

# Effective Sampling for Robot Motion Planning Through the Lens of Lattices

Itai Panasoff and Kiril Solovey

Viterbi Faculty of Electrical and Computer Engineering  
Technion–Israel Institute of Technology, Haifa, Israel  
itaip@campus.technion.ac.il, kirilsol@technion.ac.il

**Abstract**—Sampling-based methods for motion planning, which capture the structure of the robot’s free space via (typically random) sampling, have gained popularity due to their scalability, simplicity, and for offering global guarantees, such as probabilistic completeness and asymptotic optimality. Unfortunately, the practicality of those guarantees remains limited as they do not provide insights into the behavior of motion planners for a finite number of samples (i.e., a finite running time). In this work, we harness lattice theory and the concept of  $(\delta, \varepsilon)$ -completeness by Tsao et al. (2020) to construct deterministic sample sets that endow their planners with strong finite-time guarantees while minimizing running time. In particular, we introduce a highly-efficient deterministic sampling approach based on the  $A_d^*$  lattice, which is the best-known geometric covering in dimensions  $\leq 21$ . Using our new sampling approach, we obtain at least an order-of-magnitude speedup over existing deterministic and uniform random sampling methods for complex motion-planning problems. Overall, our work provides deep mathematical insights while advancing the practical applicability of sampling-based motion planning. <https://github.com/MRSTechnion/lattice-sampling-mp>

## I. INTRODUCTION

Motion planning is a key ingredient in autonomous robotic systems, whose aim is computing collision-free trajectories for a robot operating in environments cluttered with obstacles [24]. Over the years, various approaches have been developed for tackling the problem, including potential fields [30], geometric methods [13], and optimization-based approaches [40, 31, 34]. In this work, we focus on sampling-based planners (SBPs), which aim to capture the structure of the robot’s free space through graph approximations that result from configuration sampling (typically in a random fashion) and connecting nearby samples. SBPs have enjoyed popularity in recent years due to their relative scalability, in terms of the number of robot degrees of freedom (DoFs), and the ease of their implementation [35].

Another key benefit is the ability of SBPs to escape local minima (unlike potential fields) and global solution guarantees (in contrast, optimization-based approaches [40], which typically provide only local guarantees). Earlier work on the theoretical foundations of SBPs has focused on deriving probabilistic completeness (PC) guarantees for methods such as PRM [19] or RRT [25, 23, 21]. PC implies that the probability of a given planner finding a solution (if one exists) converges to one as the number of samples tends to infinity. The work of Karaman and Frazzoli [18] initiated studying the quality of the solution returned by SBPs. Specifically, they

introduced the planners PRM\* and RRT\*, and proved that the solution length of those planners converges to the optimum as the number of samples tends to infinity—a property called asymptotic optimality (AO). Subsequent work has introduced even more powerful AO planners for geometric [16, 11] and dynamical [14, 28] systems.

Unfortunately, the practical relevance of the aforementioned theoretical findings remains limited due to the lack of meaningful finite-time implications. Specifically, when a solution is obtained using a finite number of samples, it is unclear to what extent its quality can be improved with additional computation time. Moreover, in cases where no solution is returned, it is uncertain whether a solution does not exist or if the algorithm simply failed to find one. Developing finite-time bounds through randomized sampling continues to be a significant challenge [7, 41].

Deterministic sampling methods such as grid sampling or Halton sequences [24], where samples are generated according to a geometric principle, can improve the performance of SBPs in practice and simplify the algorithm analysis. Specifically, some deterministic sampling procedures have a significantly lower dispersion than uniform random sampling, which implies that the former requires fewer samples to cover the search space to a desired resolution [17]. Recently, Tsao et al. [49] have leveraged deterministic sampling to disrupt the asymptotic analysis paradigm by introducing a significantly stronger notion than AO, called  $(\delta, \varepsilon)$ -completeness, that yields finite-time guarantees for PRM-based algorithms such as PRM\* [18], FMT\* [16], BIT\* [11], and GLS [33]. Informally, a *finite* sample set is  $(\delta, \varepsilon)$ -complete for a given approximation factor  $\varepsilon > 0$  and clearance parameter  $\delta > 0$ , if the corresponding planner returns a solution whose length is at most  $(1 + \varepsilon)$  times the length of the shortest  $\delta$ -clear solution. If no solution is found using a  $(\delta, \varepsilon)$ -complete sample set then no solution of clearance  $\delta$  exists.

The work of Tsao et al. [49] derived a relation between  $(\delta, \varepsilon)$ -completeness and geometric space coverage to obtain lower bounds on the number of samples necessary to achieve  $(\delta, \varepsilon)$ -completeness, as well as upper bounds accompanied with explicit (deterministic) sampling distributions. A follow-up work by Dayan et al. [4] has introduced an even more compact  $(\delta, \varepsilon)$ -complete sample distribution that is more efficient than the one proposed in [49] or rectangular grid sampling. In particular, the staggered grid [4] consists of two shifted




Fig. 1. Sample sets within a fixed disc in  $\mathbb{R}^2$ , derived from the lattices  $\mathbb{Z}^2$ ,  $D_2^*$  and  $A_2^*$ , which yield  $(\delta, \varepsilon)$ -complete guarantees for the same values of  $\delta$  and  $\varepsilon$ . The set  $\mathcal{X}_{\mathbb{Z}^2}^{\delta,\varepsilon}$  can be viewed as a tessellation of space using cubes. The set  $\mathcal{X}_{D_2^*}^{\delta,\varepsilon}$  is obtained by placing a (rescaled) standard grid, and then placing another point in the middle of each cube. The set  $\mathcal{X}_{A_2^*}^{\delta,\varepsilon}$  can be viewed as a rescaled hexagonal grid as each point is surrounded by a hexagon whose vertices are points in the set. Note that the density of  $\mathcal{X}_{\mathbb{Z}^2}^{\delta,\varepsilon}$  and  $\mathcal{X}_{D_2^*}^{\delta,\varepsilon}$  is the same, and higher than the density of  $\mathcal{X}_{A_2^*}^{\delta,\varepsilon}$ .

and rescaled copies of the rectangular grid (see Figure 1 and Figure 2).

However, the work [4] still leaves a significant gap between the lower bound in [49] and the upper bound obtained with the staggered grid. In practice, this gap limits the applicability of the  $(\delta, \varepsilon)$ -completeness theory to relatively low dimensions (up to dimension 6) due to the large number of samples currently needed to satisfy this property, which can lead to excessive running times.

**Contribution.** In this work, we develop a theoretical framework for obtaining highly-efficient  $(\delta, \varepsilon)$ -complete sample sets by leveraging the foundational theory of lattices<sup>1</sup> [3], which has been instrumental in diverse areas from number theory [43], coding theory [8], and crystallography [38]. Specifically, we show that lattices can be transformed to obtain  $(\delta, \varepsilon)$ -complete sample sets (Theorem 1) and develop tight theoretical bounds on their size (Theorem 2), which allows to compare between different sample sets qualitatively. Using this machinery, we not only refine and generalize previous results on the staggered grid [4] but also introduce a new highly efficient  $(\delta, \varepsilon)$ -complete sample set that is based on the  $A_d^*$  lattice, which is famous for its minimalist coverage properties [3]. We also initiate the study of a new property, which estimates the computational cost resulting from using a given sample set in a more informative manner than sample complexity. In particular, the property called collision-check complexity captures the amount of collision checks, which is typically a computational bottleneck.

From a practical perspective, when solving motion-planning problems using lattice-based sample sets, we show that our  $A_d^*$ -based sample sets can result in at least order-of-magnitude improvement in terms of running time over staggered-grid

samples and two orders of magnitude improvements over rectangular grids. Moreover,  $A_d^*$ -based sample sets are vastly superior in practice to the widely-used uniform random sampling, which is evident in improved running times, success rates, and solution quality.

**Organization.** In Section II we review basic definitions on motion planning and  $(\delta, \varepsilon)$ -completeness, and formally define our objectives. In Section III, we develop a general tool for transforming lattices into  $(\delta, \varepsilon)$ -complete sample sets. We obtain sample-complexity bounds for lattice-based sample sets in Section IV, and generalize those bounds to collision-check complexity in Section V. We evaluate the practical implications of our theory in Section VII, and conclude with a discussion of limitations and future directions in Section VIII.

## II. PRELIMINARIES

The motion-planning problem concerns computing a collision-free path for a robot in an environment cluttered with obstacles. We consider a holonomic robot with a configuration space  $\mathcal{C} = \mathbb{R}^d$ . The dimension  $d \geq 2$  represents the DoF and is finite. A motion planning problem is a tuple  $\mathcal{M} := (\mathcal{C}_f, q_s, q_g)$ , where  $\mathcal{C}_f \subseteq \mathcal{C}$  is the free space (the set of collision-free configurations), and  $q_s, q_g \in \mathcal{C}_f$  are the start and goal configurations, respectively. A solution for  $\mathcal{M}$  is a continuous collision-free path  $\pi : [0, 1] \rightarrow \mathcal{C}_f$  that begins at  $\pi(0) = q_s$  and ends at  $\pi(1) = q_g$ .

Two critical properties of a given path  $\pi$  for a problem  $\mathcal{M} = (\mathcal{C}_f, q_s, q_g)$ , are its length  $\ell(\pi) \geq 0$ , and its clearance. For a given value  $\delta \geq 0$ , we say that the path  $\pi$  is  $\delta$ -clear if  $\bigcup_{0 \leq t \leq 1} \mathcal{B}_\delta(\pi(t)) \subseteq \mathcal{C}_f$ , where  $\mathcal{B}_\rho(p)$  is the  $d$ -dimensional closed Euclidean ball with radius  $\rho > 0$  centered at  $p \in \mathbb{R}^d$ . We denote  $\mathcal{B}_\rho := \mathcal{B}_\rho(o)$ , where  $o$  is the origin of  $\mathbb{R}^d$ .

### A. Probabilistic roadmaps and completeness

We present a formal definition of the Probabilistic Roadmap (PRM) method [19], which constructs a discrete graph captur-

<sup>1</sup>Lattices are point sets exhibiting a regular geometric structure, which are obtained by transforming the integer lattice  $\mathbb{Z}^d$ . For instance, the aforementioned rectangular grid and the staggered grid can be viewed as lattices.

ing the connectivity of  $\mathcal{C}_f$  through sampling. Albeit sampling usually refers to a randomized process, here we consider deterministic sampling, as was recently done in [49, 4].<sup>2</sup> We emphasize that our analysis below is not confined to PRMs, and applies to various PRM-based planners, as mentioned above.

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

$$\begin{aligned} V &:= (\mathcal{X} \cup \{q_s, q_g\}) \cap \mathcal{C}_f, \\ E &:= \{\{v, u\} \in V \times V : \|v - u\| \leq r, \overline{uv} \subset \mathcal{C}_f\}. \end{aligned} \quad (1)$$

In this work, we are interested in obtaining sample sets and connection radii for PRM that achieve a desired solution quality in terms of path length. Unlike most theoretical results for SBP, which consider asymptotic guarantees, here we rely on a stronger deterministic notion.

**Definition 1** ( $(\delta, \varepsilon)$ -completeness [49]). Given a sample set  $\mathcal{X} \subset \mathcal{C}$  and connection radius  $r > 0$ , the pair  $(\mathcal{X}, r)$  is  $(\delta, \varepsilon)$ -complete for a clearance parameter  $\delta > 0$  and stretch factor  $\varepsilon > 0$ , if for every  $\delta$ -clear problem  $\mathcal{M} = (\mathcal{C}_f, q_s, q_g)$ , the graph  $G_{\mathcal{M}(\mathcal{X}, r)}$  contains a path from  $q_s$  to  $q_g$  with length at most  $(1 + \varepsilon)$  times the optimal  $\delta$ -clear length, denoted by  $\text{OPT}_\delta$ . That is, it holds that

$$\ell(G_{\mathcal{M}(\mathcal{X}, r)}, q_s, q_g) \leq (1 + \varepsilon)\text{OPT}_\delta,$$

where  $\ell(G_{\mathcal{M}(\mathcal{X}, r)}, q_s, q_g)$  denotes the length of the shortest path from  $q_s$  to  $q_g$  in the graph  $G_{\mathcal{M}(\mathcal{X}, r)}$ .

The property of  $(\delta, \varepsilon)$ -completeness has several key advantages over asymptotic notions, such as PC and AO. First, there exists (with probability 1) a *finite* sample set  $\mathcal{X}$  and radius  $r \in (0, \infty)$  that jointly guarantee  $(\delta, \varepsilon)$ -completeness. Second, if a solution is not found using a  $(\delta, \varepsilon)$ -complete pair  $(\mathcal{X}, r)$ , then no  $\delta$ -clear solution exists. That is,  $(\delta, \varepsilon)$ -completeness can be used for deterministic infeasibility proofs [26]. Third, the computational complexity of constructing a PRM graph can be tuned according to the desired values of  $\delta$  and  $\varepsilon$ .

## B. Problem definition

A  $(\delta, \varepsilon)$ -complete sample set  $\mathcal{X}$  and radius  $r > 0$  can be obtained by constructing a so-called  $\beta$ -cover [4].

**Definition 2.** For a given  $\beta > 0$ , a sample set  $\mathcal{X} \subset \mathbb{R}^d$  is a  $\beta$ -cover if for every point  $p \in \mathcal{C} = \mathbb{R}^d$  there exists a sample  $x \in \mathcal{X}$  such that  $\|p - x\| \leq \beta$ .

<sup>2</sup>With deterministic sampling, the term “probabilistic” in PRM can seem misleading. Nevertheless, we choose to stick to PRM considering its popularity and the underlying graph structure it represents, which is a key to our analysis.

Note that the coverage property above is defined with respect to the whole configuration space, rather than a specific free space. The connection between  $(\delta, \varepsilon)$ -completeness and  $\beta$ -cover is established in the following lemma.

**Lemma 1** (Completeness-cover relation [49]). Fix  $\delta > 0$  and  $\varepsilon > 0$ . Suppose that a sample set  $\mathcal{X}$  is a  $\beta^*$ -cover, where

$$\beta^* = \beta^*(\delta, \varepsilon) := \frac{\delta\varepsilon}{\sqrt{1 + \varepsilon^2}}. \quad (2)$$

Then  $(\mathcal{X}, r^*)$  is  $(\delta, \varepsilon)$ -complete, where

$$r^* = r^*(\delta, \varepsilon) := \frac{2\delta(1 + \varepsilon)}{\sqrt{1 + \varepsilon^2}}. \quad (3)$$

The lemma prescribes an approach for constructing sample sets and connection radii that satisfy the  $(\delta, \varepsilon)$ -completeness requirement. Considering that we will fix the radius  $r^* := r^*(\delta, \varepsilon)$  throughout this work, we will say that a sample set  $\mathcal{X}$  is  $(\delta, \varepsilon)$ -complete if the pair  $(\mathcal{X}, r^*)$  is  $(\delta, \varepsilon)$ -complete.

Lemma 1 still leaves a critical unresolved question: How do we find a  $(\delta, \varepsilon)$ -complete sample  $\mathcal{X}$  minimizing the computation time of the PRM graph and subsequent methods? In this work, we are interested in the following two related problems. The first deals with finding a sample set of minimal size, which can be used as a proxy for computation time.

**Problem 1** (Sample complexity). For a given  $\delta > 0, \varepsilon > 0$ , find a  $(\delta, \varepsilon)$ -complete sample set  $\mathcal{X}$  of minimal *sample complexity*, i.e.,  $|\mathcal{X} \cap \mathcal{B}_{r^*}|$ .

Previous work [49, 4] has considered a slightly different notion for sample complexity aiming to minimize the global expression  $|\mathcal{X} \cap \mathcal{C}|$  with a bounded  $\mathcal{C}$ . We believe that our local notion, within an  $r^*$ -ball, is more informative and optimistic, as typical problems have obstacles, and the solution lies in a small subset of the search space. Furthermore, it would allow us to obtain tighter analyses by exploiting methods from discrete geometry that reason about ball structures.

Previous work has introduced several sample sets and analyzed their sample complexity [49, 4]. In this work, we consider more compact sets. Moreover, we introduce a new notion that better captures the computational complexity of constructing a PRM graph. In particular, we leverage the observation that the computational complexity of sampling-based planning is typically dominated by the amount of collision checks performed [20]. Furthermore, nearest-neighbor search, which is another key contributor to the algorithm’s computational complexity, can be eliminated for deterministic samples, assuming that they have a regular structure (as for lattice-based samples, which we describe below).

Collision checks are run both on the PRM vertices and edges, where edge checks are usually performed via dense sampling of configurations along the edges and individually validating each configuration. Thus, the total number of collision checks is proportional to the total edge length of the graph. We use this observation to develop a more accurate proxy for computational complexity. In particular, we will estimate the length of edges adjacent to the origin point

$o \in \mathbb{R}^d$ , which is a vertex in all the sample sets introduced below. Moreover, due to the regularity of the sets, the attribute below is equal across all vertices (in the absence of obstacles).

**Problem 2** (Collision-check complexity). For given  $\delta > 0, \varepsilon > 0$ , find a  $(\delta, \varepsilon)$ -complete sample set  $\mathcal{X}$  of minimal *collision-check complexity*, i.e., minimizing the expression

$$CC_{\mathcal{X}} := \sum_{x \in \mathcal{X} \cap \mathcal{B}_{r,*}} \|x\|.$$

### III. LATTICE-BASED SAMPLE SETS

We derive sample sets optimizing sample complexity (Problem 1) and collision-check complexity (Problem 2) both in theory and experiments. We focus on sample sets induced by lattices.

A lattice is a point set in Euclidean space with a regular structure [3].

**Definition 3** (Lattice). A lattice  $\Lambda$  is defined as all the linear combinations (with integer coefficients) of a basis<sup>3</sup>  $E_{\Lambda} = \{e_i \in \mathbb{R}^N\}_{i=1}^m$  of rank  $1 \leq m \leq N$ , i.e.,

$$\Lambda := \left\{ \sum_{i=1}^m a_i e_i \mid a_i \in \mathbb{Z}, e_i \in E_{\Lambda} \right\}.$$

It would be convenient to view lattices through their generator matrices.

**Definition 4** (Lattice generator). The generator matrix  $G_{\Lambda}$  of a lattice  $\Lambda$  with basis  $E_{\Lambda} = \{e_i \in \mathbb{R}^N\}_{i=1}^m$  is an  $m \times N$  matrix such that for every  $1 \leq i \leq m$ , the row  $i$  is equal to  $e_i$ . Note that  $\Lambda = \{a \cdot G_{\Lambda} \mid a \in \mathbb{Z}^{1 \times m}\}$ . Additionally, define  $\det(\Lambda) := \det(G_{\Lambda} G_{\Lambda}^t)$ .

#### A. Useful lattices

We describe three lattices, visualized in Figures 1 and 2. The first lattice is a simple rectangular grid, which is provided to benchmark more complicated and efficient lattices. Below, we fix the dimension  $d \geq 2$ .

**Definition 5** ( $\mathbb{Z}^d$  lattice). The  $\mathbb{Z}^d$  lattice is defined by the identity generator matrix  $I \in \mathbb{R}^{d \times d}$ , with  $\det(\mathbb{Z}^d) = 1$  [3, p. 106].

More efficient sample sets can be generated via the  $D_d^*$  lattice [3, p120]. This lattice was also presented in [4], where it was called a “staggered grid”. In this work, we provide improved sample complexity bounds for lattice-based sample sets (including for the  $D_d^*$  lattice and the  $A_d^*$  lattice defined later on), following [3], and develop theoretical bounds for collision-check complexity.

<sup>3</sup>A basis can be of full rank ( $m = N$ ) or subdimensional ( $m < N$ ). A basis can be non-unique.

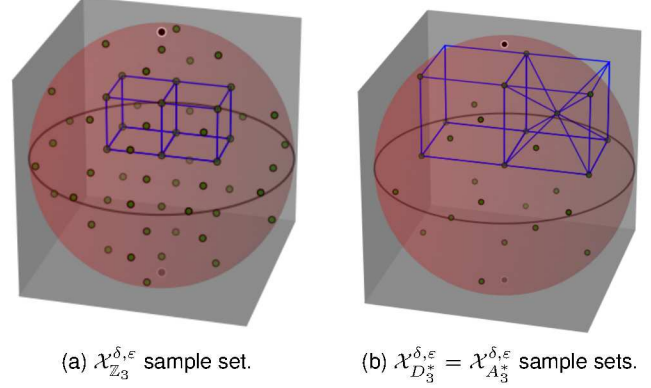


Fig. 2.  $(\delta, \varepsilon)$ -complete sample sets in  $\mathbb{R}^3$  derived from the lattices  $\mathbb{Z}^3$ ,  $D_3^*$  and  $A_3^*$ . Note the sets  $\mathcal{X}_{D_3^*}^{\delta, \varepsilon}, \mathcal{X}_{A_3^*}^{\delta, \varepsilon}$  coincide for  $d = 3$ , and diverge for  $d \geq 4$ . Note that the density of  $\mathcal{X}_{D_3^*}^{\delta, \varepsilon}$  and  $\mathcal{X}_{A_3^*}^{\delta, \varepsilon}$  (also known as the Body-Centered Cubic structure in crystallography), and is lower than the density of  $\mathcal{X}_{\mathbb{Z}^3}^{\delta, \varepsilon}$ .

**Definition 6** ( $D_d^*$  lattice). The  $D_d^*$  lattice is defined by the generator matrix

$$G_{D_d^*} = \begin{pmatrix} 1 & 0 & 0 & \dots & 0 & 0 \\ 0 & 1 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & 0 \\ 0 & 0 & 0 & \dots & 1 & 0 \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \dots & \frac{1}{2} & \frac{1}{2} \end{pmatrix} \in \mathbb{R}^{d \times d},$$

with  $\det(D_d^*) = \frac{1}{4}$  [3, p. 120].

The following  $A_d^*$  lattice [3, p115] leads to even more efficient sample sets. This lattice is also called a “hexagonal grid”, and was previously used for 2D path planning [1, 27]. This work is the first to consider its application in dimensions  $d \geq 3$ , and moreover, in the context of  $(\delta, \varepsilon)$ -completeness guarantees.

**Definition 7** ( $A_d^*$  lattice). The  $A_d^*$  lattice is defined through the generator matrix

$$G_{A_d^*} = \begin{pmatrix} 1 & -1 & 0 & 0 & \dots & 0 & 0 \\ 1 & 0 & -1 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & 0 \\ 1 & 0 & 0 & 0 & \dots & -1 & 0 \\ \frac{-d}{d+1} & \frac{1}{d+1} & \frac{1}{d+1} & \frac{1}{d+1} & \dots & \frac{1}{d+1} & \frac{1}{d+1} \end{pmatrix}_{d \times (d+1)},$$

with  $\det(A_d^*) = \frac{1}{\sqrt{d+1}}$  [3, p. 115].

Note that  $A_d^*$  is contained in  $\mathbb{R}^{d+1}$  (due to the number of rows of the generator matrix), but the lattice itself is  $d$ -dimensional as it lies in a  $d$ -dimensional hyperplane (for any lattice point  $(x_1, \dots, x_{d+1}) \in A_d^*$  it holds that  $\sum_{i=1}^{d+1} x_i = 0$ ). Our motivation for considering  $A_d^*$  is its low *density*, defined as the average number of spheres (centered on lattice points) containing a point of the space [3]. In particular,  $A_d^*$  is the best lattice covering (and best covering in general) in terms of density for dimension  $d \leq 5$  (see [37]) and overall the best *known* covering for  $d \leq 21$ .



### B. From lattices to $(\delta, \varepsilon)$ -complete sample sets

We derive sample sets from the lattices above by transforming the lattices such that the resulting point sets lie in  $\mathbb{R}^d$  and form  $\beta^*$ -covers (Lemma 1). To achieve that, we will leverage the geometry of the lattices and their covering radius, which is defined below.

**Definition 8.** (Covering radius [3]) For a point set  $\mathcal{X} \subset \mathbb{R}^d$ , a covering radius is defined to be

$$f_{\mathcal{X}} = \sup_{y \in \mathbb{R}^d} \inf_{x \in \mathcal{X}} \|x - y\|.$$

When considering the covering radius of  $A_d^*$ , which lies in  $\mathbb{R}^{d+1}$ , we will abuse the above definition to refer to its covering radius in the  $d$ -dimensional plane  $\sum_{i=1}^{d+1} x_{d+1} = 0$ . Note that in order to cover  $\mathbb{R}^d$  with balls of radius  $\rho > 0$  centered at the points of a set  $\mathcal{X}$ , it must hold that  $\rho \geq f_{\mathcal{X}}$ .

**Theorem 1.** Fix  $\delta > 0, \varepsilon > 0$ , and take  $\beta^*$  as defined in Lemma 1. Then the following sample sets are  $(\delta, \varepsilon)$ -complete:

- 1)  $\mathcal{X}_{\mathbb{Z}^d}^{\delta, \varepsilon} := \left\{ \frac{2\beta^*}{\sqrt{d}} \cdot v, v \in \mathbb{Z}^d \right\} = \frac{2\beta^*}{\sqrt{d}} \cdot \mathbb{Z}^d$ .
- 2)  $\mathcal{X}_{D_d^*}^{\delta, \varepsilon} := \begin{cases} d \text{ is odd: } \left\{ \frac{4\beta^*}{\sqrt{2d-1}} G_{D_d^*}^t \cdot v, v \in \mathbb{Z}^d \right\} = \frac{4\beta^*}{\sqrt{2d-1}} \cdot D_d^*, \\ d \text{ is even: } \left\{ \sqrt{\frac{8}{d}} \beta^* G_{D_d^*}^t \cdot v, v \in \mathbb{Z}^d \right\} = \sqrt{\frac{8}{d}} \beta^* \cdot D_d^*. \end{cases}$
- 3)  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon} := \left\{ \sqrt{\frac{12(d+1)}{d(d+2)}} \beta^* T \cdot v, v \in \mathbb{Z}^d \right\}$ , where
$$T := \begin{pmatrix} 1 & 1 & \dots & 1 & a-1 \\ -1 & 0 & \dots & 0 & a \\ 0 & -1 & \dots & 0 & a \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & -1 & a \end{pmatrix} \in \mathbb{R}^{d \times d} \text{ and } a = \frac{1}{d+1-\sqrt{d+1}}.$$

Each of the new sample sets can be viewed as a lattice in  $\mathbb{R}^d$ , according to Definition 3. For instance,  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  is a lattice with the generator matrix  $\sqrt{\frac{12(d+1)}{d(d+2)}} \beta^* T^t$ . Also, note that the result above for  $\mathcal{X}_{D_d^*}^{\delta, \varepsilon}$  is a tightening of the result in [4] for odd dimensions.

*Proof:* The lattices  $\mathbb{Z}^d, D_d^*, A_d^*$  require a transformation to achieve  $(\delta, \varepsilon)$ -completeness for a given value of  $\delta$  and  $\varepsilon$ . For a given lattice  $\Lambda$ , we compute a rescaling factor  $w_{\Lambda} > 0$  such that the covering radius of the lattice  $w_{\Lambda} \Lambda$  is not bigger than  $\beta^*$ . This would imply that  $w_{\Lambda} \Lambda$  is  $(\delta, \varepsilon)$ -complete according to Lemma 1. In particular,  $w_{\Lambda} = \beta^* f_{\Lambda}^{-1}$  where  $f_{\Lambda}$  is the covering radius of  $\Lambda$ . Next, we consider each of the three lattices individually.

*The lattice  $\mathbb{Z}^d$ .* This lattice can be viewed as a set of axis-aligned unit hypercubes whose vertices are the lattice points. The center point of a cube is located at a distance of  $\sqrt{d}/2$  from the cube's vertices, which yields the covering radius  $f_{\mathbb{Z}^d} = \sqrt{d}/2$ , and the rescaling factor  $w_{\mathbb{Z}^d} = \beta^* / f_{\mathbb{Z}^d} = 2\beta^* / \sqrt{d}$ . As a result, the sample set  $\mathcal{X}_{\mathbb{Z}^d}^{\delta, \varepsilon} := \frac{2\beta^*}{\sqrt{d}} \mathbb{Z}^d$ , is a  $\beta^*$ -cover, and  $(\delta, \varepsilon)$ -complete due to Lemma 1.




Fig. 3. Visualization of embedding the lattice  $A_d^*$  originally defined in  $\mathbb{R}^3$  onto  $\mathbb{R}^2$  via the mapping  $T$ . The blue rectangle represents the plane  $H$ , where the corresponding  $A_d^*$  lattice points are drawn in red. The points are generated by taking integer vectors in  $\mathbb{R}^d$  and applying the mapping  $G^t$ .  $H$  and  $A_d^*$  is reflected onto the plane  $H_0 = \{x_3 = 0\}$  using the mapping  $PG^t$  (denoted by the green rectangle). The third dimension is removed via the mapping  $E$  to yield the embedding of  $A_d^*$  in  $\mathbb{R}^2$ .

*The lattice  $D_d^*$ .* We use the covering radius of  $D_d^*$ , which depends on whether the dimension  $d$  is odd or even [3, page 120]. In particular,  $f_{D_{d-\text{odd}}^*} = \frac{\sqrt{2d-1}}{4}$ , and  $f_{D_{d-\text{even}}^*} = \frac{\sqrt{2d}}{4}$ . This immediately implies the definition of  $\mathcal{X}_{D_d^*}^{\delta, \varepsilon}$ .

*The lattice  $A_d^*$ .* Recall that  $A_d^*$  has the generator matrix  $G := G_{A_d^*} \in \mathbb{R}^{d \times (d+1)}$ . As this is a mapping from  $d$  to  $d+1$ , we start with the process of embedding  $A_d^*$  in  $\mathbb{R}^d$  (see Figure 3). Afterward, we show that the embedding in  $\mathbb{R}^d$  shares the same covering radius as the original set in  $\mathbb{R}^{d+1}$ .

Any row  $i$  of  $G$ , denoted by  $G_i := (g_{i1}, g_{i2}, \dots, g_{i(d+1)})$ , lies on the (hyper)plane  $H : \{\sum_{j=1}^{d+1} g_{ij} = 0\}$ . Thus,  $A_d^*$  itself is contained in that  $d$ -dimensional plane. It remains to find a transformation of  $A_d^*$  such that the dimension is reduced to  $d$  while maintaining the structure of the points in  $A_d^*$ .

In the first step, we reflect  $A_d^*$  lattice points onto the plane  $H_0 := \{x_{d+1} = 0\}$  using the Householder matrix

$$P := \left( \begin{array}{c|c} I_d - \frac{1}{D-\sqrt{D}} \mathbb{1} & \begin{matrix} \frac{1}{\sqrt{D}} \\ \vdots \\ \frac{1}{\sqrt{D}} \end{matrix} \\ \hline \frac{1}{\sqrt{D}} & \dots & \frac{1}{\sqrt{D}} & \frac{1}{\sqrt{D}} \end{array} \right), \quad (4)$$

where  $D := d+1$ ,  $I_d$  is an  $d \times d$  identity matrix, and  $\mathbb{1}$  is the  $d \times d$  matrix with 1s in all its entries. That is, we reflect a lattice point  $p \in \mathbb{R}^{d+1}$  by computing the value  $v = P \cdot p$ . (See the derivation of  $P$  in the supplementary material.)

It remains to eliminate the  $(d+1)$ th dimension of the points  $v := P \cdot p$ . This is accomplished by the mapping

$$E = \begin{pmatrix} 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & 0 \\ 0 & 0 & \dots & 1 & 0 \end{pmatrix}_{d \times (d+1)}.$$

We finish by computing the an explicit mapping that yields the embedding of  $A_d^*$  to  $\mathbb{R}^d$ . In particular, we have the embedding  $T(g) := EPG^t(g)$ , for  $g \in \mathbb{Z}^d$ , where  $T$  is as specified in the statement of this theorem. (See the derivation of  $T$  in in the supplementary material.)

For the final part, we wish to derive an  $(\delta, \varepsilon)$ -complete sample set  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$ . Here, we first recall that the covering radius  $f_{A_d^*}$  is equal to  $\sqrt{\frac{d(d+2)}{12(d+1)}}$  [3, page 115]. Considering that reflections and embeddings are isometries, i.e., they preserve distances between pairs of points, we can use the same covering radius after mapping the points of  $A_d^*$  using  $T$ . Thus, we obtain the rescaling coefficient  $w_{A_d^*} := \beta^* f_{A_d^*}^{-1}$ , which implies that  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon} := \{w_{A_d^*} T \cdot v, v \in \mathbb{Z}^d\}$  is  $(\delta, \varepsilon)$ -complete. ■

#### IV. SAMPLE COMPLEXITY

We derive the following lower and upper bounds on the sample complexity of the sets  $\mathcal{X}_{\mathbb{Z}^d}^{\delta, \varepsilon}$ ,  $\mathcal{X}_{D_d^*}^{\delta, \varepsilon}$ , and  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$ .

**Theorem 2** (Sample-complexity bounds). *Consider a lattice  $\Lambda \in \{\mathbb{Z}^d, D_d^*, A_d^*\}$  with a covering radius  $f_\Lambda$ , which yields the  $(\delta, \varepsilon)$ -complete set  $\mathcal{X}_\Lambda^{\delta, \varepsilon}$  for some  $\delta > 0, \varepsilon > 0$ . Then,*

$$|\mathcal{X}_\Lambda^{\delta, \varepsilon} \cap \mathcal{B}_{r^*}| = \frac{\partial(\mathcal{B}_1)}{\sqrt{\det(\Lambda)}} \bar{\theta}_{r^*}^d + P_d(\bar{\theta}_{r^*}), \quad (5)$$

where  $\bar{\theta}_{r^*} := 2f_\Lambda (1 + \frac{1}{\varepsilon})$ , and  $P_d(\alpha) \in \mathbb{R}$  is the discrepancy function [15].<sup>4</sup>

*Proof:* We estimate the size of the sample set  $\mathcal{X}_\Lambda^{\delta, \varepsilon}$  induced by the lattice  $\Lambda$  and the scaling factor  $f_\Lambda$  by exploiting the relation between  $\mathcal{X}_\Lambda^{\delta, \varepsilon}$  and the grid lattice  $\mathbb{Z}^d$ . In particular, a ball with respect to  $\mathcal{X}_\Lambda^{\delta, \varepsilon}$  can be viewed as a rescaled ball for the  $\mathbb{Z}^d$  lattice. This allows the use of bounds on the number of  $\mathbb{Z}^d$  points within an ellipse.

Fix a ball radius  $R > 0$ . Due to the rescaling performed in Theorem 1 we transition from the lattice  $\Lambda$ , which is a  $f_\Lambda$ -cover for  $\mathbb{R}^d$ , into the set  $\mathcal{X}_\Lambda^{\delta, \varepsilon}$ , which is a  $\beta^*$ -cover for  $\mathbb{R}^d$ . In particular, the rescaling factor is  $w_\Lambda := \beta^*/f_\Lambda$ , which is multiplied by  $\Lambda$  to obtain  $\mathcal{X}_\Lambda^{\delta, \varepsilon}$  (for  $A_d^*$  we also applied an isometric transformation, but this does not change the scale reasoning). Thus, the ball  $\mathcal{B}_R$  with respect to  $\mathcal{X}_\Lambda^{\delta, \varepsilon}$  can be viewed as the ball  $\mathcal{B}_{\bar{\theta}}$ , where  $\bar{\theta}_R = R w_\Lambda = \frac{R}{\beta^*} f_\Lambda$ . Thus,  $|\mathcal{X}_\Lambda^{\delta, \varepsilon} \cap \mathcal{B}_R| = |\Lambda \cap \mathcal{B}_{\bar{\theta}_R}|$ . For the remainder of the proof, we wish to bound the expression  $|\Lambda \cap \mathcal{B}_{\bar{\theta}_R}|$ .

Let  $v = aG_\Lambda$  be a lattice  $\Lambda$  point, where  $a \in \mathbb{Z}^d$ . By definition of the Gram matrix  $\text{Gr}_\Lambda := G_\Lambda G_\Lambda^t$ , we obtain

$$a \text{Gr}_\Lambda a^t := a G_\Lambda G_\Lambda^t a^t = \|a G_\Lambda\|^2 = \|v\|^2.$$

This leads to the relation

$$\mathcal{B}_{\bar{\theta}_R} \cap \Lambda = \{v \in \Lambda \mid \|v\| \leq \bar{\theta}_R\} = \{a \in \mathbb{Z}^d \mid a \text{Gr}_\Lambda a^t \leq \bar{\theta}_R^2\}$$

<sup>4</sup>For a value  $\alpha > 0$ ,  $|P_3(\alpha)| = \Omega_+(\sqrt{\alpha \log(\alpha)})$  and  $|P_3(\alpha)| = O(\alpha^{\frac{21}{32} + \varepsilon})$ ,  $|P_4(\alpha)| = \Omega(\alpha^2)$ ,  $|P_4(\alpha)| = O(\alpha \log^{2/3} \alpha)$ , and  $|P_d(\alpha)| = \Theta(\alpha^{d-2})$  for  $d > 4$ . Notice that  $P_d$  can be negative. For two functions  $f, g$ , the notation  $f(\alpha) = O(g(\alpha))$  (or  $f(\alpha) = \Omega(g(\alpha))$ ) means that there exists a constant  $m_u > 0$  (or  $m_l > 0$ ) such that for a large enough  $\alpha$  it holds that  $f(\alpha) \leq m_u(g(\alpha))$  (or  $f(\alpha) \geq m_l(g(\alpha))$ ).

$$= E_{\bar{\theta}_R^2}(\text{Gr}_\Lambda) \cap \mathbb{Z}^d,$$

where  $E_s(A) := \{x \in \mathbb{R}^d \mid x A x^t \leq s\}$  for a matrix  $A$  and  $s > 0$ .

Next, observe that the Gram matrix  $\text{Gr}_\Lambda$  of the lattices  $\mathbb{Z}^d, D_d^*, A_d^*$  contains only rational numbers, and hence defines a *rational quadratic form*, which allows us employ *rational ellipsoid bounds* [15] for  $|E_{\bar{\theta}_R}(\text{Gr}_\Lambda) \cap \mathbb{Z}^d|$ . Hence,

$$\begin{aligned} |\mathcal{X}_\Lambda^{\delta, \varepsilon} \cap \mathcal{B}_R| &= |\mathcal{B}_{\bar{\theta}_R} \cap \Lambda| = |E_{\bar{\theta}_R}(\text{Gr}_\Lambda) \cap \mathbb{Z}^d| \\ &= \frac{\partial(\mathcal{B}_{\bar{\theta}_R})}{\sqrt{\det(\Lambda)}} + P_d(\bar{\theta}_R) = \frac{\partial(\mathcal{B}_1)}{\sqrt{\det(\Lambda)}} \bar{\theta}_R^d + P_d(\bar{\theta}_R). \end{aligned} \quad (6)$$

The expression in Equation (5) immediately follows by plugging

$$\bar{\theta}_{r^*} = \frac{r^*}{\beta^*} f_\Lambda = \frac{2\delta(1+\varepsilon)}{\sqrt{1+\varepsilon^2}} \cdot \frac{\sqrt{1+\varepsilon^2}}{\delta\varepsilon} f_\Lambda = \frac{2(1+\varepsilon)}{\varepsilon} f_\Lambda. \quad \blacksquare$$

**Discussion.** The expression in Equation (5) is identical for our three sample sets, except for the value of the covering radius  $f_\Lambda$ . This highlights the fact that a smaller covering radii leads to a lower sample complexity. Also, notice that the value  $f_\Lambda$  is raised to the power of  $d$  in Equation (5), which emphasizes the difference between the sets in terms of sample complexity. See a plot of the sample complexity in theory and practice in Figure 5, wherein  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  has the lowest sample complexity (except in dimensions 3 where it coincides with  $\mathcal{X}_{D_d^*}^{\delta, \varepsilon}$ ). Notice that the theoretical bounds are well-aligned with the practical values.

For example, consider  $d = 12$ , where the  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  sample set is  $\approx 4$  times smaller than  $\mathcal{X}_{D_d^*}^{\delta, \varepsilon}$ . Even when this difference is not as big, e.g., in dimension 6 where the size of  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  is  $\approx 1.63$  times smaller than  $\mathcal{X}_{D_d^*}^{\delta, \varepsilon}$ , we observe tremendous impact in terms of the running of the motion-planning algorithm in experiments. This follows from the fact that the sample complexity studied here corresponds to the branching factor of the underlying search algorithm, which is known to have a significant impact on the running time. In the next session, we show that the superiority of  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  is maintained also for the collision-check complexity metric.

#### V. COLLISION-CHECK COMPLEXITY

In this section, we seek to analyze the collision-check complexity of our sample sets  $\mathcal{X}_{\mathbb{Z}^d}^{\delta, \varepsilon}, \mathcal{X}_{D_d^*}^{\delta, \varepsilon}, \mathcal{X}_{A_d^*}^{\delta, \varepsilon}$ , as the length of the graph edges is a better proxy for the computational complexity of constructing the resulting PRM than sample complexity.

Recall that the collision-check complexity of a lattice-based sample set  $\mathcal{X}_\Lambda$  is  $CC_{\mathcal{X}_\Lambda} = \sum_{x \in \mathcal{X}_\Lambda \cap \mathcal{B}_{r^*}} \|x\|$ . A naive approach to upper-bound  $CC_{\mathcal{X}_\Lambda}$  is to multiply the number of points in a  $r^*$ -ball (which corresponds to the sample complexity of  $\mathcal{X}_\Lambda$ )




Fig. 4. A sample-complexity plot for the sample sets  $\mathcal{X}_{\mathbb{Z}_d^{\delta,\varepsilon}}$ ,  $\mathcal{X}_{D_d^{\delta,\varepsilon}}$ , and  $\mathcal{X}_{A_d^{\delta,\varepsilon}}$  with  $\delta = 1$  and  $\varepsilon = 2$ . The dashed line represents the theoretical approximation (Equation (5)), where the asymptotic error term  $P_d$  is excluded. The solid line depicts the practical value, i.e., the number of lattice points within the  $r^*$ -ball in practice. Missing values are due to memory limitations.

with the radius  $r^*$ , i.e.,

$$CC\mathcal{X}_\Lambda \leq r^* \cdot |X_\Lambda \cap \mathcal{B}_{r^*}| = r^* \cdot \frac{\partial(B_1)}{\sqrt{\det(\Lambda)}} \bar{\theta}_{r^*}^d + r^* \cdot P_d(\bar{\theta}_{r^*}). \quad (7)$$

The following is a tighter bound, which reduces the coefficient of the dominant factor of the naive bound.

**Theorem 3** (CC complexity bound). *Consider a lattice  $\Lambda \in \{\mathbb{Z}_d^d, D_d^*, A_d^*\}$  with a covering radius  $f_\Lambda$ , which yields the  $(\delta, \varepsilon)$ -complete set  $\mathcal{X} := \mathcal{X}_\Lambda^{\delta,\varepsilon}$  for some  $\delta > 0, \varepsilon > 0$  and dimension  $d \geq 2$ . Then,*

$$CC\mathcal{X} \leq \zeta \cdot r^* \cdot \frac{\partial(B_1)}{\sqrt{\det(\Lambda)}} \bar{\theta}_{r^*}^d + r^* \cdot P_d(\bar{\theta}_{r^*}), \quad (8)$$

where  $\bar{\theta}_{r^*} := 2f_\Lambda \left(1 + \frac{1}{\varepsilon}\right)$ ,  $\zeta := \left(1 - \frac{\xi^{d+2} - \xi}{d\xi - (d+1)}\right) < 1$  and  $\xi := \left(\frac{d}{d+1}\right)^d$ .

*Proof:* We improve the naive bound in Equation (7) by partitioning the  $r$ -ball into annuli. Consider a sequence of  $k+1 \geq 2$  radii  $0 < r_k < \dots < r_0 = r^*$ , where  $r_i := \tilde{d}^i r$  and  $\tilde{d} := \frac{d}{d+1}$ . This leads to the bound

$$\begin{aligned} CC\mathcal{X} &\leq \sum_{i=0}^{k-1} r_i |\mathcal{X} \cap (\mathcal{B}_{r_i} \setminus \mathcal{B}_{r_{i+1}})| + r_k |\mathcal{X} \cap \mathcal{B}_{r_k}| \\ &= \sum_{i=0}^{k-1} r_i (|\mathcal{X} \cap \mathcal{B}_{r_i}| - |\mathcal{X} \cap \mathcal{B}_{r_{i+1}}|) + r_k |\mathcal{X} \cap \mathcal{B}_{r_k}| \\ &= r^* |\mathcal{X} \cap \mathcal{B}_r| + \sum_{i=1}^k (r_i - r_{i-1}) |\mathcal{X} \cap \mathcal{B}_{r_i}| \end{aligned}$$

$$\begin{aligned} &= r^* \left( \frac{\partial(B_1)}{\sqrt{\det(\Lambda)}} \bar{\theta}_r^d + P_d(\bar{\theta}_r^{d-2}) \right) \\ &+ \sum_{i=1}^k (r_i - r_{i-1}) \left( \frac{\partial(B_1)}{\sqrt{\det(\Lambda)}} \bar{\theta}_{r_i}^d + P_d(\bar{\theta}_{r_i}^{d-2}) \right) \\ &= \frac{\partial(B_1)}{\sqrt{\det(\Lambda)}} \left( \underbrace{r^* \bar{\theta}_{r^*}^d + \sum_{i=1}^k (r_i - r_{i-1}) \bar{\theta}_{r_i}^d}_{:=\gamma} \right) + r^* P_d(\bar{\theta}_{r^*}). \end{aligned} \quad (9)$$

Next, it can be shown that

$$\gamma = r^* \bar{\theta}_{r^*}^d \left( 1 - \frac{\xi^{d+2} - \xi}{d\xi - (d+1)} \right) = r^* \bar{\theta}_{r^*}^d \zeta,$$

where  $\bar{\theta}_{r_i} = r_i \frac{\bar{\theta}_{r^*}}{r^*}$ ,  $k = d$ , and  $\xi := \tilde{d}^d = \left(\frac{d}{d+1}\right)^d$ . See the supplementary material for the full details of this derivation, including the motivation behind the definitions of  $r_i$  and  $k$ .

By plugging the value of  $\gamma$  in Equation (9), Equation (8) follows. Notice that  $\tilde{d} < 1$ . Therefore,  $\tilde{d}^{d(d+1)} - 1 < 0$  and  $\tilde{d}^{d+1} - 1 < 0$ , so our expression in the parenthesis is less than 1, which implies an improvement over the trivial bound (Equation (7)). ■

**Discussion.** Theorem 3 leads to a constant-factor improvement over the naive bound concerning the coefficient of the main bound term. The value of  $\gamma$  ranges from  $\approx 0.751$  for  $d = 2$  and monotonically increases towards 1 as  $d \rightarrow \infty$  (see a plot in the supplementary material). Although our result currently leads to modest improvement, we hope it will pave the way for tighter bounds in the future. Nevertheless, our analysis emphasizes the impact of the coverage quality of lattices on the running time, which is more substantial than what the sample-complexity bound suggests.

See a plot of the theoretical bound, along with the practical values in Figure 5, where both values show similar trends. With that said, the discrepancy between the leading value in Equation (8) and the practical value is bigger than in the sample-complexity case, which suggests that the result in Theorem 3 could be tightened. Finally, note that collision-check complexity is roughly an order of magnitude larger than the sample complexity, particularly in higher dimensions.

## VI. ALGORITHMIC IMPLICATIONS

In this section, we examine the algorithmic implications of our theoretical results. We discuss the types of sampling-based planners (SBPs) that can benefit from lattice-based sampling (LBS), as defined in Theorem 1. We then discuss how to efficiently generate the points of a given LBS  $\mathcal{X}_\Lambda^{\delta,\varepsilon}$ . Finally, we provide a pseudo-code for an SBP that we use in the experimental results, while highlighting algorithmic aspects that leverage the regularity of LBS.




Fig. 5. A collision-check complexity plot for the sample sets  $\mathcal{X}_{\mathbb{Z}^d}^{\delta, \varepsilon}$ ,  $\mathcal{X}_{D^*}^{\delta, \varepsilon}$ , and  $\mathcal{X}_{A^*}^{\delta, \varepsilon}$  with  $\delta = 1$  and  $\varepsilon = 2$ . The dashed line represents the theoretical approximation (Equation (8)), where the asymptotic error term  $P_d$  is excluded.

#### A. Batch planners

LBS is well suited for batch SBPs that generate a priori a batch (or several batches) of samples, implicitly or explicitly (see below), and process them simultaneously, where the resulting graph structure is agnostic to ordering between the samples, as in Equation (1). In our context, an LBS  $\mathcal{X}_{\Lambda}^{\delta, \varepsilon}$  is regarded as a sample batch. The multi-query planners PRM [19], PRM\* [18], and SPARS [6], and the single-query planners BIT\* [11], AIT\* [45], TMIT\* [48], IRIS [9], LRA\* [32], GLS [33], FMT\* [16] and LazySP [5], and the multi-robot planners dRRT [44] and dRRT\* [42], to name just a few examples, fit this description.

LBS does not fit, in its current form, within incremental SBPs (e.g., RRT [24], RRG, or RRT\* [18]), which require the ability to add new samples one by one as time permits (as in, e.g., uniform random sampling or Halton sequences). The resulting graph structure in those algorithms depends not only on the physical location of the samples but also on the order in which they were processed. (We do believe that LBS could be adapted for *incremental* processing, as we discuss in the final section.)

#### B. Construction of lattice points

Next, we provide details for generating the lattice  $\Lambda$  points within an  $R$ -ball in Algorithm 1. Note that if  $R = r^*$ , the above point set can be viewed as the neighbor set of the origin vertex of the graph  $\mathcal{G}_{\mathcal{M}(\Lambda, r^*)}$  (assuming that  $\mathcal{B}_{r^*} \subset \mathcal{C}_f$ ). To compute  $\Lambda \cap R$ , we run a BFS search over the integer vectors  $\mathbb{Z}^d$  starting from the origin vertex, and compute for each vector  $v$  the resulting lattice point through  $x := vG_{\Lambda}$  while discarding points outside of  $R$ , where  $G_{\Lambda}$  is the generator matrix of  $\Lambda$ .

---

#### Algorithm 1: getLatticeNeighbors( $G_{\Lambda}, R$ )

---

```

1 OPEN = {0} // initialize search with center vertex
2 OPENnext = ∅ // next layer of OPEN
3 VISITED = ∅;
4 while !OPEN.empty() ∨ !OPENnext.empty() do
5   if OPEN.empty() then
6     /* switch to the next layer of neighbors */
7     OPEN = OPENnext;
8     OPENnext = ∅;
9   p = OPEN.pop();
10  if p ∈ VISITED then
11    continue
12  VISITED.insert(p);
13  /* iterate over basis integer vectors {e1, ..., ed},
14   where ei contains "1" in the i'th coordinate, and
15   "0" elsewhere, in both directions */
16  for i ∈ {1, ..., d}, sign ∈ {±1} do
17    e = sign · ei;
18    pnew = p + e · GΛ // obtain next lattice point
19    if pnew ∉ VISITED ∪ OPEN ∪ OPENnext and
20      ||p|| ≤ R then
21      OPENnext.insert(pnew);
22 return VISITED

```

---

#### C. Implicit A\*

To demonstrate the impact of LBS in practice, we use a single-query planner where an implicitly-represented PRM graph  $\mathcal{G}_{\mathcal{M}(\mathcal{X}_{\Lambda}^{\delta, \varepsilon}, r^*)}$  is explored using a search heuristic, similar to BIT\* [11], and GLS [33]. Those state-of-the-art approaches are well-suited for settings where samples are generated in large batches, as LBS facilitates. We focus on the single-query setting, as it allows us to experiment with more complex problem scenarios (e.g., in terms of dimensions and tightness) than in a multi-query setting, where the entire configuration space needs to be explored, which requires additional memory and compute time.

The planner we use, which is termed for simplicity *implicit A\** (iA\*), can be viewed as a simplified version of BIT\* with a single sample batch searched using the A\* algorithm. iA\* generates a sample set  $\mathcal{X}$  from a given sample distribution ( $\mathcal{X}_{\mathbb{Z}^d}^{\delta, \varepsilon}$ ,  $\mathcal{X}_{D^*}^{\delta, \varepsilon}$ ,  $\mathcal{X}_{A^*}^{\delta, \varepsilon}$  or uniform random sampling). Instead of constructing the entire PRM graph  $G := G_{\mathcal{M}(\mathcal{X}, r)}$  resulting from  $\mathcal{X}$  and a given radius parameter  $r$ , iA\* constructs a partial graph  $G' \subset G$  in an implicit manner, where the construction is guided by the underlying A\* search. That is, when a vertex  $v$  of  $G$  is expanded, its neighbors  $N_v$  within an  $r$ -neighborhood are retrieved from  $\mathcal{X}$ , and the edges between  $v$  and every  $u \in N_v$  are collision-checked and added to the explored portion of the graph  $G'$ . For LBS, we set  $r := r^*$ .

We consider two flavors of iA\*. In the first flavors, denoted by GLO, vertex neighbors (as  $N_v$  above) are retrieved by calling a *global* nearest-neighbor (NN) data structure. Before starting the A\* search, this data structure is initialized with



the set  $\mathcal{X}$ .

Although the benefits of LBS over randomized sampling are already apparent for the GLO flavor (especially  $\mathcal{X}_{A_d}^{\delta, \varepsilon}$ ), as will be shown in the next section, the performance of iA\* can be further improved by exploiting the *local* regular structure of lattices. In the second flavor of iA\*, denoted by LOC, which only applies to lattice-based sets, vertex neighbors are efficiently retrieved without NN data structures. Given a lattice-based sample set  $\mathcal{X}$  and a connection radius  $r > 0$ , denote by  $N_x := \mathcal{B}_r(x) \cap \mathcal{X}$  the  $r$  nearest-neighbors of a vertex  $x \in \mathcal{X}$ . Then, for another sample  $x' \in \mathcal{X}_\Lambda$  it holds that  $N_{x'} := \mathcal{B}_r(x') \cap \mathcal{X} = (x' - x) + N_x$ , i.e.,  $N_{x'}$  is a translation of  $N_x$ . Hence, the computation of the neighbor set  $N_0$ , discussed above, is only performed from scratch once at the beginning of the run of the search algorithm, where the  $N_v$  is easily obtained from vector operations.

We provide pseudo-code for iA\*-LOC in Algorithm 2. The algorithm follows the structure of the standard A\* algorithm by extracting vertices from the open list, with some small modifications. The algorithm returns a path from  $q_s$  to  $q_g$  on the graph  $G_{\mathcal{M}}(\mathcal{X}, r^*)$  (or reports that none exists otherwise), where  $\mathcal{X} := \mathcal{X}_\Lambda^{\delta, \varepsilon}$ . Without loss of generality, we assume that  $q_s \in \mathcal{X}$ , and moreover,  $q_s$  is the origin. Notice that  $\mathcal{X}$  is given to the algorithm as an implicit parameter, as is the free space  $\mathcal{C}_f$ . As mentioned earlier, the neighbor set of the origin vertex is computed once (line 5), and this information is then reused to compute the neighbor set of a general vertex  $q$  (line 10). This is where iA\*-LOC differs from iA\*-GLO, where the latter computes the neighbor set by running a nearest-neighbor search algorithm, as is common in SBPs (i.e., iA\*-GLO requires an explicit representation of the full sample set).

From an implementation standpoint, we mention that additional improvements can be obtained for iA\*-LOC by associating every sample  $q \in \mathcal{X}_\Lambda$  with the integer vector  $v \in \mathbb{Z}^d$  such that  $q := vG_\Lambda$ . This allows for efficient tracking of explored vertices. For simplicity, we omit those fine details from the pseudo-code, which can be found in our code repository.

## VII. EXPERIMENTAL RESULTS

We study practical aspects of our theoretical findings in motion planning for multi-robot systems and a manipulator arm. We report comparisons between the three lattice-based sets, as well as uniform random sampling. An additional experiment studying the effect of the parameters  $\delta$  and  $\varepsilon$  is reported in the supplementary material.

### A. Implementation details and planners

We evaluate the performance of the iA\* planner for LBS using the LOC and GLO flavors, and for uniform random sampling using the GLO variant. The experiments were performed on an ASUS Vivobook 16x laptop equipped with an Intel Core i9-13900H CPU, 32GB DDR4 RAM, and SSD storage, running Ubuntu 22.04.5 LTS OS.

The planners were implemented in C++ within OMPL [46], with FCL [36] for collision detection, and GNAT [12] for nearest-neighbor search (where applicable). The manipulator

---

### Algorithm 2: iA\*-LOC ( $\mathcal{C}_f, \mathcal{X}, r^*, q_s, q_g$ )

---

```

1  $g(q_s) = 0$  // cost-to-go of  $q_s$ 
2  $OPEN = \{q_s\}$ ;
3  $VISITED = \{q_s\}$ ;
4  $PREV[q_s] = \emptyset$  // Predecessor vertices along shortest
   path to  $q_s$ 
5  $N_0 \leftarrow \text{getLatticeNeighbors}(\mathcal{X}, r^*)$  // neighbors of  $q_s$ 
6 while  $OPEN$  not empty do
   /* extract vertex  $q$  minimizing  $g(q) + \|q_g - q\|$  */
7    $q \leftarrow OPEN.\text{pop\_min}()$ ;
8   if  $q == q_g$  then
9     | return path reconstructed using PREV
   /* compute neighbors of  $q$  */
10   $N_q = N_0 + q$ ;
11  if  $\|q - q_g\| < r^*$  then
12    |  $N_q \leftarrow N_q \cup \{q_g\}$  // add  $q_g$  to neighbor set
13  for  $n$  in  $N_q$  do
    /* collision-check edge from  $q$  to  $n$  */
14    if  $\overline{qn} \not\subset \mathcal{C}_f$  then
15      | continue
16     $g_{\text{new}}(n) = g(q) + \|q - n\|$  // potential
      cost-to-come of  $n$ 
17     $g_{\text{cur}}(n) = \infty$  // current cost-to-come of  $n$ 
18    if  $n \in VISITED$  then
19      |  $g_{\text{cur}}(n) = g(n)$ ;
20    if  $g_{\text{new}}(n) < g_{\text{cur}}(n)$  then
21      | if  $n \notin VISITED$  then
22        |  $VISITED.\text{insert}(n)$ 
23        |  $g(n) = g_{\text{new}}(n)$ ;  $PREV[n] = q$ ;
24      | if  $n \notin OPEN$  then
25        |  $OPEN.\text{insert}(n)$ ;
26 return  $\emptyset$ 

```

---

simulations also rely on the VAMP library [47], which significantly speeds up collision checking via fine-grain parallelization of edge collision checking.

### B. Scenarios

We test the planners on various motion-planning problems in two settings.

**Multi-robot system.** The first setting considers multi-robot systems with a total number of degrees of freedom  $d \in \{4, 6, 8, 10, 12\}$ , where an  $m$  disc-robot system is viewed as a single-robot system whose configuration space is  $\mathbb{R}^{2m}$ . As such, we construct sample sets in  $\mathbb{R}^{2m}$  by relying on  $2m$ -dimensional lattices (or uniform random sampling).

Each scenario consists of a multi-robot system of  $m$  labeled planar disc robots that need to (simultaneously) exchange positions (i.e., robot  $i \in \{0, \dots, m-1\}$  moves to the start position of robot  $i+1 \bmod m$ ) while avoiding collisions with each other and static obstacles.

This setting is considered for two reasons. First, such multi-robot systems can be viewed as Euclidean systems, which allows applying our theoretical results (see discussion

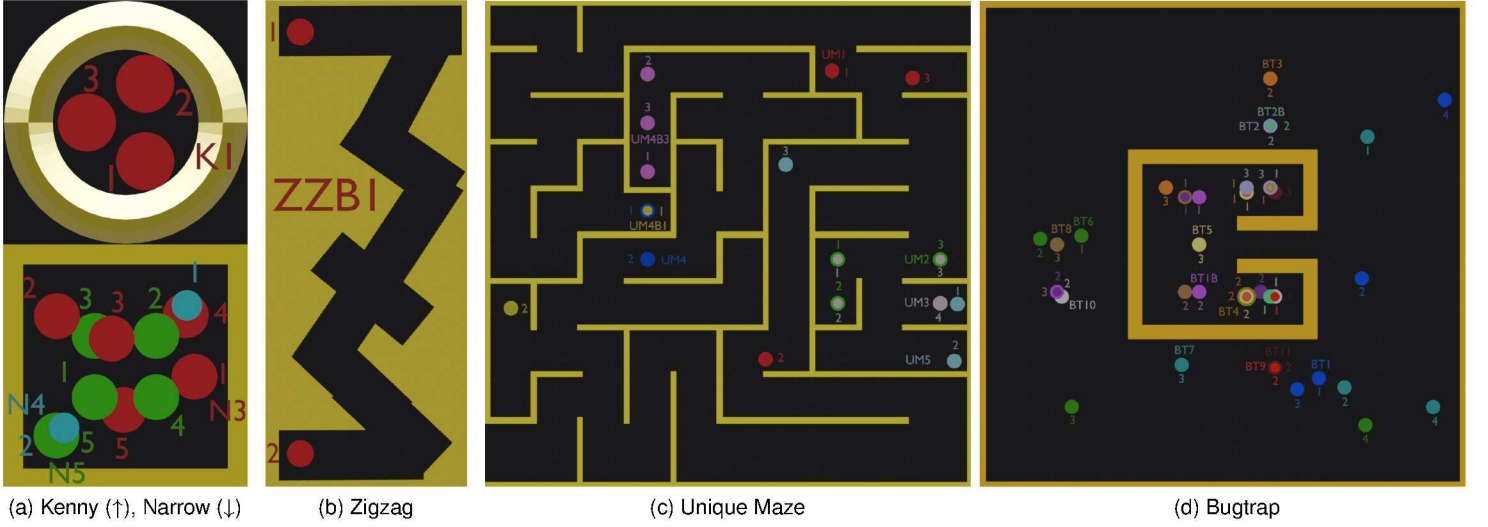


Fig. 6. A subset of the scenarios used in the experimental results for the multi-robot setting. Some of the figures depict several scenarios, where each scenario consists of a workspace environment, along with an initial configuration specifying the number of robots, their initial positions, and a permutation of their target positions. Each configuration is drawn in a different color, where the scenario name is indicated by the first letters of the workspace name along with a number. (a) (Top) For the "Kenny" workspace, there is a single scenario K1, where robot 1 starts in the bottom position, robot 2 in the top position, and robot 3 in the remaining position, where the disc size corresponds to the robot geometry. In this map, the robots are forced to perform a simultaneous movement. (a) (Bottom) For the "Narrow Room" workspace, we illustrate the scenarios N3, N4, and N5. Due to overlap, the discs in N4 are slightly shrunk for better visualizations. (b) The Zigzag scenario contains narrow, winding passages in which one of the robots can hide in a pocket to let the other robot pass. Swaps can also occur in the top or bottom corners, albeit with greater care. (c), (d) Those workspaces are taken from OMPL [46].

in Section VIII of extensions to more general systems). (Specifically, the configuration space of an  $m$ -robots system is  $\mathbb{R}^{2m}$ .) Second, this setting allows us to test systems of various dimensions while still being able to visualize the problem setting (which is in 2D). Third, it provides a simple approach for determining the value  $\delta$  (see below).

A subset of the multi-robot test scenarios is found in Figure 6 (additional scenarios are found in the supplementary material along with a detailed description). The scenarios present various difficulty levels for the planners, where the most difficult scenarios consist of narrow passages for the individual robots and a significant amount of coordination between the robots, giving rise to narrow passages in the full configuration space.

Unless stated otherwise, the parameters  $\delta$  and  $\varepsilon$  were specified in the following manner. We assigned  $\varepsilon := 10$  to focus on running time scalability rather than solution quality. The parameter  $\delta$  was initially set to be the clearance of the start configuration (capturing both distances from obstacles and between robots), which was decreased until a solution using  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  was found. We discuss the automatic tuning of  $\delta$  and  $\varepsilon$  in Section VIII.

**Manipulator arm** The second setting considers a single Panda manipulator arm with  $d = 7$  (see scenarios in Figure 7). Here, we demonstrate the integration of lattice-based sampling within the novel VAMP library, which significantly speeds up collision checking [47].

Below, we evaluate several aspects of the lattice-based sampling, which we first consider in the multi-robot setting. The final part of this section considers the manipulator case.

### C. Comparison between lattice-based sample sets

In the first set of experiments, we study the running time and solution quality (in terms of path length) of the three lattice-based sample sets  $\mathcal{X}_{\mathbb{Z}^d}^{\delta, \varepsilon}$ ,  $\mathcal{X}_{D_d^*}^{\delta, \varepsilon}$ , and  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  within the LOC flavor of iA\* for a selected set of scenarios. Although in some scenarios, the performance between the sample sets can be comparable, in terms of running time, especially for  $\mathcal{X}_{D_d^*}^{\delta, \varepsilon}$  and  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$ , here we highlight situations where large gaps occur.

The results are reported in Table I. (Results for additional scenarios are provided in the supplementary material.) In terms of running time,  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  outperforms the other sample sets, as predicted by our theoretical results with respect to sample complexity and collision-check complexity.  $\mathcal{X}_{\mathbb{Z}^d}^{\delta, \varepsilon}$  results in at least one order of magnitude (up to 2 orders) slower running times than the other two sample sets.  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  outperforms  $\mathcal{X}_{D_d^*}^{\delta, \varepsilon}$ , being at least  $3\times$ , and sometimes as much as  $10\times$ , faster. The notations "dnf" and "-" indicate a running time threshold of 1000 seconds was exceeded and failure to find a solution, respectively.

Regarding solution quality,  $\mathcal{X}_{\mathbb{Z}^d}^{\delta, \varepsilon}$  outperforms the other sample sets, except for the last scenario (unless it does not manage to find a solution) due to a denser graph. The difference in the solution length suggests that the completeness-cover relation Lemma 1 can be tightened by, e.g., reducing the value of  $\beta^*$ , albeit we emphasize this is a worst-case bound. From a practical perspective, the difference in the path length can be reduced via post-processing techniques with negligible computational cost. To summarize,  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$ -sampling can drastically reduce computational cost while achieving comparable solution quality to the other sample distribution.

TABLE I  
COMPARISON OF RUNNING TIME AND SOLUTION LENGTH USING LATTICES-BASED SAMPLE SETS (WHERE THE UNDERLYING LATTICE IS DENOTED IN THE TABLE) IN THE iA\*-LOC ALGORITHM. SOLUTION LENGTH IS NORMALIZED WITH RESPECT TO THE LENGTH OBTAINED USING  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$ .

Scenario (robot #)	Time (s)			Length (r)	
	$\mathbb{Z}^d$	$D_d^*$	$A_d^*$	$\mathbb{Z}^d$	$D_d^*$
N1(5)	165.35	4.59	0.36	0.65	0.79
N1B(6)	dnf	328.30	15.08	dnf	0.89
BT10(2)	-	1.20	0.30	-	0.92
BT5(3)	0.54	0.14	0.06	0.38	0.51
BT1(4)	146.69	50.81	3.51	0.95	1.03
K1(3)	32.31	4.97	1.37	0.82	0.89
UM4(2)	-	8.47	2.43	-	0.90
UM2(3)	13.35	1.22	0.04	1.04	1.52

TABLE II  
COMPARISON OF RUNNING TIME AND SOLUTION LENGTH BETWEEN  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  (USING LOC AND GLO) AND UNIFORM RANDOM SAMPLING. FOR RND WE REPORT THE AVERAGE VALUES IN TERMS OF RUNNING AND SOLUTION LENGTH (THE LATTER IS GIVEN AS NORMALIZED VALUE WITH RESPECT TO THE SOLUTION LENGTH WITH  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$ ).

Scenario (robot #)	Time (s)			Length (r)		Success (%)
	$A_d^*$ LOC	$A_d^*$ GLO	RND GLO	RND GLO	RND GLO	
N1(5)	0.36	3.05	4.16	1.48		80
N1(5)	0.36	3.05	4.16	1.48		80
N2(5)	0.41	2.67	2.74	2.43		65
N3(5)	0.59	3.83	5.44	2.02		85
BT3(3)	5.38	14.15	62.22	1.05		100
BT8(3)	12.17	19.31	169.32	1.00		100
BT8B(3)	3.17	3.60	41.63	1.16		100
UM4(2)	2.43	2.93	12.71	0.96		70
UM1(3)	6.68	58.62	49.35	0.98		100
UM2(3)	0.04	2.94	4.49	1.95		75
ZZB1(2)	0.44	0.49	10.44	0.89		100
ZZB2(2)	0.71	2.51	272.70	0.89		100
ZZB3(2)	0.47	7.69	341.84	0.88		100

For this reason, we omit comparisons with  $\mathcal{X}_{\mathbb{Z}^d}^{\delta, \varepsilon}$  and  $\mathcal{X}_{D_d^*}^{\delta, \varepsilon}$  in the remainder of this section.

#### D. Comparison with randomized sampling

We compare the performance of iA\* when using  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$ -sampling and uniform random sampling (denoted by RND). We show that the advantage of  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  stems not only from algorithmic speedups due to the regular lattice structure (as in LOC) but also from structural properties leading to a more efficient coverage of space. Hence, we run both GLO and LOC using  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$ , where RND is tested only with GLO.

For a given scenario, we fix the parameters  $\delta$  and  $\varepsilon$  used to derive  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$ . To derive a set of random samples  $\mathcal{X}_{\text{rnd}}^{\delta, \varepsilon}$ , we compute the number of  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  samples within the scenario (while ignoring collision with obstacles and between robots) and produce the same number of points via uniform random sampling. Importantly, when running GLO with  $\mathcal{X}_{\text{rnd}}^{\delta, \varepsilon}$  we specify the standard connection radius  $r_{\text{rnd}}(n) = \psi \left( \frac{\log n}{n} \right)^{1/d}$ , where  $n := |\mathcal{X}_{\text{rnd}}^{\delta, \varepsilon}|$  and  $\psi$  is a constant, such that asymptotic optimality is guaranteed [18]. The results are reported in

Table II, where RND is averaged over 20 repetitions. (Results for additional scenarios are provided in the supplementary material.)

Note that RND achieves a perfect success rate only in some of the scenarios. Furthermore, its running time is typically at least  $5\times$  slower than LOC, and in some scenarios, up to three orders of magnitude slower. This time gap can be partially attributed to RND constructing and maintaining a nearest-neighbor data structure. However, notice that in some scenarios,  $A_d^*$ -GLO significantly outperforms RND. One explanation of that is  $A_d^*$  requires a smaller connection radius than RND, which leads to a lower sample complexity and collision check complexity. Another reason is that RND, due to poor coverage of space, is forced to make detours and so considers additional vertices and collision checks. We provide further evidence for those points in the supplementary material. We also report results for RND where the connection radius  $r_{\text{rnd}}(n)$  is substituted with  $r^*$ , which further reduces success rates and emphasizes the efficiency of space coverage using  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$ .

Overall, we conclude that locality can substantially speed up performance from an algorithmic perspective, while the structural properties of  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$  also improve performance in terms of success rate, running time, and solution quality.

#### E. Tests on a Panda arm

Here, we evaluate the performance of iA\*, which is implemented in the VAMP environment. Specifically, we compare iA\*-LOC, using the sample set  $\mathcal{X}_{A_d^*}^{\delta, \varepsilon}$ ,

with iA\*-RND, where the former again proves to be significantly better (see results in Table III). We set  $\varepsilon = 10$  and lowered the parameter  $\delta$  until iA\*-LOC managed to find a solution.

In the "1x" column, we report the results for RND, where the number of samples equals  $|\mathcal{X}_{A_d^*}^{\delta, \varepsilon}|$ , as we have previously done. Here, other than the limbo scenario, RND achieves a similar running time to LOC. However, other than in two scenarios, RND achieves poor success rates.

We then increase the number of samples used by RND by 10 and 100 (reported in the columns "10x" and "100x") to find how many samples are necessary so that RND achieves near-perfect success rates. Indeed, in the window and sleeve scenarios, the success rate improves with additional samples, but with a significant running-time penalty. We also attempted this modification in the room scenario, but it was too computationally heavy even for 10x.

### VIII. LIMITATIONS AND FUTURE WORK

We leveraged foundational results in lattice theory to develop a theoretical framework for generating efficient sample sets for motion planning that endow their planners with finite-time guarantees, which are vastly superior to previous asymptotic properties. We demonstrated the practical potential of the framework in challenging motion-planning scenarios wherein our  $A_d^*$ -based sampling procedures lead to substantial speeds over previous methods (deterministic and random). Below, we






Fig. 7. Manipulator tests in the VAMP environment, depicting the start and target configurations. Window: The arm moves through a window frame. Sleeve: The arm moves through an elongated window frame. Limbo: The arm ducks under a pole. Sandwich: The arm is being constrained to move in a low-ceiling room. Room: The arm moves through a low window frame, while dodging several obstacles on the floor.

TABLE III  
MANIPULATOR TEST RESULTS, WHERE RND RESULTS ARE AVERAGED OVER 100 RUNS. RUNNING TIMES ARE REPORTED IN MILLISECONDS.

Scenario	Total (ms)	1x samples			10x samples			100x samples		
		Length (r)		Success (%)	Length (r)		Success (%)	Length (r)		Success (%)
		RND	RND		RND	RND		RND	RND	
	$A_d^*$	GLO	GLO	GLO	GLO	GLO	GLO	GLO	GLO	GLO
window	3.6	2.9	1.3	13	43.31	1.13	34	526.2	1.06	94
sleeve	160.2	108.9	0.77	26	1510.8	0.69	51	18800.2	0.66	100
limbo	82.2	675.7	0.7	96	-	-	-	-	-	-
sandwich	9.7	5.1	1.6	100	-	-	-	-	-	-
room	9718	14823	1.0	9	-	-	-	-	-	-

discuss several limitations of our work, which motivate further research directions.

**Multi-resolution search.** The choice of the sampling parameters  $\delta$  and  $\varepsilon$  can significantly impact the sample and collision-check complexity of the sampling distribution, and hence the planner’s performance. In this work, we selected those values according to a rule of thumb, which might not be generally applicable to broader problem settings, or lead to feasible solutions. We plan on exploring algorithmic extension, which would allow the automatic selection of those parameters during runtime. One promising direction is exploiting multi-resolution search [39, 10], which allows for adjusting sampling densities according to the space’s structure. This has the potential of not only simplifying the usage of our sample sets but also significantly improving the planner’s performance.

**Non-Euclidean systems.** Our current work focuses on geometric systems (whose state space is Euclidean) and path-length cost functions, which limits its applicability. In the future, we plan to extend our work to more general settings, which entails at least two challenges. First, the adaptation of our sampling distributions to non-Euclidean spaces is non-trivial due to the heterogeneity of the space coordinates, e.g., position and rotation components in  $SE(3)$  [50], or state derivatives in dynamical systems [17]. Second, those settings call for more general cost functions, which requires revising the concept of  $(\delta, \varepsilon)$ -completeness and its relation to geometric coverage (particularly, Lemma 1). We plan to address those challenges incrementally. As a first step, we plan on generalizing our theory to linear systems with quadratic costs, where we believe that the structure of LQR controllers [29] can be useful. More general approaches could be obtained from



exploiting the locally-linear structure of nonlinear systems using differential geometry [2].

**Incremental sampling.** Uniform random sampling is naturally amenable for densification, i.e., the incremental addition of new sampling points, making them applicable for anytime planner like RRT or RRT\* [25, 18]. In contrast, lattice-based sampling requires the introduction of large batches, which currently limits their applicability. One way to bridge this gap could be generalizing Halton sequences [22], a popular method for deterministic incremental low-discrepancy sampling, to the  $A_d^*$  lattice. Halton sequences ensue from transforming points from the standard grid lattice  $\mathbb{Z}^d$ , and it would be interesting to understand the structure emerging from feeding those points through the generator matrix of the lattice  $A_d^*$ . Another interesting direction is incrementally introducing the elements of lattice-based samples according to a random permutation.

#### ACKNOWLEDGMENTS

This research was partly supported by Israel’s Ministry of Innovation, Science and Technology, Ministry of Transport, National Road Safety Authority, and Netivei Israel, Grant No. 0006730.

We thank Rom Pinchasi, Roy Meshulam, Dan Halperin, and Oren Salzman, and Steve LaValle for fruitful discussions. We also thank Ido Jacobi, Roy Steinberg and Yaniv Hasidoff for comments on the manuscript.

#### REFERENCES

- [1] James P. Bailey, Alex Nash, Craig A. Tovey, and Sven Koenig. Path-length analysis for grid-based path planning. *Artificial Intelligence*, 301:103560, 2021.
- [2] Francesco Bullo and Andrew D. Lewis. *Geometric Control of Mechanical Systems*, volume 49 of *Texts in Applied Mathematics*. Springer Verlag, 2004.
- [3] John Horton Conway and Neil James Alexander Sloane. *Sphere packings, lattices and groups*, volume 290. Springer Science & Business Media, 2013.
- [4] Dror Dayan, Kiril Solovey, Marco Pavone, and Dan Halperin. Near-optimal multi-robot motion planning with finite sampling. *IEEE Transactions on Robotics*, 2023.
- [5] Christopher M. Dellin and Siddhartha S. Srinivasa. A unifying formalism for shortest path problems with expensive edge evaluations via lazy best-first search over paths with edge selectors. In *International Conference on Automated Planning and Scheduling*, pages 459–467. AAAI Press, 2016.
- [6] Andrew Dobson and Kostas E. Bekris. Sparse roadmap spanners for asymptotically near-optimal motion planning. *Int. J. Robotics Res.*, 33(1):18–47, 2014.
- [7] Andrew Dobson, George V. Moustakides, and Kostas E. Bekris. Geometric probability results for bounding path quality in sampling-based roadmaps after finite computation. In *International Conference on Robotics and Automation*, pages 4180–4186. IEEE, 2015.
- [8] Wolfgang Ebeling. *Lattices and Codes: A Course Partially Based on Lectures by Friedrich Hirzebruch*. Advanced Lectures in Mathematics. Springer, 2013.
- [9] Mengyu Fu, Alan Kuntz, Oren Salzman, and Ron Alterovitz. Asymptotically optimal inspection planning via efficient near-optimal search on sampled roadmaps. *Int. J. Robotics Res.*, 42(4-5):150–175, 2023.
- [10] Mengyu Fu, Kiril Solovey, Oren Salzman, and Ron Alterovitz. Toward certifiable optimal motion planning for medical steerable needles. *Int. J. Robotics Res.*, 42(10):798–826, 2023.
- [11] Jonathan D. Gammell, Timothy D. Barfoot, and Siddhartha S. Srinivasa. Batch informed trees (BIT\*): Informed asymptotically optimal anytime search. *Int. J. Robotics Res.*, 39(5), 2020.
- [12] Bryant Gipson, Mark Moll, and Lydia E Kavraki. Resolution independent density estimation for motion planning in high-dimensional spaces. In *2013 IEEE international conference on robotics and automation*, pages 2437–2443. IEEE, 2013.
- [13] Dan Halperin, Oren Salzman, and Micha Sharir. Algorithmic motion planning. In Csaba D. Tóth, Joseph O’Rourke, and Jacob E. Goodman, editors, *Handbook of Discrete and Computational Geometry*, pages 1311–1342. Chapman and Hall/CRC, 3rd edition, 2017.
- [14] Kris Hauser and Yilun Zhou. Asymptotically optimal planning by feasible kinodynamic planning in a state-cost space. *IEEE Trans. Robotics*, 32(6):1431–1443, 2016.
- [15] A Ivic, E Krätzel, M Kühleitner, and WG Nowak. Lattice points in large regions and related arithmetic functions: recent developments in a very classic topic. *arXiv preprint math/0410522*, 2004.
- [16] 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. *I. J. Robotics Res.*, 34(7):883–921, 2015.
- [17] Lucas Janson, Brian Ichter, and Marco Pavone. Deterministic sampling-based motion planning: Optimality, complexity, and performance. *The International Journal of Robotics Research*, 37(1):46–61, 2018.
- [18] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7):846–894, 2011.
- [19] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE transactions on Robotics and Automation*, 12(4):566–580, 1996.
- [20] Michal Kleinbort, Oren Salzman, and Dan Halperin. Collision detection or nearest-neighbor search? On the computational bottleneck in sampling-based motion planning. In *Workshop on the Algorithmic Foundations of Robotics*, pages 624–639. Springer, 2016.
- [21] Michal Kleinbort, Kiril Solovey, Zakary Littlefield, Kostas E. Bekris, and Dan Halperin. Probabilistic completeness of RRT for geometric and kinodynamic plan-

- ning with forward propagation. *Robotics and Automation Letters*, 4(2):277–283, 2019.
- [22] Lauwerens Kuipers and Harald Niederreiter. *Uniform distribution of sequences*. Courier Corporation, 2012.
  - [23] Tobias Kunz and Mike Stilman. Kinodynamic RRTs with fixed time step and best-input extension are not probabilistically complete. In *Workshop on the Algorithmic Foundations of Robotics*, pages 233–244, 2014.
  - [24] Steven M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006.
  - [25] Steven M. LaValle and James J. Kuffner. Randomized kinodynamic planning. *I. J. Robotics Res.*, 20(5), 2001.
  - [26] Sihui Li and Neil T Dantam. A sampling and learning framework to prove motion planning infeasibility. *The International Journal of Robotics Research*, 42(10):938–956, 2023.
  - [27] Teng Li, Min Xia, Jiahong Chen, Shujun Gao, and Clarence de Silva. A hexagonal grid-based sampling planner for aquatic environmental monitoring using unmanned surface vehicles. In *IEEE International Conference on Systems, Man, and Cybernetics*, pages 3683–3688, 2017.
  - [28] Yanbo Li, Zakary Littlefield, and Kostas E. Bekris. Asymptotically optimal sampling-based kinodynamic planning. *I. J. Robotics Res.*, 35(5):528–564, 2016.
  - [29] Daniel Liberzon. *Calculus of variations and optimal control theory: A concise introduction*. Princeton University Press, 2011.
  - [30] Yunhao Luo, Chen Sun, Joshua B. Tenenbaum, and Yilun Du. Potential based diffusion motion planning. In *Forty-first International Conference on Machine Learning*, 2024.
  - [31] D. Malyuta, T. P. Reynolds, M. Szmuk, T. Lew, R. Bonalli, M. Pavone, and B. Acikmese. Convex optimization for trajectory generation. *IEEE Control Systems Magazine*, 42(5):40–113, 2022.
  - [32] Aditya Mandalika, Oren Salzman, and Siddhartha S. Srinivasa. Lazy receding horizon A\* for efficient path planning in graphs with expensive-to-evaluate edges. In *International Conference on Automated Planning and Scheduling*, pages 476–484. AAAI Press, 2018.
  - [33] Aditya Mandalika, Sanjiban Choudhury, Oren Salzman, and Siddhartha S. Srinivasa. Generalized lazy search for robot motion planning: Interleaving search and edge evaluation via event-based toggles. In *International Conference on Automated Planning and Scheduling*, pages 745–753. AAAI, 2019.
  - [34] Tobia Marcucci, Mark Petersen, David von Wrangel, and Russ Tedrake. Motion planning around obstacles with convex optimization. *Science Robotics*, 8(84), 2023.
  - [35] Andreas Orthey, Constantinos Chamzas, and Lydia E. Kavraki. Sampling-based motion planning: A comparative review. *Annual Review of Control, Robotics, and Autonomous Systems*, 7(1), 2024.
  - [36] Jia Pan, Sachin Chitta, and Dinesh Manocha. FCL: A general purpose library for collision and proximity queries. In *2012 IEEE International Conference on Robotics and Automation*, pages 3859–3866, 2012.
  - [37] Sergei Sergeevich Ryshkov and Evgenii Petrovich Baranovskii. Solution of the problem of least dense lattice covering of five-dimensional space by equal spheres. In *Doklady Akademii Nauk*, volume 222, pages 39–42. Russian Academy of Sciences, 1975.
  - [38] Donald E. Sands. *Introduction to Crystallography*. Dover Publications, New York, 1994.
  - [39] Dhruv Mauria Saxena, Tushar Kusnur, and Maxim Likhachev. AMRA\*: Anytime multi-resolution multi-heuristic A\*. In *IEEE International Conference on Robotics and Automation*, pages 3371–3377, 2022.
  - [40] John Schulman, Yan Duan, Jonathan Ho, Alex X. Lee, Ibrahim Awwal, Henry Bradlow, Jia Pan, Sachin Patil, Ken Goldberg, and Pieter Abbeel. Motion planning with sequential convex optimization and convex collision checking. *Int. J. Robotics Res.*, 33(9):1251–1270, 2014.
  - [41] Seiji A Shaw, Aidan Curtis, Leslie Pack Kaelbling, Tomás Lozano-Pérez, and Nicholas Roy. Towards practical finite sample bounds for motion planning in TAMP. In *Workshop on the Algorithmic Foundations of Robotics*, 2024.
  - [42] Rahul Shome, Kiril Solovey, Andrew Dobson, Dan Halperin, and Kostas E. Bekris. dRRT\*: Scalable and informed asymptotically-optimal multi-robot motion planning. *Auton. Robots*, 44(3-4):443–467, 2020.
  - [43] Carl Ludwig Siegel. *Lectures on the Geometry of Numbers*, volume 201 of *Springer-Verlag Grundlehren der mathematischen Wissenschaften*. Springer, 1989.
  - [44] Kiril Solovey, Oren Salzman, and Dan Halperin. Finding a needle in an exponential haystack: Discrete RRT for exploration of implicit roadmaps in multi-robot motion planning. *Int. J. Robotics Res.*, 35(5):501–513, 2016.
  - [45] Marlin P. Strub and Jonathan D. Gammell. Adaptively informed trees (AIT\*) and effort informed trees (EIT\*): Asymmetric bidirectional sampling-based path planning. *Int. J. Robotics Res.*, 41(4):390–417, 2022.
  - [46] Ioan A. Şucan, Mark Moll, and Lydia E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 19(4):72–82, December 2012. <https://ompl.kavrakilab.org>.
  - [47] Wil Thomason, Zachary Kingston, and Lydia E. Kavraki. Motions in microseconds via vectorized sampling-based planning. In *IEEE International Conference on Robotics and Automation*, pages 8749–8756.
  - [48] Wil Thomason, Marlin P. Strub, and Jonathan D. Gammell. Task and motion informed trees (TMIT\*): Almost-surely asymptotically optimal integrated task and motion planning. *IEEE Robotics Autom. Lett.*, 7(4):11370–11377, 2022.
  - [49] Matthew Tsao, Kiril Solovey, and Marco Pavone. Sample complexity of probabilistic roadmaps via  $\epsilon$ -nets. In *International Conference on Robotics and Automation*, pages 2196–2202. IEEE, 2020.

- [50] Anna Yershova and Steven M. LaValle. Deterministic sampling methods for spheres and  $SO(3)$ . In *IEEE International Conference on Robotics and Automation*, pages 3974–3980, 2004.