First we show through *localization* that RGGs maintain their properties in arbitrarily-small neighborhoods. The *tessellation* stage extends these properties to complex domains which can be viewed

* This work has been supported in part by the Israel Science Foundation (grant no. 1102/11), by the German-Israeli Foundation (grant no. 1150-82.6/2011), and by the Hermann Minkowski-Minerva Center for Geometry at Tel Aviv University. Kiril Solovey is also supported by the Clore Israel Foundation.

as free spaces of motion-planning problems. Namely, the configuration space punctured by obstacles.

We demonstrate the power of the framework by providing conditions for probabilistic completeness and asymptotic (near-)optimality of Probabilistic Roadmaps (PRMs) [32]. Our proofs are (arguably) much simpler than the original proofs of Karaman and Frazzoli [31].

Furthermore, we introduce two variants of PRMs called Soft-PRM and Embedded-PRM, which perform connections in a randomized fashion, and analyze them using our framework. Using Soft-PRMs we show that the standard PRM still maintains its favorable properties even when implemented using *approximate* nearest-neighbor search queries.

Organization. In Section II we review related work. In Section III we provide formal definitions of several types of RGGs and describe their properties, which will be employed by our localization-tessellation framework. In Section IV we describe the *localization* component of the framework, that is, we show that RGGs maintain a wide range of their properties in arbitrarily-small neighborhoods. In Section V we focus on the two specific properties of connectivity and bounded stretch and show that they hold in general domains via a *tessellation* argument. In Section VI we make the transition to motion planning: we describe several planners—including the standard PRM—and study their asymptotic behavior using the framework. In Section VII we show empirically that the theoretical results obtained by the framework also hold in practice and compare the Soft-PRM and PRM algorithms. We conclude the paper with a discussion and state several future research directions (Section VIII). Due to the lack of space, we provide the main proofs and omit proofs of lesser importance. Omitted proofs are found in the extended version of the paper [64].

II. RELATED WORK

We review related work in the area of sampling-based algorithms for motion planning and random geometric graphs.

A. Sampling-based motion planning

Sampling-based algorithms, such as PRMs [32], Expansive Space Trees (EST) [27] and Rapidly-exploring Random Trees (RRT) [35], as well as their many variants, have proven to be effective tools for motion planning. These algorithms, and others were shown to be probabilistically complete. While this is a desirable property of any algorithm, in certain applications stronger guarantees are required.

In recent years we have seen an increasing interest in *high-quality*¹ motion planning. The literature contains many examples of planners that are shown empirically to produce high-quality paths (for a partial list see [4, 23, 40, 42, 56, 62, 67]). Unfortunately, they are not backed by rigorous proofs pertaining to the quality of the solution produced by the algorithm. A

¹Quality can be measured in terms of length, clearance, smoothness, energy, to mention a few criteria. However, in this paper we will restrict our focus to the standard length measure.

complementary work proves that in certain settings RRT can produce paths of arbitrarily-poor quality [47].

In their seminal work, Karaman and Frazzoli [31] develop the first rigorous analysis of quality in the setting of sampling-based motion planning: They provide conditions under which existing planners are not asymptotically optimal. More importantly, they introduce two new variants of RRT and PRM, termed RRT* and PRM*, which are shown to be asymptotically optimal, under the right choice of parameters. Following this exposition, several asymptotically-optimal algorithms have emerged (see e.g., [3, 6, 22, 29, 59]). To reduce the running time of such algorithms several asymptotically *near* optimal planners have been suggested, which trade the quality of the solution with speed of computation (see e.g., [16, 41, 58, 60]).

Although the focus of this paper is on the simplified “geometric” setting of motion planning, we mention that some planners can cope with more complex robotic systems in which uncertainty and physical constraints come into play (see, e.g., [30, 36, 39, 55, 65, 70, 71]). Some of these planners can also produce high-quality paths.

B. Random geometric graphs

The study of random geometric graphs (RGGs) was initiated by Gilbert [24] who considered the following model: a collection of points is sampled at random in a given subspace of \mathbb{R}^d , and a graph is formed by drawing edges between points that are closer than a given $r > 0$, called the *connection radius*.

An immediate question that follows is for which values of r the graph is connected (with high probability). Several works have addressed this question and showed that it is both necessary and sufficient that the connection radius will be proportional to $\left(\frac{\log n}{n}\right)^{1/d}$, where n is the number of points and the points are sampled from the unit hypercube $[0, 1]^d$ (see, e.g., [5, 34, 48]). Penrose [49] established that connectivity occurs approximately when the graph has no isolated vertices. The monograph [50] of the same author on this subject studies many more properties of RGGs, including vertex degree, clique size and coloring. The reader is also referred to a survey on the subject by Walters [69].

In recent years RGGs have attracted much attention as a tool for modeling large-scale communication networks, and in particular sensor networks: the vertices of the graph represent sensors and an edge is drawn between two sensors that are in the communication range. Gupta and Kumar used this analogy in order to deduce the transmission power necessary for the network to be connected [25]. An important parameter that arises in this context is the number of transmitters a message has to traverse in order to establish a broadcast between two given transmitters. Several works have established that this parameter is proportional to the Euclidean distance between the two nodes (see, e.g., [11, 15, 18, 20, 44, 46]).

Various alternative connection strategies for RGGs have been proposed over the years, the most studied of which is the k -nearest model (see, e.g., [8, 9, 72]). More complex models assign edges between vertices in a randomized fashion (see,

e.g., [12, 21, 51]). Some models introduce an ordering on the sampled points (see, e.g., [1, 52, 53, 61, 68]) which results in a directed graph that resembles the RRT tree [35].

III. PRELIMINARIES

We describe several models of random geometric graphs (RGGs) and mention useful properties that will be used throughout the paper. When possible, we follow the notation and conventions in the standard literature of RGGs (see, e.g., [50]). Let $\mathcal{X}_n = \{X_1, \dots, X_n\}$ be n points chosen independently and uniformly at random from the Euclidean d -dimensional cube $[0, 1]^d$. We assume that the dimension d of the domain is fixed and greater than one. Let $\|x - y\|_2$ denote the Euclidean distance between two points $x, y \in \mathbb{R}^d$ and θ_d denote the Lebesgue measure of the unit ball in \mathbb{R}^d . Finally, denote by $\mathcal{B}_r(x)$ be 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 $\sigma : [0, 1] \rightarrow \mathbb{R}^d$ denote $\mathcal{B}_r(\sigma) = \bigcup_{\tau \in [0, 1]} \mathcal{B}_r(\sigma(\tau))$.

Throughout the paper we will use the standard notation for asymptotic bounds: Let $f = f(n), g = g(n)$ be two functions. The notation $f = \omega(g)$ indicates that $\lim_{n \rightarrow \infty} f/g \rightarrow \infty$, and $f = o(g)$ indicates that $\lim_{n \rightarrow \infty} f/g \rightarrow 0$. Let 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 *almost surely* (a.s., in short) if $\lim_{n \rightarrow \infty} \Pr[B(A_n)] = 1$.

Definition 1 [50]. Given $r_n \in \mathbb{R}^+$, the *random geometric graph* (RGG) $\mathcal{G}^{\text{disk}}(\mathcal{X}_n; r_n)$ is an undirected graph with the vertex set \mathcal{X}_n . For any two given vertices $x, y \in \mathcal{X}_n$ the graph contains the edge (x, y) if $\|x - y\|_2 \leq r_n$.

We use the term RGG to refer both to the family of random geometric graphs and to the specific model described in Def. 1. This slight abuse of notation is introduced to be consistent with existing literature and the exact meaning of RGG will be clear from the context.

The following model is a generalization of RGGs. Here a pair of vertices are connected by an edge with a probability that depends on the length of the edge.

Definition 2 [51]. Let $r_n \in \mathbb{R}^+$, and let ϕ_n be a probability measure over $[0, r_n]$. The *soft random geometric graph* (SRGG) $\mathcal{G}^{\text{soft}}(\mathcal{X}_n; r_n; \phi_n)$ is an undirected graph with the vertex set \mathcal{X}_n . Denote by E the edge set of this graph. For any pair of vertices $x, y \in \mathcal{X}_n$, the following holds independently:

$$\Pr[(x, y) \in E] = \begin{cases} 0 & \text{if } \|x - y\|_2 > r_n, \\ \phi_n(\|x - y\|_2) & \text{otherwise.} \end{cases}$$

The following model can be viewed as a special case of SRGG where $r_n = \infty$ and ϕ_n is constant.

Definition 3 [21]. The *randomly-embedded geometric graph* (REGG) $\mathcal{G}^{\text{embed}}(\mathcal{X}_n; p_n)$ is an undirected graph with the vertex set \mathcal{X}_n . For every two distinct vertices $x, y \in \mathcal{X}_n$, the graph contains the edge (x, y) with probability p_n , and independently from the other edges.

A. Connectivity

We mention two results related to the connectivity of the RGG and SRGG models.

Theorem 1 [12]. Let $\mathcal{G}_n = \mathcal{G}^{\text{disk}}(\mathcal{X}_n, r_n)$ and $r_n = \gamma \left(\frac{\log n}{n}\right)^{1/d}$. Then \mathcal{G}_n is disconnected a.s. if $\gamma < \gamma^*$, and connected a.s. if $\gamma > \gamma^*$, where $\gamma^* = 2(2d\theta_d)^{-1/d}$.

The following theorem follows from [51, Theorem 2.3]. The proof is technical and omitted due to lack of space.

Theorem 2. Let $\mathcal{G}_n = \mathcal{G}^{\text{soft}}(\mathcal{X}_n; r_n; \phi_n)$ and $r_n = \gamma \left(\frac{\log n}{n}\right)^{1/d}$. Set $\gamma = \alpha\gamma^*$ for any $\alpha \geq (d+1)^{1/d}$, and define $\phi_n(z) = 1 - z/r_n$, for any $z \in \mathbb{R}^+$. Then \mathcal{G}_n is connected a.s.

B. Bounded stretch

Let \mathcal{G} be a graph whose vertices are embedded in \mathbb{R}^d . For every two vertices $x, y \in \mathcal{G}$ denote their *weighted graph distance*, i.e., the sum of lengths of the shortest path from x to y , by $\text{dist}(\mathcal{G}, x, y)$. Throughout the paper we will use the term *stretch* to denote the ratio between $\text{dist}(\mathcal{G}, x, y)$ and the length of the shortest path between x, y in the domain in which the graph is embedded. For instance, if this domain is convex, then for every $x, y \in \mathcal{G}$ the stretch is defined to be $\text{dist}(\mathcal{G}, x, y) / \|x - y\|_2$.

Theorem 3 [20]. Let $\mathcal{G}_n = \mathcal{G}^{\text{disk}}(\mathcal{X}_n; r_n)$ with $r_n = \gamma \left(\frac{\log n}{n}\right)^{1/d}$ where $\gamma > \gamma^*$ (see Theorem 1). Then there exists a constant ζ such that for every two vertices x, y in the same connected component of \mathcal{G}_n , with $\|x - y\|_2 = \omega(r_n)$, it holds that $\text{dist}(\mathcal{G}_n, x, y)$ is at most $\zeta \|x - y\|_2$ a.s.

Theorem 4 [21]. Let $\mathcal{G}_n = \mathcal{G}^{\text{embed}}(\mathcal{X}_n; p_n)$ and $p_n = \omega \left(\frac{\log^d n}{n}\right)$. Then for every two vertices $x, y \in \mathcal{X}_n$ it holds that $\text{dist}(\mathcal{G}_n, x, y)$ is at most $\|x - y\|_2 + o(1)$ a.s.

IV. LOCALIZATION OF MONOTONE PROPERTIES OF RGGs

In this section we discuss graph properties and their asymptotic behavior, when focusing on a subset of the domain $[0, 1]^d$. A property \mathcal{A} is *monotone* if for every $G = (V, E)$ and $H = (V, E')$ such that $E \subseteq E'$, it holds that $G \in \mathcal{A} \implies H \in \mathcal{A}$. Note that connectivity (Section III-A) and bounded stretch (Section III-B) are monotone.²

Definition 4. Let $\mathcal{G} = (X, E)$ be a graph embedded in $[0, 1]^d$, i.e., the vertices of X represent points in $[0, 1]^d$ and edges represent straight-line paths between the corresponding vertices. Given $\Gamma \subset [0, 1]^d$ we denote by $\mathcal{G}(\Gamma)$ the graph obtained from the intersection of \mathcal{G} and Γ . This graph consists of the vertex set $X \cap \Gamma$ and all the edges in E that are fully contained in Γ .

Definition 5. Let \mathcal{G}_n be an RGG, SRGG or an REGG, defined over the vertex set \mathcal{X}_n . Then \mathcal{G}_n is *localizable* for a property \mathcal{A}

² Additional examples of monotone properties for a graph \mathcal{G} are: \mathcal{G} is Hamiltonian, \mathcal{G} contains a clique of size t , \mathcal{G} is not planar, the clique number of \mathcal{G} is larger than that of its complement, the diameter of \mathcal{G} is at most s , etc.

if for every constant $0 < \varepsilon \leq 1$ and every d -dimensional axis-aligned cube $B_\varepsilon \subseteq [0, 1]^d$ with side length of ε it holds that $\mathcal{G}_n(B_\varepsilon) \in \mathcal{A}$ a.s.

Lemma 1. *Let \mathcal{A} be a monotone property and $\gamma_{\mathcal{A}}$ some constant. Let $\mathcal{G}_n = \mathcal{G}^{\text{disk}}(\mathcal{X}_n; r_n)$ be an RGG such that $\mathcal{G}_n \in \mathcal{A}$ a.s., for every $r_n = \gamma \left(\frac{\log n}{n}\right)^{1/d}$, where $\gamma > \gamma_{\mathcal{A}}$. Then \mathcal{G}_n is localizable for \mathcal{A} .*

Proof: For simplicity of presentation, we will use the notation \mathcal{G} to refer to $\mathcal{G}^{\text{disk}}$ throughout the proof. Recall that \mathcal{X}_n is a collection of n points chosen independently and uniformly at random from $[0, 1]^d$. We will also use $\mathcal{Y}_m^\varepsilon = \{Y_1, \dots, Y_m\}$ to denote a collection of m points chosen independently and uniformly at random from B_ε . Without loss of generality, assume that $B_\varepsilon = [0, \varepsilon]^d$.

Observe that there exists a constant $\alpha > 1$ independent of n such that $\mathcal{G}(\mathcal{X}_n; r'_n) \in \mathcal{A}$ a.s., where $r_n = \alpha r'_n$. The role of α is purely technical and will become clear shortly. Now,

$$\begin{aligned} \Pr[\mathcal{G}(\mathcal{X}_n \cap B_\varepsilon; r_n) \notin \mathcal{A}] &= \Pr[\mathcal{G}(\mathcal{X}_n \cap B_\varepsilon; \alpha r'_n) \notin \mathcal{A}] \\ &= \sum_{m=0}^n \Pr \left[\mathcal{G}(\mathcal{X}_n \cap B_\varepsilon; \alpha r'_n) \notin \mathcal{A} \mid |\mathcal{X}_n \cap B_\varepsilon| = m \right] \\ &\quad \cdot \Pr[|\mathcal{X}_n \cap B_\varepsilon| = m] \\ &= \sum_{m=0}^n \Pr[\mathcal{G}(\mathcal{Y}_m^\varepsilon; \alpha r'_n) \notin \mathcal{A}] \cdot \Pr[|\mathcal{X}_n \cap B_\varepsilon| = m]. \end{aligned}$$

Denote

$$\sigma(i, j) = \sum_{m=i}^j \Pr[\mathcal{G}(\mathcal{Y}_m^\varepsilon; \alpha r'_n) \notin \mathcal{A}] \cdot \Pr[|\mathcal{X}_n \cap B_\varepsilon| = m],$$

and by definition we have that for $1 \leq \ell \leq n$

$$\Pr[\mathcal{G}(\mathcal{X}_n \cap B_\varepsilon; \alpha r'_n) \notin \mathcal{A}] = \sigma(0, n) = \sigma(0, \ell - 1) + \sigma(\ell, n).$$

We show that for $\ell = \alpha^{-d} \varepsilon^d n$ both $\lim_{n \rightarrow \infty} \sigma(0, \ell - 1) = 0$ and $\lim_{n \rightarrow \infty} \sigma(\ell, n) = 0$ which will conclude the proof of the lemma (for simplicity we assume that $\ell \in \mathbb{N}$). We start with the former expression:

$$\begin{aligned} \sigma(0, \ell - 1) &= \sum_{m=0}^{\ell-1} \Pr[\mathcal{G}(\mathcal{Y}_m^\varepsilon; \alpha r'_n) \notin \mathcal{A}] \cdot \Pr[|\mathcal{X}_n \cap B_\varepsilon| = m] \\ &\leq \sum_{m=0}^{\ell-1} \Pr[|\mathcal{X}_n \cap B_\varepsilon| = m] = \Pr[|\mathcal{X}_n \cap B_\varepsilon| < \ell] \\ &= \Pr[|\mathcal{X}_n \cap B_\varepsilon| < \alpha^{-d} \mathbb{E}[|\mathcal{X}_n \cap B_\varepsilon|]] \\ &\leq \exp\{-n\varepsilon^d(1 - \alpha^{-d})^2\}. \end{aligned}$$

The last inequality follows from the *Chernoff inequality* (see, e.g., [17, Theorem 1.1]), the application of which is made possible due to the α^{-d} factor.

We now focus on showing that $\lim_{n \rightarrow \infty} \sigma(\ell, n) = 0$. For any two integers n, m such that $\ell \leq m \leq n$ we have that

$$\begin{aligned} \Pr[\mathcal{G}(\mathcal{Y}_m^\varepsilon; \alpha r'_n) \notin \mathcal{A}] &\stackrel{(1)}{=} \Pr[\mathcal{G}(\mathcal{X}_m; \alpha \varepsilon^{-1} r'_n) \notin \mathcal{A}] \\ &\stackrel{(2)}{\leq} \Pr[\mathcal{G}(\mathcal{X}_m; r'_m) \notin \mathcal{A}]. \end{aligned}$$

where the transitions are made possible due to (1) a scaling of the graph from $[0, \varepsilon]^d$ to $[0, 1]^d$; (2) the monotonicity of \mathcal{A} and the fact that $r'_m \leq \frac{\alpha}{\varepsilon} r'_n$. To show that indeed $r'_m \leq \frac{\alpha}{\varepsilon} r'_n$, note that $\alpha \varepsilon^{-1} > 1$ and that $\lim_{n \rightarrow \infty} r'_n = 0$. Thus,

$$\begin{aligned} r'_m \leq r'_\ell &= \alpha^{-1} r_n = \alpha^{-1} \gamma \left(\frac{\log \alpha^{-d} \varepsilon^d n}{\alpha^{-d} \varepsilon^d n} \right)^{1/d} \\ &= \varepsilon^{-1} \gamma \left(\frac{\log \alpha^{-d} \varepsilon^d n}{n} \right)^{1/d} \leq \varepsilon^{-1} \gamma \left(\frac{\log n}{n} \right)^{1/d} = \frac{\alpha}{\varepsilon} r'_n. \end{aligned}$$

Furthermore, set $m^* = \operatorname{argmax}_{m \in [\ell, n]} (\Pr[\mathcal{G}(\mathcal{X}_m, r'_m) \notin \mathcal{A}])$. It follows that

$$\begin{aligned} \sigma(\ell, n) &= \sum_{m=\ell}^n \Pr[\mathcal{G}(\mathcal{Y}_m^\varepsilon; \alpha r'_n) \notin \mathcal{A}] \cdot \Pr[|\mathcal{X}_n \cap B_\varepsilon| = m] \\ &\leq \sum_{m=\ell}^n \Pr[\mathcal{G}(\mathcal{X}_m, r'_m) \notin \mathcal{A}] \cdot \Pr[|\mathcal{X}_n \cap B_\varepsilon| = m] \\ &\leq \Pr[\mathcal{G}(\mathcal{X}_{m^*}, r'_{m^*}) \notin \mathcal{A}] \sum_{m=\ell}^n \Pr[|\mathcal{X}_n \cap B_\varepsilon| = m] \\ &= \Pr[\mathcal{G}(\mathcal{X}_{m^*}, r'_{m^*}) \notin \mathcal{A}] \cdot \Pr[|\mathcal{X}_n \cap B_\varepsilon| \geq \ell] \\ &\leq \Pr[\mathcal{G}(\mathcal{X}_{m^*}, r'_{m^*}) \notin \mathcal{A}]. \end{aligned}$$

Note that $\lim_{n \rightarrow \infty} \Pr[\mathcal{G}(\mathcal{X}_{m^*}, r'_{m^*}) \notin \mathcal{A}] = 0$, which concludes the proof. \blacksquare

The following are the SRGG and REGG equivalents of Lemma 1. The proofs are very similar to that of the previous lemma, and are omitted due to lack of space.

Lemma 2. *Let \mathcal{A} be a monotone property and let $\mathcal{G}_n = \mathcal{G}^{\text{soft}}(\mathcal{X}_n; r_n; \phi_n)$ be an SRGG such that $\mathcal{G}_n \in \mathcal{A}$ a.s., where r_n, ϕ_n are as defined in Theorem 2. Then \mathcal{G}_n is localizable for \mathcal{A} .*

Lemma 3. *Let \mathcal{A} be a monotone property and let $\mathcal{G}_n = \mathcal{G}^{\text{embed}}(\mathcal{X}_n; p_n)$ be an REGG such that $\mathcal{G}_n \in \mathcal{A}$ a.s., where p_n is non-decreasing. Then \mathcal{G}_n is localizable for \mathcal{A} .*

V. PROPERTIES OF RGGs IN GENERAL DOMAINS VIA TESSELLATION

In the previous section we considered three models of RGGs defined over the *convex* domain $[0, 1]^d$. We discussed the necessary conditions such that random graphs will be localizable for any monotone property \mathcal{A} . In this section we consider the specific monotone properties of *connectivity* and *bounded stretch* for general domains.

A region $\Gamma \subset [0, 1]^d$ is said to be ρ -safe for some $\rho > 0$ if $\mathcal{B}_\rho(\Gamma) \subset [0, 1]^d$, namely if the Minkowski sum of Γ with a ball of radius ρ is contained in $[0, 1]^d$.

A. Connectivity

Denote by $\mathcal{A}_{\text{conn}}$ the connectivity property. We show that for any random graph \mathcal{G}_n which is an RGG, SRGG or REGG that is localizable for $\mathcal{A}_{\text{conn}}$ it also holds that \mathcal{G}_n is connected over any ρ -safe region $\Gamma \subset [0, 1]^d$. Note that we make no additional assumptions on Γ in this section.

Theorem 5. Let $\Gamma \subset [0, 1]^d$ be a ρ -safe region for some constant $\rho > 0$ independent of n and let \mathcal{G}_n be a random graph that is localizable for $\mathcal{A}_{\text{conn}}$. Then any two points $x, y \in \Gamma \cap \mathcal{X}_n$ that lie in the same connected component of Γ are connected in $\mathcal{G}_n(\mathcal{B}_\rho(\Gamma))$ a.s..

In the proof of Theorem 5 we will place two partially-overlapping grids over Γ and use the localization of \mathcal{G}_n in each grid cell (see Fig. 1). We now proceed to define the grids and state several of their properties which, in turn, will allow us to formally prove Theorem 5.

Let \mathbb{H}_ε be a grid partition of $[0, 1]^d$ into axis-aligned hypercubes with side length of $\varepsilon = \frac{2}{3\sqrt{d}}\rho$. Furthermore, denote by $\mathbb{H}_\varepsilon(\Gamma)$ the subset of cells of \mathbb{H}_ε whose intersection with Γ is non-empty. Namely, $\mathbb{H}_\varepsilon(\Gamma) = \{H \in \mathbb{H}_\varepsilon \mid H \cap \Gamma \neq \emptyset\}$. Let $\tilde{\mathbb{H}}_\varepsilon$ be a grid partition of $[0, 1]^d$ into axis-aligned hypercubes with side length of ε obtained by shifting \mathbb{H}_ε by $\varepsilon/2$ along every axis and let $\tilde{\mathbb{H}}_\varepsilon(\Gamma) = \{H \in \tilde{\mathbb{H}}_\varepsilon \mid H \cap \mathbb{H}_\varepsilon(\Gamma) \neq \emptyset\}$. We have the following claim.

Claim 1. Let $H \in \mathbb{H}_\varepsilon(\Gamma) \cup \tilde{\mathbb{H}}_\varepsilon(\Gamma)$. Then $H \subset \mathcal{B}_\rho(\Gamma)$.

We introduce some more terminology. Every two cells $H, H' \in \mathbb{H}_\varepsilon(\Gamma)$ are called *neighbors* if they share a $(d-1)$ -dimensional face. We now consider a refinement of each grid cell H of $\mathbb{H}_\varepsilon(\Gamma)$ (or of $\tilde{\mathbb{H}}_\varepsilon(\Gamma)$) into 2^d sub-cells obtained by splitting H by two along each axis through the middle point of H . This induces the set of (refined) grid cells $\mathbb{H}_{\varepsilon/2}(\Gamma)$ (or $\tilde{\mathbb{H}}_{\varepsilon/2}(\Gamma)$, respectively). Note that the number of cells in $\mathbb{H}_{\varepsilon/2}(\Gamma)$ and $\tilde{\mathbb{H}}_{\varepsilon/2}(\Gamma)$ is fixed for the given d, ρ, Γ , and does not depend on n .

Claim 2. Let $H \in \mathbb{H}_\varepsilon(\Gamma) \cup \tilde{\mathbb{H}}_\varepsilon(\Gamma)$. Then $\mathcal{X}_n \cap H \neq \emptyset$, a.s..

We are ready for the main proof.

Proof (Theorem 5): Recall that \mathcal{G}_n is localizable for $\mathcal{A}_{\text{conn}}$. As $\bigcup_{H \in \mathbb{H}_\varepsilon(\Gamma)} H \subset \mathcal{B}_\rho(\Gamma)$, and since x and y are in the same connected component of Γ , there exists a sequence of hypercubes $H_1, \dots, H_k \in \mathbb{H}_\varepsilon(\Gamma)$ such that (i) $x \in H_1$, (ii) $y \in H_k$ and (iii) H_i and H_{i+1} are neighbors for $1 \leq i < k$. By Claim 1 each H_i is contained in $\mathcal{B}_\rho(\Gamma)$.

Claim 2 ensures, using the fact that Γ is ρ -safe, that each H_i contains a vertex of \mathcal{G}_n a.s. Let $x = x_1, \dots, x_k = y$ denote such a set of vertices where $x_i \in H_i$. We will show (using the localization of monotone properties) that x_i and x_{i+1} are connected in $\mathcal{G}_n(\mathcal{B}_\rho(\Gamma))$ which will conclude our proof.

Let $\tilde{H} \in \tilde{\mathbb{H}}_\varepsilon(\Gamma)$ be a hypercube that intersects both H_i and H_{i+1} (there are always 2^{d-1} such hypercubes). By Claim 2, both $\tilde{H} \cap H_i$ and $\tilde{H} \cap H_{i+1}$ contain a vertex of \mathcal{G}_n a.s., since both of these intersection represent hypercubes in $\mathbb{H}_{\varepsilon/2}(\Gamma)$. Let z_i and z_{i+1} be these vertices, respectively (see Fig. 2).

Now, using Lemmas 1-3 we have that x_i and z_i are connected in H_i , that z_i and z_{i+1} are connected in \tilde{H} , and that z_{i+1} and x_{i+1} are connected in H_{i+1} a.s. This must hold for every $1 \leq i < k$ in order to ensure that x and y are connected in $\mathcal{G}_n(\mathcal{B}_\rho(\Gamma))$. Due to the fact that k can be at most the number of cells in $\mathbb{H}_\varepsilon(\Gamma_\rho)$, which is independent of n , we deduce that




Fig. 1. Visualization of Γ (green), $\mathcal{B}_\rho(\Gamma)$ (purple) and the grid \mathbb{H}_ε used for the proof of Theorem 5. The boundary of the set of grid cells $\mathbb{H}_\varepsilon(\Gamma)$ is depicted using dark red lines.




Fig. 2. Visualization of the proof of Theorem 5. Hypercubes H_i, H_{i+1} and \tilde{H} of side length ε are depicted in solid blue lines and dashed red lines, respectively. A path connecting $x_i \in H_i$ to $x_{i+1} \in H_{i+1}$ via intermediate points $z_i \in H_i \cap \tilde{H}$ and $z_{i+1} \in H_{i+1} \cap \tilde{H}$ is depicted by a purple line.

indeed x, y are connected in $\mathcal{G}_n(\mathcal{B}_\rho(\Gamma))$ a.s. ■

B. Bounded stretch

Given $\zeta \geq 1$ denote by $\mathcal{A}_{\text{str}}^\zeta$ the property indicating that a given geometrically-embedded graph has a *bounded stretch* of ζ , for any two vertices. Formally, let \mathcal{G} be a graph defined over a vertex set $X \subset [0, 1]^d$. The notation $\mathcal{G} \in \mathcal{A}_{\text{str}}^\zeta$ indicates that for every $x, y \in X$ it holds that $\text{dist}(\mathcal{G}, x, y) \leq \zeta \|x - y\|_2$. The proof of the following theorem is very similar to that of Theorem 5.

Theorem 6. Let $\Gamma \subset [0, 1]^d$ be a ρ -safe region for some constant $\rho > 0$ independent of n . Let \mathcal{G}_n be a random graph that is localizable for $\mathcal{A}_{\text{str}}^\zeta$, for some $\zeta \geq 1$. Additionally, let $x, y \in \mathcal{X}_n$ be two points that lie in the same connected component of Γ . Then $\text{dist}(\mathcal{G}_n(\mathcal{B}_\rho(\Gamma)), x, y) \leq \zeta \|x - y\|_\Gamma + o(1)$ a.s., where $\|x - y\|_\Gamma$ denotes the length of the shortest path between x and y that is fully contained in Γ .

VI. APPLICATION TO SAMPLING-BASED MOTION PLANNING

We now move to the setting of motion planning in which a robot operates in the configuration space $\mathcal{C} = [0, 1]^d$, and whose free space is denoted by $\mathcal{F} \subseteq \mathcal{C}$. Recall that the problem consists of finding a continuous path between two configurations (points) $s, t \in \mathcal{F}$, that is fully contained in \mathcal{F} .

The reason why we cannot apply results on RGGs to motion planning directly is that \mathcal{F} is not the full hypercube $[0, 1]^d$ but rather could be a geometrically and topologically very complicated subset of this hypercube. However, the localization-tessellation approach that we have devised enables us to fairly directly adapt results from the theory of RGGs to this more involved setting, as we do in this section.

Specifically, we start by introducing the Soft-PRM algorithm and Embedded-PRM algorithms which are extensions of SRGG and REGG to the setting of motion planning. We remark that Soft-PRM is very similar to a technique that was studied experimentally by McMahan et al. [43]; here we provide theoretical analysis for it. We then continue to provide proofs for probabilistic completeness and asymptotic optimality of these methods.

A. Motion-planning algorithms

We introduce the Soft-PRM algorithm. The description of PRM and Embedded-PRM immediately follow, as they are special cases of Soft-PRM. Recall that SRGG is defined for a connection radius r_n and the function $\phi_n : \mathbb{R}^+ \rightarrow [0, 1]$: two vertices $x, y \in \mathcal{X}_n$ for which $\|x - y\|_2 \leq r_n$ are connected with an edge with probability $\phi_n(\|x - y\|_2)$.

We use the following standard procedures: `sample(n)` returns n configurations that are sampled uniformly and randomly from \mathcal{C} ; `nearest_neighbors(x, V, r)` returns all the configurations from V that are found within a distance of r from x ; `collision_free(x, y)` tests whether the straight-line segment connecting x and y is contained in \mathcal{F} ; `random_variable()` selects uniformly at random a real number in the range $[0, 1]$.

The *preprocessing phase* of Soft-PRM is described in Alg. 1. In lines 1-4, n configurations are sampled (note that this slightly differs from some PRM descriptions in which the samples are assumed to be collision free) and for each sample, Soft-PRM retrieves the neighboring samples which are within a distance of at most r_n from it. For each sample point x and each candidate neighbor y it decides with probability $\phi_n(\|x - y\|_2)$ whether to attempt the connection (lines 5-6). If this is the case, the edge (x, y) is tested for being collision free (line 7), and added accordingly to E .

The (standard) PRM and Embedded-PRM are identical to Alg. 1 using the parameters r_n and $\phi_n = 1$ for PRM and $r_n = \infty$ and $\phi_n = p_n$ for Embedded-PRM. Note that in the implementation of Embedded-PRM there is no need to maintain a nearest-neighbor data structure (line 3) as every pair of vertices $x, y \in \mathcal{X}_n$ is chosen with probability p_n .

In the *query stage* each of the three algorithms is given two configurations s, t , which are then connected to

Algorithm 1 Soft-PRM(n, r_n, ϕ_n)

```

1:  $V \leftarrow \{x_{\text{init}}\} \cup \text{sample}(n)$ ;  $E \leftarrow \emptyset$ ;  $\mathcal{G} \leftarrow (V, E)$ 
2: for all  $x \in V$  do
3:    $U \leftarrow \text{nearest\_neighbors}(x, V, r_n)$ 
4:   for all  $y \in U$  do
5:      $\xi \leftarrow \text{random\_variable}()$ 
6:     if  $\xi \leq \phi_n(\|x - y\|_2)$  then
7:       if collision_free( $x, y$ ) then
8:          $E \leftarrow E \cup (x, y)$ 
return  $\mathcal{G}$ 

```

their neighbors in the underlying roadmap by executing `nearest_neighbors` with the connection radius $r_{\text{query}} = \gamma \left(\frac{\log n}{n}\right)^{1/d}$, where $\gamma > \gamma^* = 2(2d\theta_d)^{-1/d}$. Naturally, every connection is tested for collision. Finally, the underlying graph is searched for the shortest path from s to t and the respective path in \mathcal{F} is returned (if exists).

Observation 1. Denote by \mathcal{G}_n the Soft-PRM roadmap produced for n samples and the connection radius r_n . Then $\mathcal{G}_n = \mathcal{G}_n^{\text{soft}}(\mathcal{X}_n; r_n; \phi_n) \cap \mathcal{F}$. The same applies for the relation of the underlying roadmaps of PRM, Embedded PRM, and RGG, REGG, respectively.

B. Probabilistic completeness

Let (\mathcal{F}, s, t) be a motion-planning problem that consists of the free space $\mathcal{F} \subset [0, 1]^d$, and $s, t \in \mathcal{F}$ are the start and target configurations, respectively. We provide the definition of *probabilistic completeness* and state the conditions under which the aforementioned algorithms possess this property.

Definition 6 [31]. Let $\sigma : [0, 1] \rightarrow \mathcal{F}$ be a continuous path, and let $\delta > 0$. The path σ is δ -robust if $\mathcal{B}_\delta(\sigma) \subseteq \mathcal{F}$.

Definition 7 [31]. A motion-planning problem (\mathcal{F}, s, t) is *δ -robustly feasible* if there exists a δ -robust path σ connecting s to t , for some fixed $\delta > 0$.

Definition 8 [31]. A planner ALG is probabilistically complete if for any robustly-feasible (\mathcal{F}, s, t) , the probability that ALG finds a solution with n samples converges to 1 as n tends to ∞ .

Lemma 4. Let $\text{ALG} \in \{\text{PRM}, \text{Soft-PRM}, \text{Embedded-PRM}\}$ with a selection of parameters for which the corresponding random graph \mathcal{G}_n be localizable for connectivity. Then ALG is probabilistically complete.

Proof: Suppose that (\mathcal{F}, s, t) is robustly feasible. By definition, there exists a path σ connecting s to t and $\delta > 0$ for which $\mathcal{B}_\delta(\sigma) \subseteq \mathcal{F}$.

Let $\delta' := \delta/2$ and let $\Gamma := \mathcal{B}_{\delta'}(\sigma)$. Note that $s, t \in \Gamma$ and also $\mathcal{B}_{\delta'}(\Gamma) = \mathcal{B}_\delta(\sigma) \subseteq \mathcal{F} \subset [0, 1]^d$, which implies that Γ is δ' -safe. By the fact that ALG is localizable for connectivity, and by Theorem 5, we have that for every $x, y \in \mathcal{X}_n \cap \Gamma$ it follows that x, y are connected in $\mathcal{G}_n(\mathcal{B}_{\delta'}(\Gamma))$ a.s.

It remains to show that during the query stage s and t are connected to $\mathcal{G}_n(\mathcal{B}_{\delta'}(\Gamma))$. It is not hard to verify that $\mathcal{X}_n \cap \mathcal{B}_{r_{\text{query}}}(s) \neq \emptyset, \mathcal{X}_n \cap \mathcal{B}_{r_{\text{query}}}(t) \neq \emptyset$, a.s., which implies connectivity. ■

Theorem 7. Recall that $\gamma^* = 2(2d\theta_d)^{-1/d}$ and that $d \geq 2$. Then the following algorithms are probabilistically complete:

- (i) PRM(r_n), where $r_n = \gamma \left(\frac{\log n}{n}\right)^{1/d}$, for $\gamma > \gamma^*$;
- (ii) Soft-PRM($r_n; \phi_n$), where $r_n = \gamma \left(\frac{\log n}{n}\right)^{1/d}$, for $\gamma = \alpha\gamma^*$, $\alpha > (d+1)^{1/d}$, and $\phi_n(z) = 1 - z/r_n$, for any $z \in \mathbb{R}^+$;
- (iii) Embedded-PRM(p_n), where $p_n = \omega \left(\frac{\log^d n}{n}\right)$.

Item (i) results from combining Theorem 1 with the localization lemma for RGGs (i.e., Lemma 1), and Lemma 4. Items (ii),(iii) similarly follow.

Remark 1. The connection radius in (i) is smaller by a factor of $2^{-1/d}$ than the one obtained by Janson et al. [29], and smaller by a factor of $2^{-1/d}(d+1)^{-1/d}$ than the connection radius proposed by Karaman and Frazzoli [31] when no obstacles are present. We also mention that, similarly to these two works, r_n can be reduced by a factor of $|\mathcal{F}|^{1/d}$, with a slight modification to Theorem 5.

C. Asymptotic (near-)optimality

Given a path σ denote its length by $|\sigma|$. We define the property of asymptotic near-optimality and state the conditions under which PRM and Embedded-PRM have this property.

Definition 9. Suppose that (\mathcal{F}, s, t) is robustly feasible. A path σ^* connecting s to t is *robustly optimal* if it is a shortest path for which the following holds: for any $0 < \varepsilon$ there exists a δ -robust path σ such that $|\sigma| \leq (1+\varepsilon)|\sigma^*|$ for some fixed $\delta > 0$.

Definition 10. A sampling-based planner ALG is *asymptotically ζ -optimal*, for a given $\zeta \geq 1$, if for every robustly-feasible problem (\mathcal{F}, s, t) it follows that $|\sigma_n| \leq \zeta|\sigma^*| + o(1)$ a.s., where σ_n denotes the solution returned by ALG with n samples. A planner that is asymptotically 1-optimal is simply called *asymptotically optimal*.

Theorem 8. For $d \geq 2$ we have the following results:

- (i) PRM(r_n) is asymptotically ζ -optimal for $r_n = \gamma \left(\frac{\log n}{n}\right)^{1/d}$, where $\gamma > \gamma^*$, and some constant ζ ;
- (ii) Embedded PRM(p_n) is asymptotically optimal for $p_n = \omega \left(\frac{\log^d n}{n}\right)$;

The proof of this theorem is similar in nature to that of Theorem 7 (but follows from Theorem 3 for (i) and Theorem 4 for (ii)), and hence omitted due to lack of space.

VII. EVALUATION

In this section we present experiments demonstrating the behavior of RGGs and SRGGs in the absence and in the presence of obstacles. We then proceed to compare the Soft-PRM and PRM algorithms. For each model, and each algorithm,




Fig. 3. (a) Connectivity and (b) stretch in the absence of obstacles.

we use the minimal parameters that are required in order to ensure connectivity. For our experiments we used the Open Motion Planning Library (OMPL) [66] with Randomly Transformed Grids (RTG) [2] as our nearest-neighbor (NN) data structure. RTG were shown to outperform other NN libraries for several motion-planning algorithms [33]. All experiments were run on a 2.8GHz Intel Core i7 processor with 8GB of memory. Results are averaged over 100 runs, and computed for dimensions $d = 2, 6, 9$ and 12 . Additionally, when results for different dimensions behave similarly, we present only plots for $d = 2$ and $d = 12$.

Connectivity and stretch in the unit cube. We begin by reporting the number of nodes that are *not* in the largest connected component (CC) for RGG and SRGG in the absence of obstacles. Clearly, when the graph is connected, this number is zero. One can see (Fig. 3a) that as the number of nodes increases, the number of nodes not in the largest CC approaches zero. Additionally, the two models exhibit very similar trends.

We continue to assess how increasing the number of nodes affects the stretch of the graphs. For each such n , we sampled $m = 50$ vertices and computed the stretch for every pair of sampled vertices. We then report on the maximal stretch obtained among all $O(m^2)$ pairs of nodes which gives a rough approximation of the average stretch of the graph. Results are depicted in Fig. 3b. Observe that typically the stretch decreases as the number of nodes increases and that RGG and SRGG behave very similarly. Notice that each point along the plot is an average of 100 different runs. In addition, the RGG's used for each time-step are independent. Thus, it is probable that a graph drawn with n_1 vertices is connected while a graph drawn with $n_2 > n_1$ vertices is disconnected.

Connectivity and stretch of RGGs in general domains. The set of experiments come to demonstrate Theorems 5 and 6. Namely, that the asymptotic behavior of RGGs with respect to connectivity and stretch is maintained in the presence of obstacles. We constructed the following toy-scenario where we subdivided the d -dimensional unit hypercube by halving it along each axis. In the center of each one of the 2^d sub-cubes, we inserted an axis-aligned hypercube as an obstacle. The size of the obstacle was chosen such that the obstacles covered 25% of the unit hypercube. See Figure 4 for a visualization in two and three dimensions.




Fig. 4. Visualization of toy scenario for (a) two and (b) three dimensions. Obstacles are depicted in blue.




Fig. 5. (a) Connectivity and (b) stretch for RGGs in the toy scenarios.

We report on the results for RGGs (Fig. 5) and note that similar results were observed for SRGGs. Stretch was computed between the origin $(0, \dots, 0)$ and the center $(0.5, \dots, 0.5)$. Observe that for all dimensions, the graph is asymptotically connected and the stretch tends to one.

Motion-planning algorithms. Finally, we compare PRM and Soft-PRM as sampling-based planners for rigid-body motion planning on the Home scenario (Fig. 6b) provided by the OMPL distribution. This six-dimensional configuration space, SE_3 , includes both translational and rotational degrees of freedom. Thus, it is not clear if the theoretic results presented in this paper still hold in this non-Euclidean space.

To apply the results, a key question one has to address is how to choose the connection radius when using a non-Euclidean metric. Let x be a point sufficiently far from the boundary and let $r_n = \gamma (\log n/n)^{1/d}$ be the connection radius used. When using the Euclidean metric, the average number of neighbors of x is $\text{nbr}(n) = (2^{d-1}/d) \cdot \log n$. Thus, for each value of n , we sampled 100 random points and, for each one, computed the radius r for which the point had $\text{nbr}(n)$ neighbors. Finally, we used the average value over all such points in the experiments.

Figure 6a presents the cost of the solution produced by each algorithm as a function of the running time. Similar to the previous tests, both algorithms exhibit similar behavior, and the cost obtained approaches the optimum as the number of nodes increases.

VIII. DISCUSSION

We conclude this paper by describing a connection between Soft-PRM and approximate *nearest-neighbor (NN) search* in sampling-based motion planning. We then proceed to describe future research directions that follow from our work.




Fig. 6. (a) Average quality obtained by the PRM and Soft-PRM algorithms as a function of the running times in the Home scenario (b). Error bars in (a) denote the 20'th and 80'th percentile. Cost is normalized such that the unit cost represents the optimum.

Approximate NN search in motion planning. NN search is a key ingredient in the implementation of sampling-based planners (see, e.g., line 3 in Alg. 1). Typically, exact NN computation, where all the neighbors of a query point in a given area are reported, tends to be slow in high dimensions, due to the “curse of dimensionality” [26]. Thus, most implementations of motion planners involve approximate NN libraries (see, e.g., [7, 33, 45]), which are only guaranteed to return a subset of the neighbors of a given query point (see, e.g., [2, 10, 19, 28]).

However, existing proofs of probabilistic completeness and asymptotic optimality of standard planners (see, e.g., [29, 31, 32]) assume that NNs are computed exactly. Without these assumptions, the proofs no longer hold (although it may be possible to modify them to take this into account).

The analysis given in this paper bridges this gap: PRM, when implemented with approximate NN search, can be modeled as a Soft-PRM. Thus, the former algorithm is probabilistically complete by using the probabilistic completeness of the latter (see Theorem 7).

Future work. The literature of RGGs is rich and encompasses many models which were not addressed in this work due to lack of space (see, e.g., [12, 9, 68]). Such models can be used to analyze existing planners and might lead to the development of novel planners.

In this work our focus was on Euclidean configuration spaces and the standard Euclidean distance. We mention that several works on RGGs consider different metrics in the Euclidean space (see, e.g., [5, 49]). Such results can be imported to the setting of motion planning using our framework, with slight modification of the proofs.

Perhaps a more urgent issue involves the analysis of exiting planners in complex configuration spaces. To the best of our knowledge the behavior of standard planners such as PRM and RRT* is not well understood for non-Euclidean spaces, even for the simple case of a rigid-body robot translating and rotating in a three-dimensional workspace. We believe that several results involving RGGs in complex domains can shed light on this question. For instance, Penrose [48] considers the case where points are sampled on a torus, whereas Penrose and Yukich [54] study the setting of points on a manifold embedded in Euclidean space.

REFERENCES

- [1] R. R. Abhay G. Bhatt. On a random directed spanning tree. *Advances in Applied Probability*, 36(1):19–42, 2004.
- [2] D. Aiger, H. Kaplan, and M. Sharir. Reporting neighbors in high-dimensional Euclidean space. *SIAM J. Comput.*, 43(4):1363–1395, 2014.
- [3] R. Alterovitz, S. Patil, and A. 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.
- [4] N. M. Amato, O. B. Bayazit, L. K. Dale, C. Jones, and D. Vallejo. OBPRM: A obstacle-based PRM for 3D workspaces. In *International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, pages 155–168, 1998.
- [5] M. J. Appel and R. P. Russo. The connectivity of a graph on uniform points on $[0, 1]^d$. *Statistics & Probability Letters*, 60(4):351–357, 2002.
- [6] O. Arslan and P. 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.
- [7] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Y. Wu. An optimal algorithm for approximate nearest neighbor searching in fixed dimensions. *J. ACM*, 45(6):891–923, 1998.
- [8] A. Bagchi and S. Bansal. Nearest-neighbor graphs on random point sets and their applications to sensor networks. In *ACM Symposium on Principles of Distributed Computing*, pages 434–434, New York, NY, USA, 2008. ACM.
- [9] P. N. Balister, B. Bollobás, A. Sarakar, and M. Walters. A critical constant for the k-nearest-neighbour model. *Advances in Applied Probability*, 41(1):1–12, 2009.
- [10] J. L. Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18:509–517, 1975.
- [11] M. Bradonjic, R. Elsässer, T. Friedrich, T. Sauerwald, and A. Stauffer. Efficient broadcast on random geometric graphs. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1412–1421, 2010.
- [12] N. Broutin, L. Devroye, N. Fraiman, and G. Lugosi. Connectivity threshold of bluetooth graphs. *Random Struct. Algorithms*, 44(1):45–66, 2014.
- [13] J. Canny. *The complexity of robot motion planning*. MIT press, 1988.
- [14] H. Choset, K. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun. *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press, 2005.
- [15] J. Díaz, D. Mitsche, G. Perarnau, and X. P. Giménez. On the relation between graph distance and Euclidean distance in random geometric graphs. *Advances in Applied Probability*, 2015. to appear.
- [16] A. Dobson and K. E. Bekris. Sparse roadmap spanners for asymptotically near-optimal motion planning. *International Journal of Robotics Research*, 33(1):18–47, 2014.
- [17] D. P. Dubhashi and A. Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.
- [18] R. B. Ellis, J. L. Martin, and C. Yan. Random geometric graph diameter in the unit ball. *Algorithmica*, 47(4):421–438, 2007.
- [19] J. H. Friedman, J. L. Bentley, and R. A. Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Trans. Math. Softw.*, 3(3):209–226, 1977.
- [20] T. Friedrich, T. Sauerwald, and A. Stauffer. Diameter and broadcast time of random geometric graphs in arbitrary dimensions. *Algorithmica*, 67(1):65–88, 2013.
- [21] A. Frieze and W. Pegden. Traveling in randomly embedded random graphs. *arXiv:1411.6596*, 2014.
- [22] J. D. Gammell, S. S. Srinivasa, and T. 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.
- [23] R. Geraerts and M. H. Overmars. Creating high-quality paths for motion planning. *International Journal of Robotics Research*, 26(8):845–863, 2007.
- [24] E. N. Gilbert. Random plane networks. *Journal of the Society for Industrial & Applied Mathematics*, 9(4):533–543, 1961.
- [25] P. Gupta and P. Kumar. Critical power for asymptotic connectivity in wireless networks. In W. M. McEneaney, G. Yin, and Q. Zhang, editors, *Stochastic Analysis, Control, Optimization and Applications*, Systems & Control: Foundations & Applications, pages 547–566. Birkhuser Boston, 1999.
- [26] S. Har-Peled, P. Indyk, and R. Motwani. Approximate nearest neighbor: Towards removing the curse of dimensionality. *Theory of Computing*, 8(1):321–350, 2012.
- [27] D. Hsu, J.-C. Latombe, and R. Motwani. Path planning in expansive configuration spaces. *International Journal of Computational Geometry and Applications*, 9(4/5), 1999.
- [28] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *ACM Symposium on the Theory of Computing (STOC)*, pages 604–613, 1998.
- [29] L. Janson, E. Schmerling, A. A. Clark, and M. 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.
- [30] S. Karaman and E. Frazzoli. Optimal kinodynamic motion planning using incremental sampling-based methods. In *IEEE Conference on Decision and Control (CDC)*, pages 7681–7687, 2010.

- [31] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *International Journal of Robotics Research*, 30(7):846–894, 2011.
- [32] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high dimensional configuration spaces. *IEEE Transactions on Robotics*, 12(4):566–580, 1996.
- [33] M. Kleinbort, O. Salzman, and D. Halperin. Efficient high-quality motion planning by fast all-pairs r-nearest-neighbors. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2985–2990, 2015.
- [34] G. Kozma, Z. Lotker, and G. Stupp. On the connectivity threshold for general uniform metric spaces. *Information Processing Letters*, 110(10):356–359, 2010.
- [35] J. J. Kuffner and S. M. LaValle. RRT-Connect: An efficient approach to single-query path planning. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 995–1001, 2000.
- [36] A. M. Ladd and L. E. Kavraki. Fast tree-based exploration of state space for robots with dynamics. In *Algorithmic Foundations of Robotics*, pages 297–312, 2004.
- [37] J.-C. Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, Norwell, MA, USA, 1991.
- [38] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, Cambridge, U.K., 2006.
- [39] Y. Li, Z. Littlefield, and K. E. Bekris. Sparse methods for efficient asymptotically optimal kinodynamic planning. In *Algorithmic Foundations of Robotics*, pages 263–282, 2014.
- [40] J.-M. Lien, S. L. Thomas, and N. M. Amato. A general framework for sampling on the medial axis of the free space. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4439–4444, 2003.
- [41] Z. Littlefield, Y. Li, and K. E. Bekris. Efficient sampling-based motion planning with asymptotic near-optimality guarantees for systems with dynamics. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1779–1785, 2013.
- [42] R. Luna, I. A. Şucan, M. Moll, and L. E. Kavraki. Any-time solution optimization for sampling-based motion planning. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 5053–5059, 2013.
- [43] T. McMahan, S. A. Jacobs, B. Boyd, L. Tapia, and N. M. Amato. Local randomization in neighbor selection improves PRM roadmap quality. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4441–4448, 2012.
- [44] D. Mitsche and G. Perarnau. On the treewidth and related parameters of random geometric graphs (stacs). In *International Symposium on Theoretical Aspects of Computer Science*, pages 408–419, 2012.
- [45] M. Muja and D. G. Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. In *International Conference on Computer Vision Theory and Applications (VISSAPP)*, pages 331–340. INSTICC Press, 2009.
- [46] S. Muthukrishnan and G. Pandurangan. Thresholding random geometric graph properties motivated by ad hoc sensor networks. *Journal of Computer and System Sciences*, 76(7):686–696, 2010.
- [47] O. Nechushtan, B. Raveh, and D. 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.
- [48] M. D. Penrose. The longest edge of the random minimal spanning tree. *The annals of applied probability*, pages 340–361, 1997.
- [49] M. D. Penrose. On k -connectivity for a geometric random graph. *Random Structures & Algorithms*, 15(2):145–164, 1999.
- [50] M. D. Penrose. *Random geometric graphs*. Oxford University Press, 2003.
- [51] M. D. Penrose. Connectivity of soft random geometric graphs. *arXiv:1311.3897*, 2013.
- [52] M. D. Penrose and A. Wade. Random directed and on-line networks. In W. S. Kendall and I. Molchanov, editors, *New Perspectives in Stochastic Geometry*, pages 248–264. Oxford University Press, 2010.
- [53] M. D. Penrose and A. R. Wade. Limit theorems for random spatial drainage networks. *Advances in Applied Probability*, 42(3):659–688, 2010.
- [54] M. D. Penrose and J. E. Yukich. Limit theory for point processes in manifolds. *The Annals of Applied Probability*, 23(6):2161–2211, 2013.
- [55] A. Perez, R. Platt, G. Konidaris, L. Kaelbling, and T. Lozano-Pérez. LQR-RRT*: Optimal sampling-based motion planning with automatically derived extension heuristics. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2537–2542, 2012.
- [56] B. Raveh, A. Enosh, and D. Halperin. A little more, a lot better: Improving path quality by a path-merging algorithm. *IEEE Transactions on Robotics*, 27(2):365–371, 2011.
- [57] J. H. Reif. Complexity of the movers problem and generalization. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 421–427, 1979.
- [58] O. Salzman and D. Halperin. Asymptotically near-optimal RRT for fast, high-quality, motion planning. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4680–4685, 2014.
- [59] O. Salzman and D. Halperin. Asymptotically-optimal motion planning using lower bounds on cost. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4167–4172, 2015.
- [60] O. Salzman, D. Shaharabani, P. K. Agarwal, and D. Halperin. Sparsification of motion-planning roadmaps by edge contraction. *International Journal of Robotics Research*, 33(14):1711–1725, 2014.
- [61] M. Schulte and C. Thaele. Central limit theorems for the radial spanning tree. *arXiv preprint arXiv:1412.7462*,

2014.

- [62] T. Siméon, J. Laumond, and C. Nissoux. Visibility-based probabilistic roadmaps for motion planning. *Advanced Robotics*, 14(6):477–493, 2000.
- [63] K. Solovey and D. Halperin. On the hardness of unlabeled multi-robot motion planning. In *Robotics: Science and Systems (RSS)*, 2015.
- [64] K. Solovey, O. Salzman, and D. Halperin. New perspective on sampling-based motion planning via random geometric graphs. *CoRR*, abs/1602.05460, 2016.
- [65] I. Şucan and L. Kavraki. A sampling-based tree planner for systems with complex dynamics. *IEEE Transactions on Robotics*, 28(1):116–131, 2012.
- [66] I. A. Şucan, M. Moll, and L. E. Kavraki. The open motion planning library. *IEEE Robotics & Automation Magazine*, 19(4):72–82, 2012.
- [67] C. Urmson and R. G. Simmons. Approaches for heuristically biasing RRT growth. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1178–1183, 2003.
- [68] A. R. Wade. Asymptotic theory for the multidimensional random on-line nearest-neighbour graph. *Stochastic Processes and their Applications*, 119(6):1889 – 1911, 2009.
- [69] M. Walters. Random geometric graphs. In R. Chapman, editor, *Surveys in Combinatorics 2011*, chapter 8, pages 365–401. Cambridge University Press, 2011.
- [70] D. J. Webb and J. P. van den Berg. Kinodynamic RRT*: Asymptotically optimal motion planning for robots with linear dynamics. In *IEEE Conference on Robotics and Automation (ICRA)*, pages 5054–5061, 2013.
- [71] C. Xie, J. P. van den Berg, S. Patil, and P. Abbeel. Toward asymptotically optimal motion planning for kinodynamic systems using a two-point boundary value problem solver. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4187–4194, 2015.
- [72] F. Xue and P. R. Kumar. The number of neighbors needed for connectivity of wireless networks. *Wireless Networks*, 10(2):169–181, 2004.