Rods and Rings: Soft Subdivision Planner for $\mathbb{R}^3 \times S^2$

Ching-Hsiang Hsu
Department of Computer Science, Courant Institute, New York University, New York, NY, USA
chhsu@nyu.edu

Yi-Jen Chiang
Department of Computer Science and Engineering, Tandon School of Engineering, New York University, Brooklyn, NY, USA
chiang@nyu.edu

Chee Yap
Department of Computer Science, Courant Institute, New York University, New York, NY, USA
yap@cs.nyu.edu

Abstract

We consider path planning for a rigid spatial robot moving amidst polyhedral obstacles. Our robot is either a rod or a ring. Being axially-symmetric, their configuration space is $\mathbb{R}^3 \times S^2$ with 5 degrees of freedom (DOF). Correct, complete and practical path planning for such robots is a long standing challenge in robotics. While the rod is one of the most widely studied spatial robots in path planning, the ring seems to be new, and a rare example of a non-simply-connected robot. This work provides rigorous and complete algorithms for these robots with theoretical guarantees. We implemented the algorithms in our open-source Core Library. Experiments show that they are practical, achieving near real-time performance. We compared our planner to state-of-the-art sampling planners in OMPL [31].

Our subdivision path planner is based on the twin foundations of ε-exactness and soft predicates. Correct implementation is relatively easy. The technical innovations include subdivision atlases for S^2, introduction of Σ_2 representations for footprints, and extensions of our feature-based technique for "opening up the blackbox of collision detection".

2012 ACM Subject Classification Theory of computation → Computational geometry; Computing methodologies → Robotic planning

Keywords and phrases Algorithmic Motion Planning, Subdivision Methods, Resolution-Exact Algorithms, Soft Predicates, Spatial Rod Robots, Spatial Ring Robots

Digital Object Identifier 10.4230/LIPIcs.SoCG.2019.43

Related Version Full version [16] hosted on arXiv as `arXiv:1903.09416 [cs.CG]`, available at `https://arxiv.org/abs/1903.09416`. Also available at `http://cse.poly.edu/chiang/rod-ring18.pdf`.

Funding Supported in part by NSF Grants #CCF-1423228 and #CCF-1563942.
Ching-Hsiang Hsu: Supported by NSF Grant #CCF-1423228.
Chee Yap: Supported in part by NSF Grants #CCF-1423228 and #CCF-1563942.

1 Introduction

Motion planning [18, 5] is a fundamental topic in robotics because the typical robot is capable of movement. Such algorithms are increasingly relevant with the current surge of interest in inexpensive commercial mobile robots, from domestic robots that vacuum the floor to drones that deliver packages. We focus on what is called **path planning** which, in its elemental form, asks for a collision-free path from a start to a goal position, assuming a known environment. Path planning is based on robot kinematics and collision-detection only, and the variety of such problems are surveyed in [14]. The output of a "path planner" is either a path or a

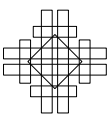


NO-PATH, signifying that no path exists. Remarkably, the single bit of information encoded by NO-PATH is often missing in discussions. The standard definitions of correctness for path planners (**resolution completeness** and **probabilistic completeness**) omit this bit [32]. The last 30 years have seen a flowering of practical path planning algorithms. The dominant algorithmic paradigm of these planners has been variants of the **Sampling Approach** such as PRM, EST, RRT, SRT, etc (see [5, p. 201]). Because this bit of information is not built into the specification of such algorithms, it has led to non-termination issues and a large literature addressing the "narrow passage problem" (e.g., [22, 8]). Our present paper is based on the **Subdivision Approach**. This approach has a venerable history in robotics – see [3, 40] for early planners based on subdivision.

Exact path planning has many issues including a serious gap between theory and implementability. In [32, 36], we introduced a theoretical framework based on subdivision to close this gap. This paper demonstrates for the first time that our framework is able to achieve rigorous state-of-the-art planners in 3D. Figure 1 shows our rod robot in an environment with 100 random tetrahedra. Figure 6 shows our ring robot in an environment with pillars and L-shaped posts. See a video demo from `http://cs.nyu.edu/exact/gallery/rod-ring/rod_ring.html`.

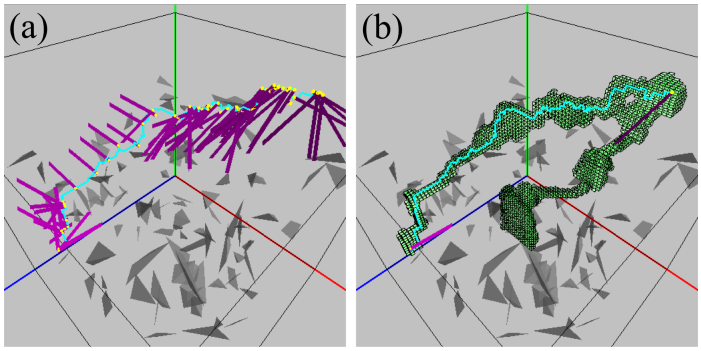

Figure 1 Rod robot amidst 100 random tetrahedra: (a) trace of a found path; (b) subdivision of translational boxes on the path.

In this paper, we consider a rigid spatial robot R_0 that has an axis of symmetry. See Figure 2(a) for several possibilities for R_0: rod ("ladder"), cone ("space shuttle"), disc ("frisbee") and ring ("space station"). Our techniques easily allow these robots to be "thickened" by Minkowski sum with balls (see [33]). The **configuration space** may be taken to be $C_{space} = \mathbb{R}^3 \times S^2$ where S^2 is the unit 2-sphere. We identify R_0 with a closed subset of $\mathbb{R}^3$, called its "canonical footprint". E.g., if R_0 is a rod (resp., ring), then the canonical footprint is a line segment (resp., circle) in $\mathbb{R}^3$. Each configuration $\gamma \in C_{space}$ corresponds to a rotated translated copy of the canonical footprint, which we denote by $Fp(\gamma)$. Path planning involves another input, the **obstacle set** $\Omega \subseteq \mathbb{R}^3$ that the robot must avoid. We assume that Ω is a closed polyhedral set. Say γ is **free** if $Fp(\gamma) \cap \Omega$ is empty. The **free space** comprising all the free configurations is an open set by our assumptions, and is denoted $C_{free} = C_{free}(\Omega)$. A parametrized continuous curve $\mu : [0, 1] \to C_{space}$ is called a **path** if the range of μ is in C_{free}. Path planning amounts to finding such paths. Following [40], we need to classify boxes $B \subseteq C_{space}$ into one of three types: FREE, STUCK or MIXED. Let $C(B)$ denote the classification of B: $C(B) =$ FREE if $B \subseteq C_{free}$, and $C(B) =$ STUCK if B is in the interior of $C_{space} \setminus C_{free}$. Otherwise, $C(B) =$ MIXED. One of our goals is to introduce classifications $\widetilde{C}(B)$ that are "soft versions" of $C(B)$ (see the full paper [16] Appendix A).

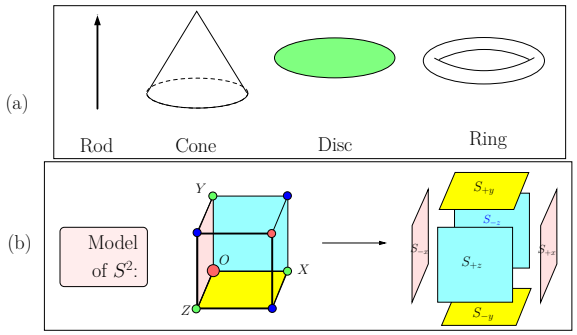

Figure 2 (a) 3D rigid robots with $C_{space} = \mathbb{R}^3 \times S^2$. (b) Subdivision atlas for S^2 via $\widehat{S^2}$.

We present four desiderata in path planning:

(G0) the planner must be mathematically rigorous and complete;
(G1) it must have correct implementations which are also:
(G2) relatively easy to achieve and
(G3) practically efficient.

In (G0), we use the standard Computer Science notion of an algorithm being **complete** if (a) it is partially complete[1] and (b) it halts. The notions of **resolution completeness** and **probabilistic completeness** in robotics have requirement (a) but not (b). In probabilistic-complete algorithms, halting with `NO-PATH` is achieved heuristically by putting limits on time and/or number of samples. But such limits are not intrinsic to the input instance. In resolution-complete algorithms, `NO-PATH` halting is based on width w of subdivision box being small enough (say $w < \varepsilon$). One issue is that the width of a box is a direct measure of clearance (but there is a nontrivial correlation); secondly, box predicates are numerical and "accurate enough" (σ-effective in our theory). These issues are exacerbated when algorithms do not use box predicates, but perform sampling at grid points of the subdivision. In contrast, our `NO-PATH` guarantees an intrinsic property: there is no path of clearance $K\varepsilon$ (see below).

But desideratum (G0) is only the base line. A (G0)-planner may not be worth much in a practical area like robotics unless it also has implementations with properties (G1-G3). E.g., the usual exact algorithms satisfy (G0) but their typical implementations fail (G1). With proper methods [30], it is possible to satisfy (G1); Halperin et al [13] give such solutions in 2D using `CGAL`. Both (G0) and (G1) can be formalized (see next), but (G2) and (G3) are informal. The robotics community has developed various criteria to evaluate (G2) and (G3). The accepted practice is having an implementation (proving (G2)) that achieves "real time" performance on a suite of non-trivial instances (proving (G3)).

The main contribution of this paper is the design of planners for spatial robots with 5 DOFs that have the "good" properties (G0-G3). This seems to be the first for such robots. To achieve our results, we introduce theoretical innovations and algorithmic techniques that may prove to be more widely applicable.

In path planning and in Computational Geometry, there is a widely accepted interpretation of desideratum (G0): it is usually simply called "exact algorithms". But to stress our interest in alternative notions of exactness, we refer to the standard notion as **exact (unqualified)**. Planners that are exact (unqualified) are first shown in [26]; this can be viewed as a fundamental result on decidability of connectivity in semi-algebraic sets [1]. The curse of

[1] Partial completeness means the algorithm produces a correct output *provided* it halts.

exact (unqualified) algorithms is that the algorithm must detect any degeneracies in the input and handle them explicitly. But exact (unqualified) algorithms are rare, mainly because degeneracies are numerous and hard to analyze: the usual expedient is to assume "nice" (non-degenerate) inputs. So the typical exact (unqualified) algorithms in the literature are **conditional** algorithms, i.e., its correctness is conditioned on niceness assumptions. Such gaps in exact (unqualified) algorithms are not an issue as long as they are not implemented. For non-linear problems beyond 2D, complete degeneracy analysis is largely non-existent. This is vividly seen in the fact that, despite long-time interest, there is still no exact (unqualified) algorithm for the Euclidean Voronoi diagram of a polyhedral set (see [15, 12, 11, 34]). For similar reasons, unconditional exact (unqualified) path planners in 3D are unknown.

We now address (G1-G3). The typical implementation is based on machine arithmetic (the IEEE standard), which may satisfy (G2) but almost certainly not (G1). We regard this as a (G1-G2) trade-off. In fact, our implementations here as well as in our previous papers [32, 20, 33] are such machine implementations. This follows the practice in the robotics community, in order to have a fair comparison against other implementations. Below, we shall expand on our claims about (G1-G3) including how to achieve theoretically correct implementation (G1). What makes this possible is our replacement of "exact (unqualified)" planners by "exact (up to resolution)" planners, defined below:

Resolution-Exact Path Planning for robot R_0:
Input: $(\alpha, \beta, \Omega; B_0, \varepsilon)$
where $\alpha, \beta \in C_{space}(R_0)$ is the start and goal, $\Omega \subseteq \mathbb{R}^3$
the obstacle set, $B_0 \subseteq C_{space}(R_0)$ is a box, and $\varepsilon > 0$.
Output: Halt with either an Ω-free path from α to β in B_0,
or `NO-PATH` satisfying the conditions (P) and (N) below.

The **resolution-exact planner** (or, ε-exact planner) has an **accuracy constant** $K > 1$ (independent of input) such that its output satisfies two conditions:

- (P) If there is a path (from α to β in B_0) of clearance $K\varepsilon$, the output *must* be[2] a path.
- (N) If there is no path in B_0 of clearance ε/K, the output *must* be `NO-PATH`.

Here, clearance of a path is the minimum separation of the obstacle set Ω from the robot's footprint on the path. Note that the preconditions for (P) and (N) are not exhaustive: in case the input fails both preconditions, our planner may either output a path or `NO-PATH`. This indeterminacy is essential to escape from exact computation (and arguably justified for robotics [36]). The constant $K > 1$ is treated in more detail in [32, 37]. But resolution-exactness is just a definition. How do we design such algorithms? We propose to use subdivision, and couple with **soft predicates** to exploit resolution-exactness. We replace the classification $C(B)$ by a soft version $\widetilde{C}(B)$ [32]. This leads to a general resolution-exact planner which we call **Soft Subdivision Search** (SSS) [36, 37] that shares many of the favorable properties of sampling planners (and unlike exact planners). We demonstrated in [32, 20, 33] that for planar robots with up to 4 DOFs, our planners can consistently outperform state-of-the-art sampling planners.

1.1 What is New: Contributions of This Paper

In this work, we design ε-exact planners for rods and rings, with accompanying implementation that addresses the desiderata (G0-G3). This fulfills a long-time challenge in robotics. We are able to do this because of the twin foundations of resolution-exactness and soft-predicates.

[2] For simplicity, we do not require the output path to have any particular clearance, but we could require clearance $\geq \varepsilon/K$ as in [32].

Although we had already used this foundation to implement a variety of planar robots [32, 20, 33, 39] that can match or surpass state-of-the-art sampling methods, it was by no means assured that we can extend this success to 3D robots. Indeed, the present work required a series of technical innovations: **(I)** One major technical difference from our previous work on planar robots is that we had to give up the notion of "forbidden orientations" (which seems "forbidding" for 3D robots). We introduced an alternative approach based on the "safe-and-effective" approximation of footprint of boxes. We then show how to achieve such approximations for the rod and ring robots separately. **(II)** The approximated footprints of boxes are represented by what we call Σ_2-sets (Sec. 4.1); this representation supports desideratum (G2) for easy implementation. One side benefit of Σ_2-sets is that they are very flexible; thus, we can now easily extend our planners to "thick" versions of the rod or ring. In contrast, the forbidden orientation approach requires non-trivial analysis to justify the "thick" version [33]. The trade-off in using Σ_2-sets is a modest increase in the accuracy constant K. **(III)** We also need good representations of the 5-DOF configuration space. Here we introduce the square model of S^2 to avoid the singularities in the usual spherical polar coordinates [19], **and also** to support subdivision in non-Euclidean spaces. **(IV)** Not only is the geometry in 3D more involved, but the increased degree of freedom requires new techniques to further improve efficiency. Here, the search heuristic based on Voronoi diagrams becomes critical to achieve real-time performance (desideratum (G3)).

Overview of the Paper

Section 2 is a brief literature review. Section 3 explains an essential preliminary to doing subdivision in S^2. Sections 4–6 describe our techniques for computing approximate footprints of rods and rings. We discuss efficiency and experimental results in Section 7. We conclude in Section 8. All proofs and some background are given in the full paper [16] (in appendices).

2 Literature Review

Halperin et al [14] gave a general survey of path planning. An early survey is [35] where two universal approaches to exact path planning were described: cell-decomposition [25] and retraction [24, 23, 4]. Since exact path planning is a semi-algebraic problem [26], it is reducible to general (double-exponential) cylindrical algebraic decomposition techniques [1]. But exploiting path planning as a connectivity problem yields singly-exponential time (e.g, [10]). The case of a planar rod (called "ladder") was first studied in [25] using cell-decomposition. More efficient (quadratic time) methods based on the retraction method were introduced in [28, 29]. On-line versions for a planar rod are also available [7, 6].

Spatial rods were first treated in [27]. The combinatorial complexity of its free space is $\Omega(n^4)$ in the worst case and this can be closely matched by an $O(n^{4+\epsilon})$ time algorithm [17]. The most detailed published planner for a 3D rod is Lee and Choset [19]. They use a retraction approach. The paper exposes many useful and interesting details of their computational primitives (see its appendices). In particular, they follow a Voronoi edge by a numerical path tracking. But like most numerical code, there is no a priori guarantee of correctness. Though the goal is an exact path planner, degeneracies are not fully discussed. Their two accompanying videos have no timing or experimental data.

One of the few papers to address the non-existence of paths is Zhang et al [38]. Their implementation work is perhaps the closest to our current work, using subdivision. They noted that "no good implementations are known for general robots with higher than three DOFs". They achieved planners with 3 and 4 DOFs (one of which is a spatial robot). Although their planners can detect `NO-PATH`, they do not guarantee detection (this is impossible without exact computation).

3 Subdivision Charts and Atlas for S^2

Terminology. We fix some terminology for the rest of the paper. The fundamental **footprint map** Fp from configuration space $C_{space} = C_{space}(R_0)$ to subsets of $\mathbb{R}^3$ was introduced above. If $B \subseteq C_{space}$ is any set of configurations, we define $Fp(B)$ as the union of $Fp(\gamma)$ as γ ranges over B. Typically, B is a "box" of C_{space} (see below for its meaning in non-Euclidean space S^2). We may assume $\Omega \subseteq \mathbb{R}^3$ is regular (i.e., equal to the closure of its interior). Although Ω need not be bounded (e.g., it may be the complement of a box), we assume its boundary $\partial(\Omega)$ is a bounded set. Then $\partial(\Omega)$ is partitioned into a set of **(boundary) features**: **corners** (points), **edges** (relatively open line segments), or **walls** (relatively open triangles). Let $\Phi(\Omega)$ denote the set of features of Ω. The (minimal) set of corners and edges is uniquely defined by Ω, but walls depend on a triangulation of $\partial\Omega$. If $A, B \subseteq \mathbb{R}^3$, define their **separation** $\mathrm{Sep}(A, B) := \inf\{\|a - b\| : a \in A, b \in B\}$ where $\|a\|$ is the Euclidean norm. The **clearance** of γ is $\mathrm{Sep}(Fp(\gamma), \Omega)$. Say γ is **Ω-free** (or simply **free**) if it has positive clearance. Let $C_{free} = C_{free}(\Omega)$ be the set of Ω-free configurations. The **clearance** of a path $\mu : [0, 1] \to C_{space}$ is the minimum clearance attained by $\mu(t)$ as t ranges over $[0, 1]$.

Subdivision in Non-Euclidean Spaces. Our C_{space} has an Euclidean part ($\mathbb{R}^3$) and a non-Euclidean part (S^2). We know how to do subdivision in $\mathbb{R}^3$ but it is less clear for S^2. Non-Euclidean spaces can be represented either (1) as a submanifold of $\mathbb{R}^m$ for some m (e.g., $SO(3) \subseteq \mathbb{R}^9$ viewed as orthogonal matrices) or (2) as a subset of $\mathbb{R}^m$ subject to identification (in the sense of quotient topology [21]). A common representation of S^2 (e.g., [19]) uses a pair of angles (i.e., spherical polar coordinates) $(\theta, \phi) \in [0, 2\pi] \times [-\pi/2, \pi/2]$ with the identification $(\theta, \phi) \equiv (\theta', \phi')$ iff $\{\theta, \theta'\} = \{0, 2\pi\}$ or $\phi = \phi' = \pi/2$ (North Pole) or $\phi = \phi' = -\pi/2$ (South Pole). Thus an entire circle of values θ is identified with each pole, causing severe distortions near the poles which are singularities. So the numerical primitives in [19, Appendix F] have severe numerical instabilities.

To obtain a representation of S^2 without singularities, we use the map [37]

$$q \in \mathbb{R}^3 \mapsto \widehat{q} := q/\|q\|_\infty$$

whose range is the boundary of a 3D cube $\widehat{S^2} := \partial([-1, 1]^3)$. This map is a bijection when its domain is restricted to S^2, with inverse map $q \in \widehat{S^2} \mapsto \overline{q} := q/\|q\|_2 \in S^2$. Thus $\overline{\widehat{q}}$ is the identity for $q \in S^2$. We call $\widehat{S^2}$ the **square model** of S^2. We view S^2 and $\widehat{S^2}$ as metric spaces: S^2 has a natural metric whose geodesics are arcs of great circles. The geodesics on S^2 are mapped to the corresponding polygonal geodesic paths on $\widehat{S^2}$ by $q \mapsto \widehat{q}$. Define the constant

$$C_0 := \sup_{p \neq q \in S^2} \left\{ \max \left\{ \frac{d_2(p, q)}{\widehat{d}_2(\widehat{p}, \widehat{q})}, \frac{\widehat{d}_2(\widehat{p}, \widehat{q})}{d_2(p, q)} \right\} \right\}$$

where d_2 and $\widehat{d}_2$ are the metrics on S^2 and $\widehat{S^2}$ respectively. Clearly $C_0 \geq 1$. Intuitively, C_0 is the largest distortion factor produced by the map $q \mapsto \widehat{q}$ (by definition the inverse map has the same factor).

▶ **Lemma 1.** $C_0 = \sqrt{3}$.

The proof in the full paper [16] Appendix B.1 also shows that the worst distortion is near the corners of $\widehat{S^2}$. The constant C_0 is one of the 4 constants that go into the ultimate accuracy constant K in the definition of ε-exactness (see [37] for details).

It is obvious how to do subdivision in $\widehat{S^2}$. This is illustrated in Figure 2(b). After the first subdivision of $\widehat{S^2}$ into 6 faces, subsequent subdivision is just the usual quadtree subdivision of each face. We interpret the subdivision of $\widehat{S^2}$ as a corresponding subdivision of S^2. In [37], we give the general framework using the notion of **subdivision charts and atlases** (borrowing terms from manifold theory).

4 Approximate Footprints for Boxes in $\mathbb{R}^3 \times S^2$

We focus on soft predicates because, in principle, once we have designed and implemented such a predicate, we already have a rigorous and complete planner within the **Soft Subdivision Search** (SSS) framework [32, 37]. (The SSS framework is summarized in the full paper [16] Appendix A.) As noted in the introduction, our soft predicate $\widetilde{C}$ classifies any input box $B \subseteq C_{space}$ into one 3 possible values. A key idea of our 2-link robot work [20, 33] is the notion of "forbidden orientations" (of a box B, in the presence of Ω). The same concept may be attempted for $\mathbb{R}^3 \times S^2$, except that the details seem to be formidable to analyze and to implement. Instead, this paper introduces a direct approximation of the **footprint of a box**, $Fp(B) := \bigcup\{Fp(\gamma) : \gamma \in B\}$. We now introduce $\widetilde{Fp}(B) \subseteq \mathbb{R}^3$ as the **approximate footprint**, and discuss its properties. This section is abstract, in order to expose the mathematical structure of what is needed to achieve resolution-exactness for our planners. The reader might peek at the next two sections to see the instantiations of these concepts for the rod/ring robot.

To understand what is needed of this approximation, recall that our approach to soft predicates is based on the "method of features" [32]. The idea is to maintain a set $\widetilde{\phi}(B)$ of approximate features for each box B. We softly classify B as $\widetilde{C}(B) = \texttt{MIXED}$ as long as $\widetilde{\phi}(B)$ is non-empty; otherwise, we can decide whether $\widetilde{C}(B) = C(B)$ is $\texttt{FREE}$ or $\texttt{STUCK}$. This decision is relatively easy in 2D, but is more involved in 3D and detailed in the full paper [16] Appendix B.2. For correctness of this procedure, we require

$$\widetilde{\phi}(B/\sigma) \subseteq \phi(B) \subseteq \widetilde{\phi}(B). \tag{1}$$

Here $\sigma > 1$ is some global constant and "B/σ" denotes the box B shrunk by factor σ. Basically, (1) guarantees that our soft predicate $\widetilde{C}(B)$ is conservative and σ-effective (i.e., if B is free then $\widetilde{C}(B/\sigma) = \texttt{FREE}$). For computational efficiency, we want the approximate feature sets to have **inheritance property**, i.e.,

$$\widetilde{\phi}(B) \subseteq \widetilde{\phi}(parent(B)). \tag{2}$$

We now show what this computational scheme demands of our approximate footprint. Define the **exact feature set** of box B as usual: $\phi(B) := \{f \in \Phi(\Omega) : f \cap Fp(B) \neq \emptyset\}$ and (tentatively) the **approximate feature set** of box B as

$$\widetilde{\phi}(B) := \left\{f \in \Phi(\Omega) : f \cap \widetilde{Fp}(B) \neq \emptyset\right\}. \tag{3}$$

The important point is that $\widetilde{Fp}(B)$ is defined prior to $\widetilde{\phi}(B)$. We need the fundamental inclusions

$$\widetilde{Fp}(B/\sigma) \subseteq Fp(B) \subseteq \widetilde{Fp}(B). \tag{4}$$

Note that this immediately implies (1). Unfortunately, (3) and (4) together do not guarantee inheritance, i.e., (2). Instead, we define $\widetilde{\phi}'(B)$ recursively as follows:

$$\widetilde{\phi}'(B) := \begin{cases} \left\{f \in \Phi(\Omega) : f \cap \widetilde{Fp}(B) \neq \emptyset\right\} & \text{if } B \text{ is the root,} \\ \left\{f \in \widetilde{\phi}'(parent(B)) : f \cap \widetilde{Fp}(B) \neq \emptyset\right\} & \text{else.} \end{cases} \tag{5}$$

Notice that this only defines $\widetilde{\phi}'(B)$ when B is an aligned box (i.e., obtained by recursive subdivision of the root box). But B/σ is never aligned when B is aligned, and thus $\widetilde{\phi}'(B/\sigma)$ is not captured by (5). Therefore we introduce a parallel definition:

$$\widetilde{\phi}'(B/\sigma) := \begin{cases} \left\{ f \in \Phi(\Omega) : f \cap \widetilde{Fp}(B/\sigma) \neq \emptyset \right\} & \text{if } B \text{ is the root,} \\ \left\{ f \in \widetilde{\phi}'(parent(B)/\sigma) : f \cap \widetilde{Fp}(B/\sigma) \neq \emptyset \right\} & \text{else.} \end{cases} \tag{6}$$

Now, $\widetilde{\phi}'(B)$ satisfies (2). But does it satisfy (1), which is necessary for correctness? This is answered affirmatively by the following lemma (proved in the full paper [16] Appendix B.3):

▶ **Lemma 2.** *If the approximate footprint $\widetilde{Fp}(B)$ satisfies Eq. (4), then $\widetilde{\phi}'(B)$ satisfies Eq. (1), i.e.,*

$$\widetilde{\phi}'(B/\sigma) \subseteq \phi(B) \subseteq \widetilde{\phi}'(B).$$

Since $\widetilde{\phi}'(B)$ has all the properties we need, we have no further use for the definition of $\widetilde{\phi}(B)$ given in (3). Henceforth, we simply write "$\widetilde{\phi}(B)$" to refer to the set $\widetilde{\phi}'(B)$ defined in (5) and (6).

Geometric Notations. We will be using planar concepts like circles, squares, etc, for sets that lie in some plane of $\mathbb{R}^3$. We shall call them **embedded** circles, squares, etc. By definition, if X is an embedded object then it defines a unique plane $Plane(X)$ (unless X lies in a line). Let $Ball(r, c) \subseteq \mathbb{R}^3$ denote a ball of radius r centered at c. If c is the origin, we simply write $Ball(r)$. Suppose $X \subseteq \mathbb{R}^3$ is any non-empty set. Let $Ball(X)$ denote the **circumscribing ball** of X, defined as the smallest ball containing X. Next, if $c \notin X$ then $Cone(c, X)$ denotes the union of all the rays from c through points in X, called the **cone** of X with **apex** c. We consider two cases of X in this cone definition: if X is a ball, then $Cone(c, X)$ is called a **round cone**. If the radius of ball X is r and the distance from the center of $Ball(X)$ to c is $h \geq r$, then call $\arcsin(r/h)$ the **half-angle** of the cone; note that the angle at the the apex is twice this half-angle. If X is an embedded square, we call $Cone(c, X)$ a **square cone**, and the ray from c through the center of the square is called the **axis** of the square cone. If P is any plane that intersects the axis of a square cone $Cone(c, X)$, then $P \cap Cone(c, X)$ is a square iff P is parallel to square X. A **ring** (resp., **cylinder**) is the Minkowski sum of an embedded circle (resp., a line) with a ball. Finally consider a box $B = B^t \times B^r \subseteq \mathbb{R}^3 \times \widehat{S^2}$ where B^t and B^r are the **translational** and **rotational** components of B, and B^r is either $\widehat{S^2}$ or a subsquare of a face of $\widehat{S^2}$. We let m_B and r_B denote the center and radius (distance from the center to any corner) of B^t. The **cone of** B, denoted $Cone(B)$, is the round cone $Cone(m_B, Ball(m_B + B^r))$. If the center of square $m_B + B^r$ is c and width of B^r is w, then $Cone(B)$ is just $Cone(m_B, Ball(c, w/\sqrt{2}))$.

4.1 On Σ_2-Sets

Besides the above inclusion properties of $\widetilde{Fp}(B)$, we also need to decide if $\widetilde{Fp}(B)$ intersects a given feature f. We say $\widetilde{Fp}(B)$ is "nice" if there are intersection algorithms that are easy to implement (desideratum G2) and practically efficient (desideratum G3). We now formalize and generalize some "niceness" properties of $\widetilde{Fp}(B)$ that were implicit in our previous work ([32, 20, 33], especially [39]).

An **elementary set** (in $\mathbb{R}^3$) is defined to be one of the following sets or their complements: half space, ball, ring, cone or cylinder. Let $\mathcal{E}$ (or $\mathcal{E}_3$) denote the set of elementary sets in $\mathbb{R}^3$. In $\mathbb{R}^2$, we have a similar notion of elementary sets $\mathcal{E}_2$ comprising half-planes, discs or their complements. All these elementary sets are defined by a single polynomial inequality – so

technically, they are all "algebraic half-spaces". The sets in $\mathcal{E}$ are evidently "nice" (niceness of a ring has some subtleties – see Sec. 6). We next extend our collection of nice sets: define a Π_1**-set** to be a finite intersection of elementary sets. We regard a Π_1-set $S = \cap_{i=1}^{n} S_i$ to be "nice" because we can easily check if a feature f intersects S by a simple while-loop (see below). Notice that Π_1 contains all convex polytopes in $\mathbb{R}^3$. Our definitions of $\widetilde{Fp}(B)$ in [32, 20, 33] are all Π_1-sets. But in [39], we make a further extension: define a Σ_2**-set** to be a finite union of the Π_1-sets, i.e., each Σ_2-set S has the form

$$S = \bigcup_{i=1}^{n} \bigcap_{j=1}^{m_i} S_{ij} \tag{7}$$

where S_{ij}'s are elementary sets. We still say such an S is "nice" since checking if a feature f intersects S can be written in a doubly-nested loop (see below). Although this intersection is more expensive to check than with a Π_1-set, it may result in fewer subdivisions and better efficiency in the overall algorithm. Thus, there is an accuracy-efficiency trade-off. *Good approximations of footprints are harder to do accurately in 3D, and the extra power of* Σ_2 *seems critical.*

We can put all these in the framework of a well-known[3] construction of an infinite hierarchy of sets, starting from some initial collection of sets. If Δ is any collection of sets, let $\Pi(\Delta)$ denote the collection of finite intersections of sets in Δ; similarly, $\Sigma(\Delta)$ denotes the collection of finite unions of sets in Δ. Then, starting with any collection Δ_1 of sets, define the infinite hierarchy of sets:

$$\Sigma_i, \Pi_i, \Delta_i \qquad (i \geq 1) \tag{8}$$

where $\Sigma_i := \Sigma(\Delta_i)$, $\Pi_i := \Pi(\Delta_i)$, and $\Delta_{i+1} := \Sigma_i \cup \Pi_i$. An element of Σ_i or Π_i is simply called a Σ_i-set or a Π_i-set.

We call (7) a Σ_2**-decomposition** of S, where $\Delta_1 := \mathcal{E}$. Note that this decomposition may not be unique, but in the cases arising from our simple robots, there is often an obvious optimal description. Moreover, n and m_i's are small constants. We can construct new sets by manipulating such a decomposition, e.g., replacing each S_{ij} by its τ**-expansion**, i.e., $S_{ij} \oplus Ball(\tau)$ (where $\oplus$ denotes the Minkowski sum), which remains elementary. Under certain conditions, the corresponding set is a reasonable approximation to $S \oplus Ball(\tau)$. If so, we can generalize the corresponding soft predicate to robots with thickness τ.

Once we have a Σ_2-decomposition of $\widetilde{Fp}(B)$, we can implement the intersection test with relative ease (G2) and quite efficiently (G3). For instance we can test intersection of the set S in (7) with a feature f by writing a doubly nested loop. At the beginning of the inner loop, we can initialize a set f_0 to f. Then the inner loop amounts to the update "$f_0 \leftarrow f_0 \cap S_{ij}$" for $j = 1, \ldots, m_i$. If ever f_0 becomes empty, we know that the set $S_i = \bigcap_{j=1}^{m_i} S_{ij}$ has empty intersection with f. The possibility of such representations is by no means automatic but in the next two sections we verify that they can be achieved for our rod and ring robots. These sections make our planners fully "explicit" for an implementation.

5 Soft Predicates for a Rod Robot

In this section, R_0 is a rod with length r_0; we choose one endpoint of the rod as the rotation center. Let $B = B^t \times B^r \subseteq \mathbb{R}^3 \times \widehat{S^2}$ be a box. Our main goal is to define approximate footprint $\widetilde{Fp}(B)$, and to prove the basic inclusions in Eqs. (4) and (1). This turns out to be a Π_1-set (we also indicate a more accurate Σ_2-set.)

[3] From mathematical analysis, constructive set theory and complexity theory.

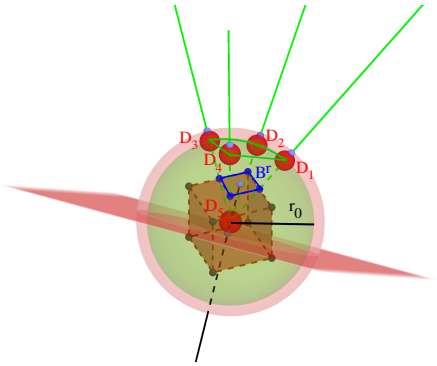

Figure 3 Rod Robot: $Fp(B) = Fp_0(B) \oplus (B^t - m_B)$ where $Fp_0(B)$ is indicated by the four green rays.

It is useful to define the **inner footprint** of B, $Fp_0(B)$, as $Fp(m_B \times B^r)$.
This set is the intersection of a ball and a square cone:

$$Fp_0(B) = Ball(r_0, m_B) \cap Cone(m_B, B^r + m_B). \tag{9}$$

The edges of this square cone is shown as green lines in Figure 3; furthermore, the brown box is $\widehat{S^2} + m_B$ (translation of $\widehat{S^2}$ so that it is centered at m_B). Note that the box footprint $Fp(B)$ is the Minkowski sum of $Fp_0(B)$ with $B^t - m_B$ (the translation of B^t to make it centered at the origin). It is immediate that

$$Fp_0(B) \subseteq Cone(B).$$

Thus we may write $Cone(m_B, B^r + m_B)$ as the intersection of four half spaces H_i $(i = 1, \ldots, 4)$. Let $Cone^{(+r_B)}(m_B, B^r + m_B)$ denote the intersection of the expanded half-spaces, $H_i \oplus Ball(r_B)$ $(i = 1, \ldots, 4)$. In general, $Cone^{(+r_B)}(m_B, B^r + m_B)$ is not a cone (it may not have a unique "apex"). Similarly we "expand" the inner footprint of (9) into

$$\text{“}\widetilde{Fp}(B)\text{”} := Ball(r_0 + r_B, m_B) \cap Cone^{(+r_B)}(m_B, B^r + m_B). \tag{10}$$

We use quotes for "$\widetilde{Fp}(B)$" in (10) because we view it as a candidate for an approximate footprint of B. Certainly, it has the desired property of containing the exact footprint $Fp(B)$. Unfortunately, this is not good enough. To see this, let θ be the half-angle of the round cone $Cone(B) = Cone(m_B, Ball(B^r + m_B))$. Then Hausdorff distance of "$\widetilde{Fp}(B)$" from $Fp(B)$ can be arbitrarily big as θ becomes arbitrarily small. Indeed θ can be arbitrarily small because it can be proportional to the input resolution ε. We conclude that such a planner is not resolution-exact. To fix this problem, we finally define

$$\widetilde{Fp}(B) := \text{“}\widetilde{Fp}(B)\text{”} \cap H_0 \tag{11}$$

where H_0 is another half space. A natural choice for H_0 is the half-space "above" the pink-color plane of Figure 3, defined as the plane normal to the axis of cone $Cone(B)$ and at distance r_B "below" m_B. We can also use the "horizontal" plane that is parallel to B^r and containing the "lower" face of B^t. We adopt this latter H_0 to have a simpler geometric structure.

This completes the description of $\widetilde{Fp}(B)$. It should be clear that checking if $\widetilde{Fp}(B)$ intersects any feature f is relatively easy (since it is even a Π_1-set). In the full paper [16] Appendix C we prove the following theorem:

▶ **Theorem 3.** *The approximate footprint $\widetilde{Fp}(B)$ as defined for a rod robot satisfies Eq. (4), i.e., there exists some fixed constant $\sigma > 1$ such that $\widetilde{Fp}(B/\sigma) \subseteq Fp(B) \subseteq \widetilde{Fp}(B)$.*

6 Soft Predicates for a Ring Robot

Let R_0 be a ring robot. Its footprint is an embedded circle of radius r_0. First we show how to compute $\mathrm{Sep}(C, f)$, the separation of an embedded circle C from a feature f. This was treated in detail by Eberly [9]. This is easy when f is a point or a plane. When f is a line, Eberly gave two formulations: they reduce to solving a system of 2 quadratic equations in 2 variables, and hence to solving a quartic equation; see the full paper [16] Appendix D.1 for more details. The predicate "Does f intersect $C \oplus Ball(r')$, a ring of thickness r'?" is needed later; it reduces to "Is $\mathrm{Sep}(C, f) \leq r'$?".

Our next task is to describe an approximate footprint, First recall the round cone of box B defined in the previous section: $Cone(B) = Cone(m_B, Ball(m_B + B^r))$. Let $\theta = \theta(B)$ be the half-angle of this cone, and c the center of B^r. Here, we think of c as a point of $\widehat{S^2}$, and define $\gamma(B) := m_B \times c$ viewed as an element of $\mathbb{R}^3 \times \widehat{S^2}$. Call $\gamma(B)$ the **central configuration** of box B. Let $Ray(B)$ be the ray from m_B through $m_B + c$. If $Plane(B)$ is the plane through m_B and normal to $Ray(B)$, then the footprint $Fp(\gamma(B))$ is an embedded circle lying in $Plane(B)$. We define the **inner footprint** of B as $Fp_0(B) := Fp(m_B \times B^r)$. The map $q \mapsto \overline{q}$ is the inverse of $q \mapsto \widehat{q}$, taking $c \in \widehat{S^2}$ to $\overline{c} \in S^2$. It is hard to work with $Fp_0(B)$. Instead consider the set $D(B)$ of all points in S^2 whose distance[4] from $\overline{c}$ is at most $\theta(B)$. So $D(B)$ is the intersection of S^2 with a round cone with ray from the origin to c. Then we have $Fp_0(B) \subseteq Fp_1(B)$ where

$$Fp_1(B) := Fp(m_B \times D(B)). \tag{12}$$

Our main computational interest is the approximate footprint of B defined as

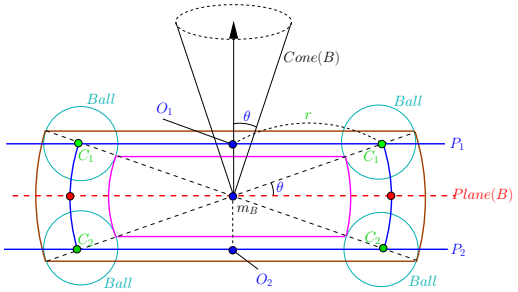

Figure 4 Ring Robot: central cross-section of $Fp_1(B)$ appears as two blue arcs. $\widetilde{Fp}(B)$ equals the union of two "thick rings" and a "truncated annulus". The axis of $Cone(B)$ is shown as a vertical ray. Each $Ball$ has radius r_B.

$$\widetilde{Fp}(B) := Fp_1(B) \oplus Ball(r_B). \tag{13}$$

Note that $Fp_1(B)$ has a simple geometric description. We illustrate this in Figure 4 using a central cross-section with a plane through m_B containing the axis of $Cone(B)$ (the axis of $Cone(B)$ is drawn vertically). The footprint of $\gamma(B)$ is a circle that appears as two red dots in the horizontal line (i.e., $Plane(B)$). Let $S^2(m_B, r_0)$ denote the 2-sphere centered at m_B with radius r_0. Then $Fp_1(B)$ is the intersection of $S^2(m_B, r_0)$ with a slab (i.e., intersection of two half-spaces whose bounding planes P_1 and P_2 are parallel to $Plane(B)$). These planes

[4] Recall that S^2 is a metric space whose geodesics are arcs of great circles.

appear as two horizontal blue lines in Figure 4. In the cross section, $Fp_1(B)$ are seen as two blue circular arcs. For $i = 1, 2$, let $C_i = P_i \cap S^2(m_B, r_0)$; it is an embedded circle that appears as a pair of green points in Figure 4. Each C_i is centered at O_i, with radius $r = r_0 \cos\theta$; see Figure 4.

We can now describe a Σ_2-decomposition of $\widetilde{Fp}(B)$: it is the union of two "thick rings", $C_1 \oplus Ball(r_B)$ and $C_2 \oplus Ball(r_B)$ (both of thickness r_B), and a shape $Ann(B)$ which we call a **truncated annulus**. First of all, the region bounded between the spheres $S^2(m_B, r_0 + r_B)$ (the brown arcs in the figure) and $S^2(m_B, r_0 - r_B)$ (the magenta arcs) is called a (solid) annulus. Let C_i^* denote the embedded disc whose relative boundary is C_i. Then we have two round cones, $Cone(m_B, C_1^*))$ and $Cone(m_B, C_2^*))$. Together, they form a *double cone* that is actually a simpler object for computation! Finally, define $Ann(B)$ to be the intersection of the annulus with the complements of the double cone.

For each thick ring $C_i \oplus Ball(r_B)$, deciding "Does a feature f intersect $C_i \oplus Ball(r_B)$?" is equivalent to "Is $\text{Sep}(C_i, f) \leq r_B$?" (see beginning of this section). In the full paper [16], Appendix D.1 discusses this computation and Appendix D.2 proves the following theorem:

▶ **Theorem 4.** *The approximate footprint $\widetilde{Fp}(B)$ as defined for a ring robot satisfies Eq. (4), i.e., there exists some fixed constant $\sigma > 1$ such that $\widetilde{Fp}(B/\sigma) \subseteq Fp(B) \subseteq \widetilde{Fp}(B)$.*

7 Practical Efficiency of Correct Implementations

We have developed ε-exact planners for rod and ring robots. We have explicitly exposed all the details necessary for a correct implementation, i.e., criterion (G1). The careful design of the approximate footprints of boxes as Σ_2-sets ensures (G2), i.e., it would be relatively easy to implement. We now address (G3) or practical efficiency. For robots with 5 or more DOFs, it becomes extremely critical that good search strategies are deployed. In this paper, we have found that some form of Voronoi heuristic is extremely effective: the idea is to find paths along Voronoi curves (in the sense of [24, 28]), and exploit subdivision Voronoi techniques based (again) on the method of features [34, 2]. There are subtleties necessitating the use of pseudo-Voronoi curves [19, 28, 29]. Since we do not rely on Voronoi heuristics for correctness, simple expedients are available. To recognize Voronoi curves, we maintain (in addition to the collision-detection feature set $\widetilde{\phi}(B)$), the **Voronoi feature set** $\widetilde{\phi}_V(B)$. These two sets have some connection but there are no obvious inclusion relationships.

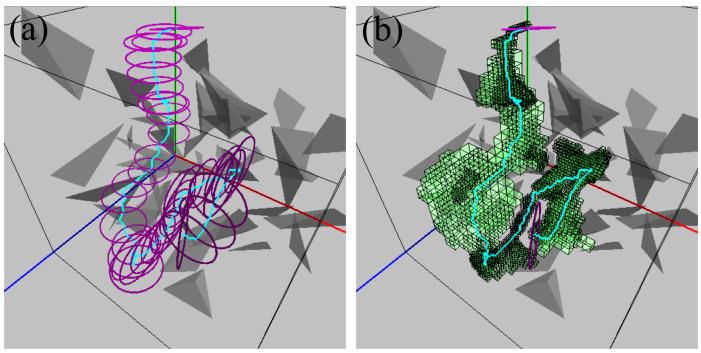

Figure 5 Ring robot amidst 40 random tetrahedra: (a) trace of a found path; (b) subdivision of translational boxes on the path.

Our current implementation achieves near real-time performance (see video `http://cs.nyu.edu/exact/gallery/rod-ring/rod_ring.html`). Table 1 summarizes ex-

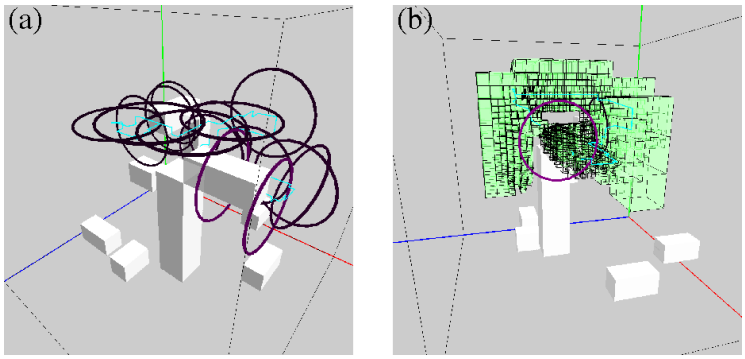

Figure 6 Ring robot amidst pillars and L-shaped posts (Posts): (a) trace of a found path; (b) subdivision of translational boxes on the path.

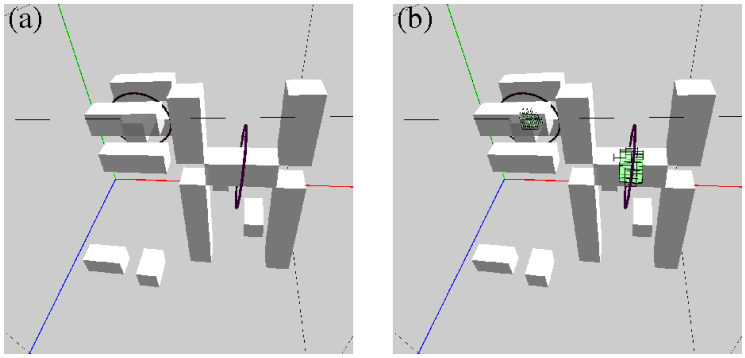

Figure 7 Ring robot amidst another set of pillars and posts (Posts2): (a) start and goal configurations (no path found); (b) subdivision of translational boxes during the search.

Table 1 Rod and Ring Experiments.

Rod Robot								
Exp. #	Envir.	Length	ε	Start Conf.	Goal Conf.	Path	Time (s)	#Boxes (K)
1/2	Rand100	120	16/8	(240, 120, 360, -0.5, -0.5, -1)	(220, 50, 80, 0.1, 0.8, 1)	Y/Y	1.05/2.82	8.1/22.1
3/4	Rand100	120	16/8	(400, 60, 380, -1, 0, 0)	(200, 200, 240, 0, 1, 0)	Y/Y	1.43/3.92	19.3/62.2
5/6	Rand40	160	16/8	(80, 32, 480, 0, 0, -1)	(240, 440, 200, 1, 0, 0)	Y/Y	16.12/90.65	244.7/ 1138.2
7/8	Rand40	160	16/8	(400, 480, 80, 0, -1, 0)	(30, 80, 480, 0.5, 0.1, -1)	Y/Y	14.54/9.4	217.5/113.0
9/10	Posts	60	16/8	(160, 480, 190, 0, 0.1, -1)	(390, 60, 420, 1, 0, 0)	Y/Y	0.07/0.13	2.1/3.7
11/12	Posts	60	16/8	(320, 120, 320, 0, 1, 0)	(200, 360, 60, 0, -1, 0)	N/N	1.77/242.7	25.6/3790.3
Ring Robot								
Exp. #	Envir.	Radius	ε	Start Conf.	Goal Conf.	Path	Time (s)	#Boxes (K)
1/2	Rand100	40	16/8	(240, 120, 360, -0.5, -0.5, -1)	(220, 50, 80, 0.1, 0.8, 1)	Y/Y	0.66/0.57	2.72/2.63
3/4	Rand100	40	16/8	(400, 60, 380, -1, 0, 0)	(160, 240, 240, 0, 1, 0)	Y/Y	0.25/0.24	0.92/0.92
5/6	Rand40	60	16/8	(80, 120, 480, 0, 0, -1)	(240, 440, 200, 1, 0, 0)	Y/Y	9.38/30.61	38.37/71.05
7/8	Rand40	60	16/8	(400, 480, 80, 0, -1, 0)	(100, 80, 480, 0.5, 0.1, -1)	Y/Y	2.07/3.76	10.61/13.90
9/10	Posts	60	16/8	(200, 320, 190, 0, 0.1, -1)	(390, 320, 320, 1, 0, 0)	Y/Y	35.68/89.7	114.3/139.3
11/12	Posts2	60	16/8	(410, 90, 190, 0, 0.1, -1)	(315, 220, 325, 0, 1, 0)	N/N	7.03/271.3	8.1/539.1

periments on our rod and ring robots. The environments Rand100, Rand40 (100 and 40 random tetrahedra), Posts and Posts2 are shown in Figs. 1, 5, 6 and 7. The dimensions of the environments are 512^3. Our implementation uses `C++` and OpenGL on the Qt platform. Our code, data and experiments are distributed[5] with our open source `Core Library`. We ran our experiments on a MacBook Pro under Mac OS X 10.10.5 with a 2.5 GHz Intel Core i7 processor, 16GB DDR3-1600 MHz RAM and 500GB Flash Storage. Details about these experiments are found in a folder in `Core Library` for this paper; a `Makefile` there can automatically run all the experiments. Thus these results are reproducible from the data there.

Table 2 (correlated with Table 1 by the Exp #'s) compares our methods with various sampling-based planners in OMPL [31], where we accepted the default parameters and each instance was run 10 times, with the "average time (in s)/standard deviation/success rate" reported. This comparison has various caveats: we simulated the rod and ring robots by polyhedral approximations. We usually outperform RRT in cases of PATH. In case of NO-PATH, we terminated in real time while all sampling methods timed out (300s).

Table 2 Comparison with Sampling Methods in OMPL (the best run-time is shown in bold).

Rod Robot									
Exp. #	Ours	PRM	Lazy PRM	RRT	Lazy RRT	RRT Connect	PDST	BFMT	Lazy Bi-KPIECE
1	1.05/Y	0.036/0.027/1	**0.017**/0.024/1	1.18/0.74/1	0.019/0.023/1	0.22/0.043/1	0.058/0.055/1	1.11/0.18/1	0.58/0.36/1
3	1.43/Y	0.05/0.047/1	0.028/0.019/1	1.73/0.82/1	**0.024**/0.024/1	0.23/0.023/1	0.1/0.056/1	1.51/0.2/1	0.59/0.28/1
5	16.12/Y	0.044/0.036/1	0.051/0.025/1	22.1/43/1	**0.036**/0.032/1	0.99/0.36/1	0.21/0.11/1	1.74/0.33/1	0.44/0.18/1
7	14.54/Y	0.077/0.04/1	**0.03**/0.02/1	10.31/6.08/1	0.033/0.023/1	1.26/0.5/1	0.26/0.2/1	1.74/0.32/1	0.5/0.21/1
9	0.07/Y	0.0058/0.002/1	0.0038/0.0044/1	1.17/0.78/1	**0.003**/0.002/1	0.084/0.017/1	0.025/0.025/1	0.3/0.053/1	0.065/0.032/1
11	**1.77**/N	300/0/0	300/0/0	300/0/0	300/0/0	300/0/0	300/0/0	300/0/0	300/0/0
Ring Robot									
Exp. #	Ours	PRM	Lazy PRM	RRT	Lazy RRT	RRT Connect	PDST	BFMT	Lazy Bi-KPIECE
1	0.66/Y	0.0057/0.0026/1	**0.0056**/0.005/1	1.15/0.93/1	0.0061/0.009/1	0.037/0.005/1	0.015/0.01/1	0.15/0.01/1	0.077/0.035/1
3	0.25/Y	0.0085/0.0067/1	**0.003**/0.002/1	24.16/54/1	0.012/0.0085/1	0.052/0.011/1	0.008/0.008/1	0.145/0.023/1	0.068/0.032/1
5	9.38/Y	0.019/0.014/1	**0.01**/0.004/1	300/0/0	150.07/10.05/0.5	0.53/0.03/1	0.057/0.023/1	0.22/0.04/1	0.093/0.022/1
7	2.07/Y	0.024/0.0066/1	**0.013**/0.006/1	3/1.49/1	0.068/0.007/1	1.46/0.14/1	0.059/0.044/1	0.27/0.048/1	0.12/0.027/1
9	35.68/Y	0.63/0.4/1	136.6/138.6/0.7	111.4/119.4/0.8	300/0/0	1.65/0.49/1	3.36/4.13/1	0.23/0.04/1	**0.097**/0.033/1
11	**7.03**/N	300/0/0	300/0/0	300/0/0	300/0/0	300/0/0	300/0/0	300/0/0	300/0/0

8 Conclusions

Path planning in 3D has many challenges. Our 5-DOF spatial robots have pushed the current limits of subdivision methods. To our knowledge there is no similar algorithm with comparable rigor or guarantees. Conventional wisdom says that sampling methods can achieve higher DOFs than subdivision. By an estimate of Choset et al [5, p. 202], sampling methods are limited to $5-12$ DOFs. We believe our approach can reach 6-DOF spatial robots. Since resolution-exactness delivers stronger guarantees than probabilistic-completeness, we expect a performance hit compared to sampling methods. But for simple planar robots (up to 4 DOFs) [32, 20, 33, 39] we observed no such trade-offs because we outperform state-of-the-art sampling methods (such as OMPL [31]) often by two orders of magnitude. But in the 5-DOF robots of this paper, we see that our performance is competitive with sampling methods. It is not clear to us that subdivision is inherently inferior to sampling (we can also do random subdivision). It is true that each additional degree of freedom is conquered only with effort and suitable techniques. This remark seems to cut across both subdivision and sampling approaches; but it hits subdivision harder because of our stronger guarantees.

[5] `http://cs.nyu.edu/exact/core/download/core/`

References

1 Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry*. Algorithms and Computation in Mathematics. Springer, 2nd edition, 2006.

2 Huxley Bennett, Evanthia Papadopoulou, and Chee Yap. Planar minimization diagrams via subdivision with applications to anisotropic Voronoi diagrams. *Eurographics Symposium on Geometric Processing*, 35(5), 2016. SGP 2016, Berlin, Germany. June 20-24, 2016.

3 Rodney A. Brooks and Tomas Lozano-Perez. A subdivision algorithm in configuration space for findpath with rotation. In *Proc. 8th Intl. Joint Conf. on Artificial intelligence - Volume 2*, pages 799–806, San Francisco, CA, USA, 1983. Morgan Kaufmann Publishers Inc.

4 John Canny. Computing roadmaps of general semi-algebraic sets. *The Computer Journal*, 36(5):504–514, 1993.

5 H. Choset, K. M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun. *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press, Boston, 2005.

6 Howie Choset, Brian Mirtich, and Joel Burdick. Sensor based planning for a planar rod robot: Incremental construction of the planar Rod-HGVG. In *IEEE Intl. Conf. on Robotics and Automation (ICRA'97)*, pages 3427–3434, 1997.

7 James Cox and Chee K. Yap. On-line motion planning: case of a planar rod. *Annals of Mathematics and Artificial Intelligence*, 3:1–20, 1991. Special journal issue. Also: NYU-Courant Institute, Robotics Lab., No.187, 1988.

8 Jory Denny, Kensen Shi, and Nancy M. Amato. Lazy Toggle PRM: a Single Query approach to motion planning. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 2407–2414, 2013. Karlsrube, Germany. May 2013.

9 David Eberly. Distance to circles in 3D, May 31 2015. Downloaded from `https://www.geometrictools.com/Documentation/Documentation.html`.

10 Mohab Safey el Din and Eric Schost. A baby steps/giant steps probabilistic algorithm for computing roadmaps in smooth bounded real hypersurface. *Discrete and Comp. Geom.*, 45(1):181–220, 2011.

11 Hazel Everett, Christian Gillot, Daniel Lazard, Sylvain Lazard, and Marc Pouget. The Voronoi diagram of three arbitrary lines in $\mathbb{R}^3$. In *25th European Workshop on Computational Geometry (EuroCG'09)*, 2009. March 2009, Bruxelles, Belgium.

12 Hazel Everett, Daniel Lazard, Sylvain Lazard, and Mohab Safey el Din. The Voronoi diagram of three lines. *Discrete and Comp. Geom.*, 42(1):94–130, 2009. See also 23rd SoCG, 2007. pp.255–264.

13 Dan Halperin, Efi Fogel, and Ron Wein. *CGAL Arrangements and Their Applications*. Springer-Verlag, Berlin and Heidelberg, 2012.

14 Dan Halperin, Oren Salzman, and Micha Sharir. Algorithmic motion planning. In Jacob E. Goodman, Joseph O'Rourke, and Csaba Toth, editors, *Handbook of Discrete and Computational Geometry*, chapter 50. Chapman & Hall/CRC, Boca Raton, FL, 3rd edition, 2017. Expanded from second edition.

15 Michael Hemmer, Ophir Setter, and Dan Halperin. Constructing the exact Voronoi diagram of arbitrary lines in three-dimensional space. In *Algorithms – ESA 2010*, volume 6346 of *Lecture Notes in Computer Science*, pages 398–409. Springer Berlin / Heidelberg, 2010.

16 Ching-Hsiang Hsu, Yi-Jen Chiang, and Chee Yap. Rods and rings: Soft subdivision planner for **R**^3 x **S**^2, 2019. Hosted on arXiv as `arXiv:1903.09416 [cs.CG]`. Also available at `http://cse.poly.edu/chiang/rod-ring18.pdf`.

17 V. Koltun. Pianos are not flat: rigid motion planning in three dimensions. In *Proc. 16th ACM-SIAM Sympos. Discrete Algorithms*, pages 505–514, 2005.

18 Steven M. LaValle. *Planning Algorithms*. Cambridge University Press, Cambridge, 2006.

19 Ji Yeong Lee and Howie Choset. Sensor-based planning for a rod-shaped robot in 3 dimensions: Piecewise retracts of $R^3 \times S^2$. *Int'l. J. Robotics Research*, 24(5):343–383, 2005.

20 Zhongdi Luo, Yi-Jen Chiang, Jyh-Ming Lien, and Chee Yap. Resolution exact algorithms for link robots. In *Proc. 11th Intl. Workshop on Algorithmic Foundations of Robotics (WAFR '14)*, volume 107 of *Springer Tracts in Advanced Robotics (STAR)*, pages 353–370, 2015. 3-5 Aug 2014, Boğazici University, Istanbul, Turkey.

21 James R. Munkres. *Topology*. Prentice-Hall, Inc, second edition, 2000.

22 Michal Nowakiewicz. MST-Based method for 6DOF rigid body motion planning in narrow passages. In *Proc. IEEE/RSJ International Conf. on Intelligent Robots and Systems*, pages 5380–5385, 2010. Oct 18–22, 2010. Taipei, Taiwan.

23 Colm Ó'Dúnlaing, Micha Sharir, and Chee K. Yap. Retraction: a new approach to motion-planning. *ACM Symp. Theory of Comput.*, 15:207–220, 1983.

24 Colm Ó'Dúnlaing and Chee K. Yap. A "Retraction" method for planning the motion of a disc. *J. Algorithms*, 6:104–111, 1985. Also, Chapter 6 in *Planning, Geometry, and Complexity*, eds. Schwartz, Sharir and Hopcroft, Ablex Pub. Corp., Norwood, NJ. 1987.

25 J. T. Schwartz and M. Sharir. On the piano movers' problem: I. The case of a two-dimensional rigid polygonal body moving amidst polygonal barriers. *Communications on Pure and Applied Mathematics*, 36:345–398, 1983.

26 Jacob T. Schwartz and Micha Sharir. On the piano movers' problem: II. General techniques for computing topological properties of real algebraic manifolds. *Advances in Appl. Math.*, 4:298–351, 1983.

27 Jacob T. Schwartz and Micha Sharir. On the piano movers' problem: V. the case of a rod moving in three-dimensional space amidst polyhedral obstacles. *Comm. Pure and Applied Math.*, 37(6):815–848, 1984. doi:10.1002/cpa.3160370605.

28 M. Sharir, C. O'D'únlaing, and C. Yap. Generalized Voronoi diagrams for moving a ladder I: topological analysis. *Communications in Pure and Applied Math.*, XXXIX:423–483, 1986. Also: NYU-Courant Institute, Robotics Lab., No. 32, Oct 1984.

29 M. Sharir, C. O'D'únlaing, and C. Yap. Generalized Voronoi diagrams for moving a ladder II: efficient computation of the diagram. *Algorithmica*, 2:27–59, 1987. Also: NYU-Courant Institute, Robotics Lab., No. 33, Oct 1984.

30 Vikram Sharma and Chee K. Yap. Robust geometric computation. In Jacob E. Goodman, Joseph O'Rourke, and Csaba Tóth, editors, *Handbook of Discrete and Computational Geometry*, chapter 45, pages 1189–1224. Chapman & Hall/CRC, Boca Raton, FL, 3rd edition, 2017. Revised and expanded from 2004 version.

31 I.A. Şucan, M. Moll, and L.E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 19(4):72–82, 2012. doi:10.1109/MRA.2012.2205651.

32 Cong Wang, Yi-Jen Chiang, and Chee Yap. On soft predicates in subdivision motion planning. *Comput. Geometry: Theory and Appl. (Special Issue for SoCG'13)*, 48(8):589–605, September 2015.

33 Chee Yap, Zhongdi Luo, and Ching-Hsiang Hsu. Resolution-exact planner for thick non-crossing 2-link robots. In *Proc. 12th Intl. Workshop on Algorithmic Foundations of Robotics (WAFR '16)*, 2016. 13-16 Dec 2016, San Francisco. The appendix in the full paper (and arXiv from `http://cs.nyu.edu/exact/` (and `arXiv:1704.05123 [cs.CG]`) contains proofs and additional experimental data.

34 Chee Yap, Vikram Sharma, and Jyh-Ming Lien. Towards exact numerical Voronoi diagrams. In *9th Int'l Symp. of Voronoi Diagrams in Science and Engineering (ISVD).*, pages 2–16. IEEE, 2012. Invited Talk. June 27-29, 2012, Rutgers University, NJ. doi:10.1109/ISVD.2012.31.

35 Chee K. Yap. Algorithmic motion planning. In J.T. Schwartz and C.K. Yap, editors, *Advances in Robotics, Vol. 1: Algorithmic and geometric issues*, volume 1, pages 95–143. Lawrence Erlbaum Associates, 1987.

36 Chee K. Yap. Soft subdivision search in motion planning. In A. Aladren et al., editor, *Proceedings, 1st Workshop on Robotics Challenge and Vision (RCV 2013)*, 2013. A Computing Community Consortium (CCC) Best Paper Award, Robotics Science and Systems Conference (RSS 2013), Berlin. In arXiv:1402.3213.

37 Chee K. Yap. Soft subdivision search and motion planning, II: Axiomatics. In *Frontiers in Algorithmics*, volume 9130 of *Lecture Notes in Comp.Sci.*, pages 7–22. Springer, 2015. Plenary Talk at 9th FAW. Guilin, China. Aug 3-5, 2015.

38 Liangjun Zhang, Young J. Kim, and Dinesh Manocha. Efficient cell labeling and path non-existence computation using C-obstacle query. *Int'l. J. Robotics Research*, 27(11–12):1246–1257, 2008.

39 Bo Zhou, Yi-Jen Chiang, and Chee Yap. Soft subdivision motion planning for complex planar robots. In *Proc. 26th European Symp. Algo.(ESA)*, pages 73:1–73:14, 2018. Helsinki, Finland, Aug 20-24, 2018.

40 D.J. Zhu and J.-C. Latombe. New heuristic algorithms for efficient hierarchical path planning. *IEEE Transactions on Robotics and Automation*, 7:9–20, 1991.