

Chapter 1

Theory of Soft Subdivision Search and Motion Planning

Chee K. Yap*
Department of Computer Science
Courant Institute of Mathematical Sciences
New York University
New York, NY 10012, USA
Email: yap@cs.nyu.edu

Abstract

We propose to design motion planning algorithms using two ingredients: the subdivision paradigm coupled with **soft predicates**. Such predicates are conservative and convergent relative to traditional exact predicates (called “hard” in this context). This leads to **resolution-exact** algorithms which can be viewed as a strong form of “resolution complete” algorithms. Resolution-exactness contains inherent indeterminacies and other subtleties. We describe an algorithmic framework, called **Soft Subdivision Search** (SSS) for designing such algorithms. There are many parallels between our framework and the well-known Probabilistic Road Maps (PRM) framework. Both frameworks lead to algorithms that are highly practical, easy to implement, have adaptive and local complexity. The critical difference is that SSS avoids the Halting Problem of PRM.

In a previous paper, we have demonstrated the ease of designing soft predicates for various motion planning problems. In this paper, we generalize and extend some of these results. We show how exact algorithms can be recovered by an extension of our framework. The SSS framework provides a theoretically sound basis for new classes of algorithms in motion planning and beyond. Such algorithms are novel, even in the exact case.

1.1 Introduction

Motion Planning is a fundamental problem in robotics. One of its origins is the “findpath problem” in Artificial Intelligence [6, 5]. In the 1980s, computational geometers began the algorithmic study of motion planning [31, 13], focusing on *exact planners*: assuming the input is exact, such planners return a path if any exists, and report “No Path” otherwise. Schwartz and Sharir [30] observed that the cell decomposition approach is a universal approach for motion planning, and in the algebraic case, is effectively reducible to Collins’s cylindrical algebraic decomposition. We introduced the concept of retraction motion planning in [27, 26]. In the first survey

* This work is supported by NSF Grant CCF-0917093.

on algorithmic robotics [31], we observed that the retraction approach is also universal (again, this is effective in the algebraic case). After the work of Canny [7], the retraction approach became popularly known as the “roadmap approach”. In the 1990’s the roadmap approach takes another turn.

¶1. **Theory.** Today, exact motion planning continues to be actively investigated (e.g., [12]). A fairly up-to-date account from the perspective of real algebraic geometry may be found in [3]. Some of these algorithms represent major theoretical advancements. Nevertheless their impact on practical robotics is quite modest: thus [36] noted that exact implementations have been limited to 3 degrees of freedom, and for simple robots only. Various sub-algorithms and supporting data structures needed in exact motion planning have been implemented in CGAL [8]. For example, the recent exact algorithm for the Voronoi diagram of lines in space is regarded as a significant advance [15]; but true goal here is Voronoi diagram of polyhedral objects. Exactness has tremendous cost in terms of computational complexity: it implicitly requires algebraic numbers. Direct manipulation of algebraic numbers is impractical. But for many basic problems, a weaker form under the paradigm of Exact Geometric Computation (EGC) is sufficient [32]. Nevertheless, the usual expedient is to replace exact arithmetic by machine arithmetic, leading to the ubiquitous problems of numerical non-robustness [33]. Even ignoring efficiency issues, there is a fundamental but less well-known barrier: *the Turing computability of exact algorithms for most non-algebraic problems is unknown* [34]. This barrier exists in most problems beyond kinematic motion planning. But see [9] for a rare case of a non-algebraic motion planning problem that is provably computable; this positive result is possible thanks to deep results in transcendental number theory.

¶2. **Practice.** Since the mid 1990’s, the method of Probabilistic Road Maps (PRM) has become dominant among roboticists. Its basic formulation comes from Kavraki, Švestka, Latombe and Overmars [19]. PRM is not a particular algorithm but a **algorithmic framework** for motion planners. Many variants of this framework are known: Expansive-Space Tree planner (EST), Rapidly-exploring Random Tree planner (RRT), Sampling-Based Roadmap of Trees planner (SRT), and many more. Quoting Choset et al [11, p.201]: “*PRM, EST, RRT, SRT, and their variants have changed the way path planning is performed for high-dimensional robots. They have also paved the way for the development of planners for problems beyond basic path planning.*”

In his invited talk at the recent workshop² on open problems in this field, J.C. Latombe stated that the major open problem of PRM is that it does not know how to terminate when there is no path. In practice, one simply times-out the algorithm, but this leads to problems such as the “Climbers Dilemma” [14, p. 4] described in the work of Bretl (2005). We call this the **Halting Problem of PRM**. This is a known issue for researchers, and is the extreme form of the so-called “Narrow Passage Problem” [11, p. 201]. Latombe’s talk suggested promising approaches

² IROS 2011 Workshop on Progress and Open Problems in Motion Planning, September 30, 2011, San Francisco.

such as Lazy PRM [4]; other lines of attack include explicit detection of the non-existence of paths [2]. The theoretical basis for PRM algorithms is that they are probabilistic complete [18]. The Halting Problem is inherent in probabilistic completeness.

¶3. **Common Ground.** We seek a common ground that provides stronger guarantees than probabilistic completeness, but avoids the inordinate demands of exactness. Fortunately for our subject, exactness is a mismatch for the needs of robotics. This is clear from the remark that physical devices and sensors have limited accuracy. Practitioners are acutely aware of this. Yet it does not absolve us from mathematical precision if we wish the theoretical development of robotic algorithms to thrive. This tension between the needs of practice and of theory has led to their divergent paths described above. So we turn to the idea of “resolution complete” algorithms, noting that the 1983 paper of Brooks and Lozano-Perez [6] was already on this track. It is known that resolution complete algorithms can avoid the Halting Problem (e.g., [36]). Unfortunately the notion of resolution completeness is seldom scrutinized, and is capable of many interpretations. In [10] we pointed out some untenable, or lacking, interpretations. As remedy, we introduced a version called **resolution-exactness**, and proved basic properties of such algorithms. Surprisingly, we show that resolution-exactness has an inherent indeterminacy, even for deterministic algorithms using exact predicates. But the indeterminacy is mild in comparison to that of probabilistic completeness. Unlike the determinacy of exact algorithms, this indeterminacy seems a perfect match for the requirements of robotics.

There are two ingredients of resolution-exact algorithms. The first is subdivision of configuration space. We organize the subdivision into a **subdivision tree**. In 2 and 3 dimensions, such trees are usually called quadtrees and octrees. Tree nodes correspond to subsets of configuration space with simple shapes such as boxes or simplices. The notion of grid search is often identified with resolution complete algorithms (e.g., see the Wikipedia entry on Motion Planning). Although grids are superficially similar to subdivision, we stress that typical grid-based methods are inadequate for resolution-exactness. The second ingredient is a **classification predicate** to decide if a node is free or not. Such predicates could be computed exactly in the algebraic case; that would be the reflex viewpoint of a computational geometer, but it is not where we want to be. Our key insight is this: *in the presence of subdivision, exact predicates can be replaced by suitable approximations*. We came by this viewpoint through a series of related work on subdivision algorithms (e.g., [28, 23, 29]). Such approximations are formalized as **soft predicates** in [10]. There we show through a series of motion planning examples, the relative ease of designing soft predicates, and claimed that they are practical.

Let us address this claim. Since the implementation of our algorithms is currently underway, our evidence for practicality is indirect: first, our subdivision infrastructure is based on well-understood and practical data structures (subdivision tree, union-find, etc). Next, the soft predicates we designed to go with subdivision are mostly reduced to estimating distances between two features, where a feature is a point, line segment or a triangle in space. Moreover, these predicates can be easily

and correctly implemented (see ¶12 below). Thus there are no implementation gaps for our algorithms. The argument so far centers on *implementability*. But how can we be sure that these implementations will be *practically efficient*? Here, we invoke the evidence of prior resolution-based work. We mention key papers such as Zhu and Latombe (1991) [37], Barhehenn and Hutchinson (1995) [1], and Zhang, Kim and Manocha (2008) [36]. Of course, since these work preceded our formulation, we must reinterpret their methods using our new perspective. In fact, it is illuminating and fruitful to revisit these papers from our current perspective. In short, through the implementability and practical efficiency of resolution-exact algorithms, we may have found a common ground for theory and practice.

¶4. **Our Goals.** The current paper aims to clearly expose the foundations of resolution exactness. There are three themes: (1) We first take a leaf from the success of PRM research: the simplicity and generality of PRM framework ensures that implementers of this framework can get easy access to a whole family of algorithms, just by modifying one or more components in the framework. This leads us to formulate an analogous framework for our approach, called **soft subdivision search** (SSS). (2) Next, we generalize the setting of our previous results [10]: for instance, the basic setting of a free space embedded in configuration space, $C_{free} \subseteq C_{space}$, can be replaced by an open subset Y of a normed linear space X . The boxes used in our subdivision trees can be replaced by other shapes such as simplices. (3) Finally, we want to revisit exact algorithms from a subdivision viewpoint: each SSS algorithm takes an input resolution parameter $\varepsilon > 0$, in addition to the normal inputs of path planning. It is essential that ε is positive. But if we admit $\varepsilon = 0$, the resolution-exact algorithm may become non-halting like PRM. We show how to fix this problem. Interestingly, such exact algorithms are novel and seems more implementable than usual exact algorithms.

All proofs are given in a separate Appendix. Cross references are hyperlinked in the pdf version of this paper.

¶5. **Preliminaries.** We establish some notations for standard concepts. To focus on the key ideas, this paper will assume the simplest formulation of the motion planning problem: point-to-point kinematic motion planning for any particular robot R_0 moving in a physical space \mathbb{R}^k ($k = 2, 3$) amidst a static obstacle $\Omega \subseteq \mathbb{R}^k$. The configuration space $C_{space} = C_{space}(R_0)$ will be appropriately embedded in \mathbb{R}^d ($d \geq k$) (see [21, p.128] for discussions of embedding issues). The **footprint map** is $Fp : C_{space} \rightarrow 2^{\mathbb{R}^k}$ where $Fp(\gamma) \subseteq \mathbb{R}^k$ is the physical space occupied by robot R_0 in configuration γ . E.g., for a rigid robot, $Fp(\gamma)$ is a rotated, translated copy of R_0 . Then a configuration $\gamma \in C_{space}$ is **free** iff $Fp(\gamma) \cap \Omega = \emptyset$. The set of free configurations $C_{free} = C_{free}(R_0, \Omega)$ is an open subset of C_{space} , assuming Ω is a closed set. But central to our theory is the boundary $\partial(C_{free})$ of C_{free} . Configurations in $\partial(C_{free})$ are said to be **semi-free**. A **motion** is a continuous function $\mu : [0, 1] \rightarrow C_{space}$, and its range $\mu[0, 1]$ is called the **trace**. The motion is **free** if its trace is contained in C_{free} . A **path** refers to a free motion. So the basic motion planning problem for robot R_0 is this: given start α and goal β configurations, and Ω (defining C_{free}), find a path from α to β if one exists, and report

“No Path” otherwise. A key tool is the **clearance function**, $C\ell : C_{space} \rightarrow \mathbb{R}_{\geq 0}$ where $C\ell(\gamma)$ is the separation of the footprint at γ from Ω , $C\ell(\gamma) := \text{sep}(Fp(\gamma), \Omega)$ where $\text{sep}(A, B) = \inf \{\|a - b\| : a \in A, b \in B\}$ is the **separation** between two sets $A, B \subseteq \mathbb{R}^k$. Thus $C\ell(\gamma) > 0$ iff $\gamma \in C_{free}$. The **clearance** of a motion μ is the minimum $C\ell(\mu(t))$ for $t \in [0, 1]$.

1.2 Two Frameworks for Motion Planning.

In this paper we use the terminology of **algorithmic framework** to discuss broad classes of algorithms, and view PRM as such a framework. An algorithm within the framework is³ just a specific instantiation, using particular data structures, strategies and subroutines. We will give a formulation of the PRM framework and our SSS framework, and compare them.

¶**6. The PRM Framework.** Here is a formulation of PRM, following LaValle [21, Section 5.4.1]: the goal is to find a path connecting $\alpha, \beta \in C_{space}$. We maintain a graph $G = (V, E)$ where $V \subseteq C_{free}$ and edges in E correspond to paths connecting the vertices of the edge. We may assume that α, β are in V . We need two predicates, $Free(u)$ to test if a configuration u is free, and $Connect(v, u)$ to test if the (straight) motion from v to u is free. Finally, assume some **termination criterion** that is comprised of two parts: success-criterion (found a path from α to β) and a failure-criterion (time-out or other condition).

PRM FRAMEWORK:
 While (termination criterion fails):
 1. Vertex Selection Method (VSM):
 Choose a vertex v in V for expansion.
 2. Configuration Generation Method (CGM):
 Generate some $u \in C_{space}$ (perhaps near v)
 3. Local Planning Method (LPM):
 If $Free(u)$,
 Add u to V
 If $Connect(v, u)$, add (v, u) to E .
 Return success or failure accordingly.

Step 1 (VSM) is usually controlled by some priority queue representing the search strategy. Step 2 (CGM) is the probabilistic step. But CGM could also be deterministic, e.g., controlled by a dense sampling sequence [20]. LaValle would call this the “Sampling Framework” to avoid any prior commitment to randomness. But we say “PRM Framework” for specificity, and in honor of the most well-known formulation of such approaches. In Step 3 (LPM), u is discarded if it is not free; another

³ To be sure, there are degrees of specificity. The most specific instantiation might be called “implementation” of some less specific “algorithm”.

method is to generate a free configuration u' such that the subpath, from v to u' , of the direct v to u path is free. As noted in [11, p.198], the practical success of PRM stems from the fact that the predicate $Free(u)$ is relatively cheap. There is a large literature on computing this predicate, under the heading of **collision detection**. Indeed, the $Connect(v, u)$ predicate is often reduced to $Free(u)$: if a “sufficiently dense” sampling of configurations from v to u is free, just assume there is a path from v to u .

By varying this simple framework, we could capture most of the known variations mentioned earlier. The original PRM is framed in terms of a road-map stage followed by a query stage; so the above version is closer to the “BasicPRM” of [18]. But our discussion of the “PRM Framework” is intended to cover such variations.

¶7. **What confers power to PRM?** The practical advantages of PRM is widely recognized, and it is natural to assume that randomness is the source of this power. LaValle et al [20] examine this question and concluded that sampling rather than randomness is the true source of power. Hsu et al [18] argue for the essential role of randomness. Independent of this debate, we offer another reason for the success of PRM: the PRM framework allows one to easily modify the constituent components (sampling strategy, connection strategy, freeness predicate, etc) to obtain a variety of algorithms that meet diverse needs. The basic infrastructure is kept relatively stable. This is possible thanks to the *simplicity* and *generality* of the PRM framework. Just as important in practice, the framework is also very *forgiving*: you could implement the constituent components approximately or even wrongly, and the software implementation⁴ may not necessarily fail (crash or loop). These properties are in sharp contrast to the usual exact algorithms which are far from simple and not too forgiving of errors [32]. In recognition of this, we would like to propose an analogous framework for the subdivision approach.

¶8. **The SSS Framework.** For a fixed robot R_0 , the motion planning input is an initial box $B_0 \subseteq C_{space}$, the obstacle $\Omega \subseteq \mathbb{R}^k$, the start and goal configurations $\alpha, \beta \in C_{space}$, and a **resolution parameter** $\varepsilon > 0$. We are interested in “resolution-restricted” search for a path from α to β *inside* B_0 . As noted in the introduction, our main data structure is a subdivision tree, \mathcal{T} . The root is B_0 and each tree node is a subbox of B_0 . The algorithm amounts to a while-loop that “grows” \mathcal{T} in each iteration by expanding some leaf until we find a path or conclude “No Path”. Here are the supporting subroutines and data structures: assume a **predicate** \tilde{C} that classifies each node in \mathcal{T} into FREE/STUCK/MIXED, with the property that $\tilde{C}(B) = \text{FREE}$ implies $B \subseteq C_{free}$ and $\tilde{C}(B) = \text{STUCK}$ implies $B \cap \overline{C_{free}} = \emptyset$. We maintain a priority queue $Q = Q_{\mathcal{T}}$ comprising those MIXED-leaves whose length $\ell(B)$ (defined below) is at least ε . Let $Q.\text{GetNext}()$ return a leaf B of highest priority. This B is given to another subroutine **Expand**(B) which subdivides B into two or more subboxes. These subboxes become the children of B (so B is no longer a leaf). For now, assume **Expand**(B) always split B into 2^d congruent subboxes. After splitting,

⁴ The hardware implementation, however, might have catastrophic consequences. But here, we rely on the fact that most robot systems are fail-safe.

\tilde{C} is immediately called to classify these subboxes. Boxes that are FREE need further processing: assume a union-find data structure D to maintain the connected components of the FREE leaves of \mathcal{T} . Say two boxes B, B' are **adjacent** if $B \cap B'$ is a $d-1$ dimensional set. This defines a graph whose vertices are the FREE boxes, and edges representing their adjacency relation. D maintains the connected components of this graph. We first insert each new FREE leaf B into D , and call $Union(B, B')$ for any FREE B' that is adjacent to B . Assume $Find(B)$ returns the connected component of B , and write “ $Box(\alpha)$ ” to denote any leaf of \mathcal{T} that contains $\alpha \in C_{space}$.

SSS FRAMEWORK

1. ▷ Initialization.
 While ($\tilde{C}(Box(\alpha)) \neq \text{FREE}$)
 If $Box(\alpha)$ has length $< \varepsilon$, Return (“No Path”)
 Else $\text{Expand}(Box(\alpha))$
 While ($\tilde{C}(Box(\beta)) \neq \text{FREE}$)
 ... do the same for β ...
2. ▷ Main Loop:
 While ($Find(Box(\alpha)) \neq Find(Box(\beta))$)
 If Q is empty, Return (“No Path”)
 $B \leftarrow Q.\text{GetNext}()$
 $\text{Expand}(B)$
3. Compute a FREE channel P from $Box(\alpha)$ to $Box(\beta)$
 Generate and return the “canonical path” \bar{P} inside P .

Resolution approaches can be wasteful when it is non-adaptive. In SSS, the resolution increases is naturally adaptive (we only expand at mixed cells). The resolution literature sometimes claimed incorrectly that the size of \mathcal{T} is (must be) exponential in the depth. A counter example is [29] where we prove that tree size is only polynomial in the depth for certain subdivision algorithms for root isolation. Our formulation can recapture the approach of Zhu and Latombe [37], Barbehenn and Hutchinson [1], or Zhang, Kim and Manocha (2008) [36] as follows: these papers expand along a “mixed channels” (i.e., path comprising FREE or MIXED boxes). We could define GetNext to expand similarly. It turns out (see [10]) that our computation of \tilde{C} could exploit the subdivision tree \mathcal{T} . LaValle observed this curious property of our method, calling it “opening up the blackbox” of collision testing.

¶9. Similarities. There are many similarities between PRM and SSS, especially in their contrasts with exact algorithms.

1. Both have two key subroutines, representing (i) the global search strategies and (ii) free-ness testing. In PRM, the two subroutines are the vertex selection method (VSM), and the $Free(u)$ predicate, respectively. In SSS, they are $\text{GetNext}()$, and the predicate $\tilde{C}(B)$.

2. An advantage of SSS and PRM is the possibility of finding paths *before* the entire C_{space} has been fully explored. Indeed, Hsu, Latombe and Kurniawati [18, p. 640] remarked that “foundational choice made in PRM planning is to avoid computing the exact shape of the free space”. Most exact methods require an expensive a pre-processing phase to compute a full description of free space.

3. Integrated path planning: both frameworks naturally compute a path, i.e., a parametrized curve in C_{free} . E.g., Step 3 of SSS converts a channel of free boxes into a path. But exact algorithms often focus on computing a symbolic path in some algebraic cell complex, assuming that some numerical subroutine will convert it into a path.
4. We have viewed PRM as a probabilistic framework, and SSS as a deterministic one. But both frameworks admit deterministic or probabilistic algorithms. In the future, we plan to explore the probabilistic side of SSS.

¶10. Differences.

- I. Foremost, SSS algorithms do not suffer from a Halting Problem.
- II. PRM needs the predicate $Connect(v, u)$ to connect two nodes. The analogue in SSS simply amounts to checking if two FREE boxes are adjacent.
- III. The search strategy in PRM resides in the Vertex Selection Method (VSM) and Configuration Generation Method (CGM). In SSS, it resides in `GetNext()` and `Expand()`. Sampling strategies is a major research question in PRM [20, 18]. Sampling in SSS seems to be more easily controlled, thanks to the nature of subdivisions. For example, a trivial randomized strategy in SSS is to pick any MIXED leaf with equal probability. Two deterministic SSS strategies are breadth-first search (BFS) and A-star/Dijkstra search [1]. We can have **hybrid strategies**: given two or more strategies, we just cycle through each one in turn. If one of them is randomized, then our hybrid will also gain any advantage of randomness.

1.3 Soft Classifiers

The $Connect(v, u)$ predicate is often implemented heuristically. LaValle [21, p. 177] discussed certified methods for this test based on Lipschitz constants. Such certified tests is a generalization of the $Free(u)$ predicate for a single configuration. We now consider a different generalization based on sets; it is basically the viewpoint of interval arithmetic [28].

We first generalize the setup in the Preliminary (¶5). Suppose X is a normed linear space with norm $\|\cdot\|$ (e.g., $X = \mathbb{R}^d$ and $\|\cdot\|$ is Euclidean norm). Fix a subset $\square X$ of the powerset 2^X (e.g., $\square X$ is the set of boxes in X). Call $\square X$ a **test domain** if it has these properties:

- Each $B \in \square X$ is a full-dimensional closed bounded polytope in X . We call B a **test cell** (or simply “cell”). We define an interior point $c(B)$ called its **center**.
- $\square X$ is closed under translation and dilation: if B is a cell, then so is $t + \sigma \cdot B$ for any $t \in X, \sigma > 0$. Here, $t + B$ denotes the translation of B by t , and $\sigma \cdot B$ denotes the dilation of B by ratio σ at the center $c(B)$.

Note that for $X = \mathbb{R}^d$, if $d = 1$, then cells are just closed intervals with distinct endpoints. For $d > 1$, we have many more possibilities.

By a **classifier** we mean any function $C : \square X \rightarrow \{\text{IN}, \text{ON}, \text{OUT}\}$. So a classifier is a special kind of predicate that “classifies” every test cell. These values⁵ correspond (respectively) to FREE/MIXED/STUCK of the previous section. Let $Y \subseteq X$ be any open subset of X . Call C a **Y -classifier** if for all $B \in \square X$,

$$\begin{cases} C(B) = \text{IN} & \implies B \subseteq Y \\ C(B) = \text{OUT} & \implies B \cap \overline{Y} = \emptyset \end{cases} \quad (1.1)$$

where \overline{Y} denote the closure of Y . Thus a trivial classifier is one that is identically ON, $C(B) \equiv \text{ON}$. If the two implications of (1.1) are replaced by logical equivalences (“if and only if” conditions) then we call C an **exact Y -classifier**, denoted by C_Y . Note that singletons $p \in X$ are not test cells, and so $C(p)$ is not defined. Nevertheless, the exact Y -classifier has a unique extension to points where $C_Y(p) = \text{IN}$ if $p \in Y$, $= \text{ON}$ if $p \in \partial Y$, and $= \text{OUT}$ otherwise. This extension is justified as follows: write “ $\lim_i B_i \rightarrow p$ ” to indicate an infinite decreasing sequence $B_1 \subseteq B_2 \subseteq \dots$ that converges to $p \in X$. It is easy to see that $\lim_i B_i \rightarrow p$ implies that the sequence $C_Y(B_1), C_Y(B_2), \dots$ eventually stabilizes to the value $C_Y(p)$. We denote this by writing “ $\lim_i C_Y(B_i) \rightarrow C_Y(p)$ ”. We can now define our key concept: a Y -classifier \tilde{C} is said to be **soft** if

$$\lim_i B_i \rightarrow p \implies \lim_i \tilde{C}(B_i) \rightarrow C_Y(p).$$

Thus, a soft predicate converges to the exact (or “hard”) predicate in the limit.

¶11. How to compute soft classifiers? Two standard ideas of resolution-based methods are (a) splitting cells to reduce complexity, and (b) using numerical approximation. Typically, (a) is determined by an arbitrary resolution parameter but [10] demonstrated the use of inherent adaptive splitting criteria. Here we focus on (b). Let $C\ell : X \rightarrow \mathbb{R}$ be a continuous function. Call $C\ell$ a (**generalized**) **clearance function** of the set $\{x \in X : C\ell(x) > 0\}$ (it is generalized because $C\ell$ can be negative). Because of the splits in (a), $C\ell$ need only be defined “locally”. Recall $\square\mathbb{R}$ is the set of closed intervals; consider an **interval function** $\square C\ell : \square X \rightarrow \square\mathbb{R}$. There is an classifier associated with $\square C\ell$, namely, $\tilde{C}(B) = \begin{cases} \text{IN} & \text{if } \square C\ell(B) > 0, \\ \text{OUT} & \text{if } \square C\ell(B) < 0, \\ \text{ON} & \text{else.} \end{cases}$

We call $\square C\ell$ a **box function** for $C\ell$ if it is **conservative** (i.e., $C\ell(B) \subseteq \square C\ell(B)$) and **convergent** (i.e., $\lim_i B_i \rightarrow p$ implies $\lim_i \square C\ell(B_i) \rightarrow C\ell(p)$). The following is straightforward.

Lemma 1. *Let $C\ell : X \rightarrow \mathbb{R}$ be a clearance function of a set Y . If $\square C\ell : \square X \rightarrow \square\mathbb{R}$ is a box function for $C\ell$, then its associated classifier is a soft Y -classifier.*

In practice, it is easier to design classifiers that focus only on the IN or the OUT decisions. So we call an interval function $\square C\ell : \square X \rightarrow \mathbb{R}$ a **positive box function** for Y if $B \subseteq Y$ implies $C\ell(B) \subseteq \square C\ell(B)$, and $\lim_i B_i \rightarrow p \in Y$ implies

⁵ These values may also be called EMPTY/MIXED/FULL, as in original Brooks-Perez paper. They reflect the 3-valued nature of geometric predicates (as opposed to 2-valued logical predicates).

$\lim_i \square C\ell(B_i) \rightarrow C\ell(p) > 0$. We similarly define **negative box function** for Y . See [10] (implicitly in [36]) for concrete examples of positive and negative classifiers.

¶12. Implementability. Correct implementation of algorithmic primitives is a central concern of EGC [8]. It remains central for SSS theory. An interesting remark is that *most papers on subdivision methods exploit cell resolution, but assume hard primitives being performed at the cell level*. Theorem 1 below illustrates this half-measure. Yet subdivision only shows its full power when we also exploit soft primitives. Of course, implementers will use approximate primitives (machine arithmetic implementation) but this is done with no guidance from theory.

We now indicate why the soft predicates which we designed in [10] are easily, efficiently and correctly implementable. To use numerical approximations, we need a dense subset \mathbb{D} of \mathbb{R} with good computational properties [34, ¶16]. A simple choice are the **dyadic numbers** $\mathbb{D} = \{m2^n : m, n \in \mathbb{Z}\}$, called BigFloats in software. To exploit hardware arithmetic, we use the technique of **(numerical) filters** in EGC [32]. Basically, filters perform machine arithmetic, but track error bounds to ensure safe decisions. The filter fails when any overflow or underflow is detected, at which point we switch to BigFloats. Using our `Core Library` [35], such filter techniques are automated so that users can write a “standard” C++ program to implement their predicates.

In a future implementation paper, we will give a careful account of the soft predicates designed in [10], but here is an overview of how to do filters using estimated error bounds. Such bounds suggest that our filters will rarely fail in the typical motion planning experiments. Assume the **Standard Model** of floating point arithmetic [16, p. 44] whereby, for any operation $x \circ y$ ($\circ \in \{+, -, \times, \div, \sqrt{\cdot}\}$) we have $x \tilde{\circ} y = (x \circ y)(1 \pm \mathbf{u})$ where $x \tilde{\circ} y$ denote the corresponding approximate arithmetic, \mathbf{u} is unit round-off error, and we use the notation “ $\tilde{x} = x(1 \pm \epsilon)$ ” to mean $\tilde{x} = x(1 + \theta)$ for some $\theta \in [-\epsilon, +\epsilon]$ (thus θ is an implicit constant). Note that $\mathbf{u} = 2^{-53} \simeq 1.11 \times 10^{-16}$ for IEEE double precision. The IEEE Standard for hardware arithmetic, and the widely available BigFloat package called MPFR [25] follows the Standard Model. For instance, to compute the distance $\|p - q\|$ between two dyadic points, assuming $p \neq q$, then in the Standard Model we can approximate $\|p - q\|$ with relative error of γ_3 where $\gamma_n := \frac{n\mathbf{u}}{1-n\mathbf{u}}$ (see [16] for this γ -analysis). We have $\gamma_n < 2n\mathbf{u}$ unless n is extremely large (e.g., $n > 2^{52}$ for IEEE double). Moreover, if $\ell(x, y) = ax + by + c = 0$ is the equation of a line, then its distance to an arbitrary dyadic point (x_0, y_0) is $|\ell(x_0, y_0)|/\sqrt{a^2 + b^2}$ and this has a relative error of γ_9 . This assumes that a, b, c are exact. But if the line is defined by two dyadic points (x_i, y_i) ($i = 1, 2$), then $a = (y_2 - y_1)$, $b = (x_2 - x_1)$, $c = y_1(x_2 - x_1) - x_1(y_2 - y_1)$. Our computation of the distance will now have a relative error of γ_{18} instead of γ_9 . The extension of such estimates to the case of rotation or in 3-D will increase n , but remains well under control. To obtain an upper or lower bound on a numerical expression such as $|\ell(x_0, y_0)|/\sqrt{a^2 + b^2}$, we just multiply its computed value by a factor of $(1 + \tilde{\gamma}_n)$ or $(1 - \tilde{\gamma}_n)$ where $\tilde{\gamma}_n$ is a machine upper bound on γ_n . If $n < 128$, say, then $\gamma_n < 2^{-45}$. Barring under or overflows, it means 45 bits of the mantissa are correct; this should suffice for typical applications.

1.4 Dyadic Subdivision Trees

Clearly subdivision trees are capable of many generalizations. So far, we assume that a node is a box in \mathbb{R}^d , and it splits into 2^d congruent children. We want to allow non-congruent shapes, and a variable number of children. One motivation is to exploit “anisotropic subdivisions”. E.g., in subdivision algorithms for isotopic approximation of curves and surfaces [23, 24, 22], we show empirically that “anisotropic subdivision” could lead to dramatic speedups.

We consider **generalized subdivision trees** whose nodes are cells from a test domain $\square X$ where $X = \mathbb{R}^d$. Two cells B, B' are **essentially disjoint** if $B \cap B'$ has dimension $\leq d - 1$. If $B \cap B'$ has dimension $d - 1$, we say the two cells are **adjacent** to each other. By a **subdivision** we mean a finite subset \mathcal{S} of $\square X$ such that any two cells are essentially disjoint. A **channel** is a sequence of cells where consecutive pairs are adjacent. We call \mathcal{S} a **subdivision of $|\mathcal{S}|$** where $|\mathcal{S}|$ denotes the union of the cells in \mathcal{S} . By a **k -split** (or split) of a cell $B \in \square X$ we mean a subdivision $\{B_1, \dots, B_k\}$ of B with $k \geq 2$ cells. We say the split is **dyadic** if each vertex of the B_i 's is either a vertex of B or the midpoint of an edge of B . A **dyadic subdivision tree** is a subdivision tree \mathcal{T} in which the children of each internal node forms a dyadic split of its parent. If \mathcal{T} is finite, then the set of leaves of \mathcal{T} forms a **dyadic subdivision** of the root. Dyadic subdivisions for boxes were exploited in [23, 24]. Why dyadic subdivision? In ¶12, we indicated the key role of dyadic numbers. Now each vertex v occurring in a dyadic subdivision tree is a linear combination of the vertices v_1, \dots, v_m of the root B_0 , $v = \sum_{i=1}^m \alpha_i v_i$ where each α_i is a dyadic number. We say v is **dyadic relative to B_0** . If B_0 is dyadic, then v is dyadic.

Our definition of a test cell B requires the concept of a center $c(B)$ in the interior of B . A candidate for $c(B)$ is the center of the **circumball**, i.e., unique smallest ball containing B . But this center may not lie in the interior of B . So we first define the **inner radius** $r_0(B)$ of B as the largest radius of a ball contained in B . Then the **incenter** $ic(B)$ comprises the centers of balls of radius $r_0(B)$ that are contained in B . E.g., the incenter of a non-square rectangle is a line segment. Clearly, $ic(B)$ is convex; the center of the circumball of $ic(B)$ is taken to be the **center** $c(B)$. Thus $c(B) \in ic(B)$, and is unique. The smallest ball centered at $c(B)$ and containing B is called the **outer ball** of B , and its radius $r(B)$ is called the **(outer) radius** of B . The **aspect ratio** is $r(B)/r_0(B)$. Let the **width** $w(B)$ (resp., **length** $\ell(B)$) refer to the minimum (resp., maximum) length of an edge of B .

We turn to **dyadic subdivision schemes**. The dyadic scheme for boxes is discussed in [23, 24, 10]. We briefly consider **dyadic simplicial schemes**. As illustrated in Figure 1.1(a), a triangle has three kinds of dyadic splits: the 4-, 3-, and 2-splits. Dyadic splits of a tetrahedron is more complicated — just three kinds are illustrated in Figure 1.1(b). See [23, 24] for a method to choose among different splits.




Fig. 1.1 Dyadic splits of (a) triangle and (b) tetrahedron

1.5 Basic Properties of SSS

We prove some general results about SSS Planners for the basic motion planning problem (¶5). An “SSS Planner” is an algorithm obtained by instantiating the various subroutines in the SSS Framework. But assume a generalization of the SSS Framework (¶8) whereby boxes are now cells in $\square X$. The 3 key subroutines are \tilde{C} , Expand, and GetNext. Our basic assumptions or **axioms** about them are:

- **(A0: Softness)** \tilde{C} is a soft classifier for C_{free} .
- **(A1: Bounded dyadic expansion)** Expand splits a cell dyadically into a bounded number of subcells, each with a bounded number of vertices, with the ratio $\ell(B)/w(B)$ bounded. Moreover, the splitting scheme is **perpetual** (it will never get stuck).
- **(A2: Clearance is Lipschitz)** There is a constant $L_0 > 0$ such that for all $\gamma, \gamma' \in C_{space}$, $|C\ell(\gamma) - C\ell(\gamma')| < L_0 \|\gamma - \gamma'\|$

We made no assumptions on GetNext here because the needed properties are embedded in the SSS framework, namely GetNext returns a MIXED-leaf with length $\ell(B) \geq \varepsilon$ as long as such leaves exist. Although our goal is soft classifiers (A0), as a proof strategy, we will initially assume \tilde{C} is the exact C_{free} -classifier. In this case, we say our planner is **exact**. Note that (A1) does not guarantee bounded aspect ratio, but it guarantees every infinite path converges. (A2) relates clearance to the norm on X . This axiom holds even with rotational degrees of freedom; see [10] for the case of C_{space} is $SE(2)$, and Appendix B.

Theorem 1. *Every SSS Planner halts.*

Thus halting is in-built, depending only on (A1). Next we give the minimal correctness property of SSS Planners.

Theorem 2 (Exact SSS). *Assuming an exact SSS Planner:*

- If there is no path, the planner outputs “No Path”.*
- If there is a path with clearance $\geq 2\varepsilon L_0$, the planner outputs a path.*

¶13. **Three Desiderata.** Loosely interpreted, current “resolution complete” algorithms provide perhaps the equivalent of Theorem 2. But there are three desiderata. First, we want to remove the assumption of exactness in \tilde{C} . As noted in ¶12, the literature invariably assumes exactness in its analysis. Second, we would like to weaken the hypothesis of Theorem 2(a) to “if there is no path with clearance ε/K ” for some input-independent $K > 1$. Third, we want to strengthen the conclusion of Theorem 2(b) so that the output path has clearance $\geq \varepsilon/K$.

¶14. **Soft Predicates and Effectivity.** The first desiderata above is to extend Theorem 2 to soft predicates. For such a result, we need a bit more of soft predicates. Call a Y -classifier \tilde{C} **effective** if it is⁶ **monotone** (i.e., $\tilde{C}(B) \neq \text{ON}$ implies $\tilde{C}(B') \neq \text{ON}$ for all $B' \subseteq B$), and there is a (effectivity) constant $\sigma > 1$ such that if $C_Y(B) = \text{IN}$ then $\tilde{C}(B/\sigma) = \text{IN}$. We remark that the explicit soft predicates which we designed in [10] are all effective. An SSS Planner is said to be **effective** if it uses an effective soft predicate for $Y = C_{free}$.

Theorem 3 (Effective SSS). *Assume an SSS Planner with effectivity $\sigma > 1$.*

(a) *If there is no path, the planner outputs “No Path”.*

(b) *If there is a path with clearance $\geq \varepsilon(1 + \sigma)L_0$, the planner outputs a path.*

¶15. **Resolution-Exactness.** The second and third desiderata lead our key definition: A planner for R_0 is said to be **resolution-exact** (or “ ε -exact”) if there exists a constant $K > 1$ such that

(i) *if there no path with clearance ε/K , it returns “No Path”;*

(ii) *if there is a path of clearance $K\varepsilon$, it returns a path with clearance ε/K .*

Call K an **accuracy constant**. Resolution-exact planners admits an indeterminacy in its output: suppose there is no path of clearance εK , but there exists one with clearance in the range $(\varepsilon/K, K\varepsilon)$. In this case, the ε -exact planner may return either a path or “No Path”. We show this indeterminacy is unavoidable in [10]. Another subtlety of this definition is revealed if we ignore the third desideratum. This amounts to replacing (ii) with the following:

(ii)’ *if there is a path of clearance $K\varepsilon$, it returns a path.*

Call the planner **weakly resolution-exact** in this case. What have we given up with this weaker requirement? We have no guarantees on the clearance of the returned μ . Nevertheless, because of (i), we know there exists a path μ' with clearance ε/K . So the third desideratum is a “constructivity requirement” (we must find such a μ'). The following development will show the highly nontrivial nature of this requirement.

To infer ε -exactness, the fundamental issue is to infer a lower bound on the clearance of a path inside a free channel. This is encoded as (A3) next:

- **(A3: Translational cells)** If $B \in \square X$ is free, then its center has clearance $C\ell(c(B)) \geq r_0(B)$ where $r_0(B)$ is the inner radius. Such cells are said to be **translational**.

⁶ Monotonicity is not strictly necessary, but it simplifies our arguments. Moreover, implementations can normally ensure monotonicity. In the interval literature, it is sometimes called “isotone”.

Like (A2), axiom (A3) relates the norm to clearance. It is a non-trivial assumption on the parametrization X of configuration space. The “translational” terminology is based on the analogy that if X is purely translational, then (A3) is true. The appendix will indicate why (A3) holds in standard motion planning formulations. So far, we have been non-specific about the “canonical path” \bar{P} that is generated in Step 3 of the SSS Framework. For ε -exactness, the nature of this canonical path is important. In particular, we must slightly modify Step 3 in SSS, encoded this in the next axiom:

- **(A4: Canonical Paths)** Assume that all cells are boxes and \tilde{C} is effective. In Step 3, if $\tilde{C}(\text{Cube}_\varepsilon(\alpha)) \neq \text{FREE}$ or $\tilde{C}(\text{Cube}_\varepsilon(\beta)) \neq \text{FREE}$, SSS will return “No Path”. (Otherwise it returns the “canonical path” \bar{P} in the channel P as usual.)

Here, $\text{Cube}_\varepsilon(\alpha)$ is the box centered at α with length and width equal to ε .

Theorem 4 (Resolution-Exact SSS).

Under (A0-A4), SSS Planners are resolution-exact.

1.6 What About Exact Algorithms?

Can the SSS framework produce⁷ exact algorithms? The answer is yes. But we first point out a non-solution, *using an Exact Planner with the resolution parameter $\varepsilon = 0$* . First of all, using Exact SSS re-introduces the need for algebraic computation. Second, by setting $\varepsilon = 0$, indeterminacy is removed, but at a high price: if there is no path, then SSS will not halt. Even if there is a path, we may not find it because of non-halting; but this could be fixed by imposing a “generalized BFS” property on `GetNext`. For these reasons, our normal⁸ formulation of SSS requires $\varepsilon > 0$.

We now present a solution that exploits resolution-exactness. It is based on the theory of constructive zero bounds [33], and does not need an Exact Planner.

Theorem 5. *If the inputs numbers describing $R_0, \Omega, \alpha, \beta$ are all algebraic numbers, there is an effectively computable number $\delta = \delta(R_0, \Omega, \alpha, \beta) > 0$ with this property: if there is a path from μ from α to β , then the clearance of μ is $> \delta$.*

One way to derive such a δ is to bound the degree and height of algebraic quantities arising in any motion planning algorithm. Then δ could be taken as the root separation for these algebraic quantities. A more careful computation of δ can proceed as follows: using the fact that if there is a path, then there is a path in some “retract” [31] (basically a Voronoi diagram augmented by some paths). This retract is algebraic, and the minimum clearance along the retract could be expressed by the solution of a suitable set of polynomial constraints involving the input data. A zero bound can be computed from a list of these constraints. These bounds depend on the representation used for input angular or rotational parameters.

⁷ We are indebted to Steve LaValle for asking this question at the IROS 2011 Workshop.

⁸ For that matter, we also assumed the accuracy constant K is strictly greater than 1.

Theorem 6. *Suppose we have a resolution-exact planner with accuracy parameter $K > 1$. If we fix the resolution parameter ε to be $\leq \delta(R_0, \Omega, \alpha, \beta)/K$, then the planner is exact.*

¶**16. Alternative approach to Exact SSS algorithms.** A more practical approach is to avoid zero bounds, and to minimize the role of algebraic computation. As in [10], we maintain a set of features $\phi(B)$ that are within an influence region of B , and another $\phi^-(B) \subseteq \phi(B)$. Our soft classifiers reduce to checking the emptiness of $\phi(B)$ or non-emptiness $\phi^-(B)$. But an exact predicate must ultimately compute the true value of $C_{C_{free}}(B)$. The idea is do this *only* when $\phi(B)$ is “simple”, otherwise we split B . Certainly, $|\phi(B)| = O(1)$ may be regarded as simple. Unfortunately, because of input degeneracies, this condition is not enough. Other options for simplicity are possible, but they depend on the nature of R_0 . This leads to new exact algorithms that seem more practical than traditional ones.

1.7 Conclusion

In this paper, we described the SSS framework for designing resolution-exact algorithms. We argued that it shares many of the attractive properties of the successful PRM framework. The ideas of resolution-limited algorithms is certainly very old. But to our knowledge, the simple⁹ properties of soft classifiers have never been isolated, nor have concepts of resolution-limited computation been carefully scrutinized. We believe focus on these “simple ideas” will open up new classes of algorithms that are practical *and* theoretically sound, not only in motion planning.

There are many open questions concerning SSS framework. Like PRM, many variations of SSS are possible. Perhaps the biggest theoretical challenge is the complexity analysis of adaptive subdivision [29]. Here are some other topics:

- Our SSS framework detects “No Path” by exhaustion. We could speed this up by looking for non-existence of M-paths [36], but it is a challenge to design efficient techniques (this is connected to issues in computational homology).
- The general study of dyadic subdivision schemes satisfying (A1) is of great interest. We also need to better understand subdivision schemes for $SE(3)$ (see [10] for $SE(2)$). Beyond kinematic spaces, good subdivision is even less understood.
- Design and analysis of adaptive search strategies, including randomized or hybrid ones. How efficiently can we update the “dynamic” A-star or Dijkstra search structures of [1]?
- Design and implement new SSS algorithms; compare them with PRM.
- An intriguing question is whether SSS match the performance of PRM in practice. Conventional wisdom says that PRM can provide practical solutions for

⁹ Several reviewers of our previous work sees only the safeness part of soft classifiers. They fail to note that previous work are silent about convergence. Of course, convergence is standard in numerical computing.

problems high¹⁰ degrees-of-freedom (DOF) while resolution methods can only reach medium DOF. This seems to be supported by current implementation. Choset [11, p. 202] suggests that the state-of-art PRM can handle DOF in the range 5 – 12. They noted that a 10 DOF planar robot from Kavraki (1995) cannot be tackled by other methods. The resolution-based algorithms of Zhang et al [36] involve planar robots (except for one 3D robot). But we believe the full potential of *adaptive* subdivision methods have hardly been reached. So this intriguing question begs for more experiments.

¶17. ACKNOWLEDGMENTS. I am indebted to Yi-Jen Chiang, Danny Halperin, Steve LaValle, and Vikram Sharma for helpful discussions.

References

1. M. Barbehenn and S. Hutchinson. Toward an exact incremental geometric robot motion planner. In *Proc. Intelligent Robots and Systems 95.*, volume 3, pp. 39–44, 1995. 1995 IEEE/RSJ Intl. Conf., 5–9, Aug 1995. Pittsburgh, PA, USA.
2. J. Basch, L. Guibas, D. Hsu, and A. Nguyen. Disconnection proofs for motion planning. In *IEEE Int’l Conf. on Robotics Animation*, pp.No Path1765–1772, 2001.
3. S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry*. Algorithms and Computation in Mathematics. Springer, 2003.
4. R. Bohlin and L. Kavraki. A randomized algorithm for robot path planning based on lazy evaluation. In P. Pardalos, S. Rajasekaran, and J. Rolim, editors, *Handbook on Randomized Computing*, pp. 221–249. Kluwer Academic Publishers, 2001.
5. M. Brady, J. Hollerbach, T. Johnson, T. Lozano-Perez, and M. Mason. *Robot Motion: Planning and Control*. MIT Press, 1982.
6. R. A. Brooks and T. Lozano-Perez. A subdivision algorithm in configuration space for findpath with rotation. In *Proc. 8th Intl. Joint Conf. on Artificial intelligence - Volume 2*, pp. 799–806, San Francisco, CA, USA, 1983. Morgan Kaufmann Publishers Inc.
7. J. F. Canny. *The complexity of robot motion planning*. ACM Doctoral Dissertation Award Series. The MIT Press, Cambridge, MA, 1988. PhD thesis, M.I.T.
8. CGAL, Computational Geometry Algorithms Library. <http://www.cgal.org>.
9. E.-C. Chang, S. W. Choi, D. Kwon, H. Park, and C. Yap. Shortest paths for disc obstacles is computable. *Int’l. J. Comput. Geometry and Appl.*, 16(5-6):567–590, 2006. Special Issue of IJCGA on Geometric Constraints. (Eds. X.S. Gao and D. Michelucci).
10. Y.-J. Chiang and C. Yap. Numerical subdivision methods in motion planning. Poster, IROS Workshop on Progress and Open Problems in Motion Planning. San Francisco, Sep 30, 2011.
11. 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.
12. M. S. el Din and É. Schost. A baby steps/giant steps probabilistic algorithm for computing roadmaps in smooth bounded real hypersurface. *DCG*, 45(1):181–220, 2011.
13. D. Halperin, L. Kavraki, and J.-C. Latombe. Robotics. In J. E. Goodman and J. O’Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 41, pp. 755–778. CRC Press LLC, 1997.
14. K. Hauser. *Motion planning for legged and humanoid robots*. PhD thesis, Stanford University, Dec 2008. Department of Computer Science.

¹⁰ Here, we do not consider specialized planning problems (molecules, snakes, humanoids, etc) where the DOF can go much higher than what is discussed here.

15. M. Hemmer, O. Setter, and D. 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*, pp. 398–409. Springer Berlin / Heidelberg, 2010.
16. N. J. Higham. *Accuracy and stability of numerical algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, second edition, 2002.
17. D. Hsu, J. Latombe, and R. Motwani. Path planning in expansive configuration spaces. In *Proc. of IEEE International Conference on Robotics and Automation*, pp. 2719–2726, 1997.
18. D. Hsu, J.-C. Latombe, and H. Kurniawati. On the probabilistic foundations of probabilistic roadmap planning. *Int'l. J. Robotics Research*, 25(7):627–643, 2006.
19. L. Kavraki, P. Švestka, C. Latombe, and M. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robotics and Automation*, 12(4):566–580, 1996.
20. S. LaValle, M. Branicky, and S. Lindemann. On the relationship between classical grid search and probabilistic roadmaps. *Int'l. J. Robotics Research*, 23(7/8):673–692, 2004.
21. S. M. LaValle. *Planning Algorithms*. Cambridge University Press, Cambridge, 2006.
22. L. Lin. *Adaptive Isotopic Approximation of Nonsingular Curves and Surfaces*. Ph.D. thesis, New York University, Sept. 2011.
23. L. Lin and C. Yap. Adaptive isotopic approximation of nonsingular curves: the parameterizability and nonlocal isotopy approach. *Discrete and Comp. Geom.*, 45(4):760–795, 2011.
24. L. Lin, C. Yap, and J. Yu. Non-local isotopic approximation of nonsingular surfaces, 2011. Submitted.
25. MPFR Homepage, Since 2000. URL <http://www.mpfr.org/>. MPFR is a C++-library for multi-precision floating-point computation with exact rounding modes.
26. C. Ó'Dúnlaing, M. Sharir, and C. K. Yap. Retraction: a new approach to motion-planning. *ACM Symp. Theory of Comput.*, 15:207–220, 1983.
27. C. Ó'Dúnlaing and C. 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.
28. S. Plantinga and G. Vegter. Isotopic approximation of implicit curves and surfaces. In *Proc. Eurographics Symp. on Geom. Processing*, pp. 245–254, New York, 2004. ACM Press.
29. M. Sagraloff and C. K. Yap. A simple but exact and efficient algorithm for complex root isolation. In I. Z. Emiris, editor, *36th ISSAC*, pp. 353–360, 2011. June 8–11, San Jose, CA.
30. J. T. Schwartz and M. 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.
31. C. K. Yap. Algorithmic motion planning. In J. Schwartz and C. Yap, editors, *Advances in Robotics, Vol. 1: Algorithmic and geometric issues*, pp. 95–143. Lawrence Erlbaum Associates, 1987.
32. C. K. Yap. Towards exact geometric computation. *Comput. Geometry: Theory and Appl.*, 7:3–23, 1997.
33. C. K. Yap. Robust geometric computation. In J. E. Goodman and J. O’Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 41, pp. 927–952. Chapman & Hall/CRC, Boca Raton, FL, 2nd edition, 2004.
34. C. K. Yap. In praise of numerical computation. In S. Albers, H. Alt, and S. Näher, editors, *Efficient Algorithms*, volume 5760 of *Lect. Notes in C.S.*, pp. 308–407. Springer-Verlag, 2009.
35. J. Yu, C. Yap, Z. Du, S. Pion, and H. Bronnimann. Core 2: A library for Exact Numeric Computation in Geometry and Algebra. In *3rd Proc. Int’l Congress on Mathematical Software (ICMS)*, pp. 121–141. Springer, 2010. LNCS No. 6327.
36. L. Zhang, Y. J. Kim, and D. Manocha. Efficient cell labelling and path non-existence computation using C-obstacle query. *Int'l. J. Robotics Research*, 27(11–12), 2008.
37. D. Zhu and J.-C. Latombe. New heuristic algorithms for efficient hierarchical path planning. *IEEE Transactions on Robotics and Automation*, 7:9–20, 1991.

1.8 APPENDIX A: Proofs

This appendix contain proofs for our theorems (for WAFR review process).

¶18. Halting and Splitting Criterion. In the SSS Framework (¶8), we use the criterion “ $\ell(B) < \varepsilon$ ” to stop splitting a cell. But in the following proofs, we will assume the variant criterion of “ $r(B) < \varepsilon$ ”. Why this difference? Generally, $\ell(B)$ is more easily computable than $r(B)$, so we expect to use $\ell(B)$ in implementations. However, the proofs are cleaner if we use $r(B)$. Note that $\ell(B)$ is just the distance between two vertices of B while $r(B)$ involves the center $c(B)$. For instance, if B is a dyadic box, then $\ell(B)$ is a dyadic number while $r(B)$ is a square-root.

Theorem 1. *Every SSS Planner halts.*

Proof. Property (A1) implies that in any infinite path $(B_i : i \geq 0)$ of an SSS subdivision tree \mathcal{T} , we have $\lim_i \ell(B_i) \rightarrow 0$. If we use the “ $\ell(B) < \varepsilon$ ” criterion for non-splitting, then this implies halting. But if we use the “ $r(B) < \varepsilon$ ” criterion, we also obtain $\lim_i r(B_i) \rightarrow 0$, because the bounded complexity of cells (A1) implies $r(B) = \Theta(\ell(B))$. **Q.E.D.**

Theorem 2 (Exact SSS).

Assuming an exact SSS Planner:

- (a) *If there is no path, the planner outputs “No Path”.*
- (b) *If there is a path with clearance $2\varepsilon L_0$, the planner outputs a path.*

Proof. Let \mathcal{T} be the subdivision tree at termination.

(a) At termination, we either report a path or output “No Path”. If we report a path, it is because we found a free channel from $B(\alpha)$ to $B(\beta)$, and this implies the existence of a path. Hence if there is no path, we will surely report “No Path”.

(b) Suppose $\mu : [0, 1] \rightarrow C_{space}$ is a path from α to β with clearance 2ε . By way of contradiction, suppose SSS outputs “No Path”. This implies that every mixed leaf satisfies $r(B) < \varepsilon$. Consider the set \mathcal{A} of leaves of \mathcal{T} that intersect $\mu[0, 1]$ (the trace of μ). If $B \in \mathcal{A}$, there exists $t \in [0, 1]$ such that $\mu(t) \in B$. This implies B is either free or mixed. We claim that B is free. If B is mixed, then $r(B) < \varepsilon$ and there is a point $p \in B$ that is semi-free. But $\|\mu(t) - p\| \leq \|\mu(t) - c(B)\| + \|c(B) - p\| < 2\varepsilon$. By (A2), $|C\ell(\mu(t)) - C\ell(p)| < 2\varepsilon L_0$. Thus $C\ell(p) > C\ell(\mu(t)) - 2\varepsilon L_0 \geq 0$, i.e., p is free. This contradicts the assumption that p is semi-free, proving our claim. Now we may form an channel of free cells from α to β using cells in \mathcal{A} . The existence of such a channel implies SSS should have reported a path. This contradicts our assumption of “No Path”. **Q.E.D.**

The application of (A2) in the above proof can be captured geometrically as follows: let $D_r(c)$ denote the **ball** in X centered at c with radius r . Then (A2) implies the following:

Lemma 2. *Let $c \in X$. If $C\ell(c) = L_0\delta > 0$ then $D_\delta(c)$ is free.*

¶19. **Soft Predicate and Effectivity.** **Theorem 3** (the soft version of **Theorem 2**) is an immediate corollary of the following:

Lemma 3.

Let \tilde{C} have effectivity constant $\sigma > 1$. If there is a path of clearance $(1 + \sigma)\varepsilon L_0$ then the SSS Planner will output a path.

Proof. The proof is similar to that of **Theorem 2(b)**. But now we use the values FREE/MIXED computed by \tilde{C} instead of the exact concepts of free/mixed. Suppose $\mu : [0, 1] \rightarrow C_{space}$ is a path of clearance $(1 + \sigma)\varepsilon L_0$ and, by way of contradiction, our algorithm outputs “No Path”. For any leaf B , if there is some $t \in [0, 1]$ such that $\mu(t) \in B$, then B is either FREE or MIXED (not STUCK). We claim B is FREE. By way of a second contradiction, assume B is MIXED. Thus $r(B) < \varepsilon$ and $\|c(B) - \mu(t)\| < \varepsilon$. By **Lemma 2**, we know $D_{(1+\sigma)\varepsilon}(\mu(t))$ is free. Then $\sigma B \subseteq D_{\sigma\varepsilon}(c(B)) \subseteq D_{(1+\sigma)\varepsilon}(\mu(t))$. Thus σB is free and hence $\tilde{C}(B) = \text{FREE}$. This contradiction proves our claim that B is FREE. Therefore the set of leaves that cover $\mu[0, 1]$ must be FREE. This contradicts the “No Path” output. **Q.E.D.**

¶20. **Resolution-Exactness.** We now prove our main theorem on ε -exactness. To provide quantitative information on the accuracy constant K , we introduce two global constants:

- There is a constant $K_0 > 1$ such that for $c \in X$, $\|c\| \leq K_0 \|c\|_\infty$ (the infinity norm). This implies that $\text{Cube}_\delta(c) \subseteq D_{K_0\delta}(c)$.
- There is a constant $K_1 > 1$ such that boxes in an SSS Planner have widths at least ε/K_1 . This constant exists because we do not subdivide a box with radius less than ε , and so $r(B) \geq \varepsilon/2$. Then axiom (A1) implies $w(B) \geq \varepsilon/K_1$ for some K_1 .

E.g., if $X = \mathbb{R}^d$, then $K_0 = \sqrt{d}$. If our boxes are restricted to cubes, then also $K_1 = \sqrt{d}$.

Theorem 4 (Resolution-Exactness).

Under (A0-A4), SSS Planners are resolution-exact.

The Planner has accuracy constant $K = \max\{4, 2K_1, (1 + \sigma)L_0K_0\}$ where σ is the effectivity constant of predicate \tilde{C} .

Recall that under (A4), we assume that our cells are actually boxes. We say a box B is **aligned (under B_0)** if B is contained in some subdivision of B_0 ; let $\text{Aligned}(B_0)$ denote the set of boxes aligned under B_0 . Clearly, the boxes appearing in an SSS Planner are all aligned under a fixed B_0 ; we can normally omit reference to this B_0 . We will prove a theorem about clearance of the “canonical path” in a free “aligned channel”:

Theorem 4A (Clearance in Free Aligned Channel). Let $P = (B_1, \dots, B_m)$ be a free channel of aligned boxes, with $\alpha \in B_1$ and $\beta \in B_m$. Assume

- (i) $\text{Cube}_\delta(\alpha)$ and $\text{Cube}_\delta(\beta)$ are free, and

(ii) the width of the B_i 's is at least ε .

Under (A3), the canonical path in P from α to β has clearance $\geq \min\{\varepsilon/2, \delta/4\}$.

The “canonical path” in SSS (Step 3) is precisely this path in Theorem 4A. Our main result on ε -exactness easily follows from Theorem 4A:

Proof of Theorem 4. We must show the two requirements of resolution-exactness:

(i) If SSS outputs a path, then this is the canonical path \overline{P} of a free channel P of aligned boxes. Each box in P has width $\geq \varepsilon/K_1$. Axiom (A4) implies $\text{Cube}_\varepsilon(\alpha)$ and $\text{Cube}_\varepsilon(\beta)$ are free. Thus Theorem 4A implies the clearance of \overline{P} is at least $\min\{\varepsilon/2K_1, \varepsilon/4\} \geq \varepsilon/K$. Therefore, if there is no path of clearance ε/K , the Planner must output “No Path”.

(ii) Suppose there is a path of clearance $(1 + \sigma)\varepsilon L_0 K_0$ where $\sigma \geq 1$ is the effectivity of \tilde{C} . The proof of Lemma 2 shows there is a FREE channel P of aligned boxes from $B(\alpha)$ to $B(\beta)$. We further claim that $\tilde{C}(\text{Cube}_\varepsilon(\alpha)) = \text{FREE}$: since the clearance of α is at least $(1 + \sigma)\varepsilon L_0 K_0$, Lemma 2 implies that $D_{(1+\sigma)\varepsilon K_0}(\alpha)$ is free. Thus $\text{Cube}_{(1+\sigma)\varepsilon}(\alpha)$, which is contained in $D_{(1+\sigma)\varepsilon K_0}(\alpha)$, is free. Since \tilde{C} has effectivity σ , $\tilde{C}(\text{Cube}_\varepsilon(\alpha)) = \text{FREE}$. Similarly, $\tilde{C}(\text{Cube}_\varepsilon(\beta)) = \text{FREE}$. Now we invoke Theorem 4A to conclude that \overline{P} has clearance $\geq \varepsilon/K$. **Q.E.D.**

We have not tried to optimize the accuracy constant K in our SSS Planner. A simple way to minimize K is to choose other criterion than “ $r(B) < \varepsilon$ ” to stop splitting a box.

¶21. Channels of Aligned Boxes. The proof of Theorem 4A is shown through a sequence of lemmas. The **canonical path** \overline{P} in the free channel $P = (B_1, \dots, B_m)$ in Theorem 4A is a fairly natural polygonal path. It is the concatenation of $m + 1$ subpaths, $\overline{P} = \mu_0; \mu_1; \dots; \mu_m$ where:

- Starting from α , we take a certain 2-step path μ_0 from α to $c(B_1)$ as described in Lemma 7(b) below.
- Then for each $i = 1, \dots, m - 1$, we continue with the i th subpath μ_i from $c(B_i)$ to $c(B_{i+1})$. Subpath μ_i is a certain canonical subpath joining the centers of the two boxes, via the center of $B_i \cap B_{i+1}$, as described in Lemma 8 below.
- Finally, we take a 2-step path μ_m from $c(B_m)$ to β . This is analogous to μ_0 .

We begin with a simple property of aligned intervals (i.e., the case where B_0 is an interval).

Lemma 4. *Let I_0 be an interval. Any two intervals $I, I' \in \text{Aligned}(I_0)$ are essentially disjoint or $I \subseteq I'$ or $I' \subseteq I$.*

Proof. We can construct a unique dyadic subdivision tree $\mathcal{T}_\infty(I_0)$ with no leaves. This infinite tree contains every aligned interval. If $I, I' \in \mathcal{T}_\infty(I_0)$ do not lie on a common path, then they are essentially disjoint. Otherwise, there is a containment relation between them. **Q.E.D.**

Let $D \subseteq \square(X)$ be any set of boxes. Its **width** $w(D)$ is the minimum of the widths of boxes in D . When $D = \{B_1, \dots, B_m\}$, we simply write $w(B_1, \dots, B_m)$ instead of “ $w(\{B_1, \dots, B_m\})$ ”.

By definition, “boxes” are full-dimensional. We now need to consider “boxes” that are less than full-dimensional. In the following, we shall assume the boxes B and B' are given by $B = \prod_{i=1}^d I_i$ and $B' = \prod_{i=1}^d I'_i$. Let $F = B \cap B'$. Assuming F is non-empty, we have $F = \prod_{i=1}^d J_i$ where $J_i = I_i \cap I'_i$. We call F a **degenerate box** if $w(J_i) = 0$ for any i . If B, B' are essentially disjoint, then F must be degenerate. Define the **width** of degenerate boxes as follows: If F is a point, we define $w(F) := 0$ and otherwise

$$w(F) := \min \{w(J_i) : w(J_i) > 0 \text{ and } i = 1, \dots, d\}.$$

So unless F is a point, we have $w(F) > 0$.

Lemma 5. *If $B, B' \in \text{Aligned}(B_0)$ and B, B' are adjacent, then $w(B \cap B') \geq w(B, B')$.*

Proof. Since B, B' are adjacent, there is a unique i such that $J_i = I_i \cap I'_i$ is degenerate. Wlog, say $w(J_1) = 0$. Then $w(J_i) > 0$ for $i = 2, \dots, d$. Since I_i, I'_i are both aligned relative to $I_i(B_0)$ (the projection of B_0 onto the i -th axis), Lemma 4 says that $I_i \subseteq I'_i$ or $I'_i \subseteq I_i$. Thus $w(J_i) = (I_i \cap I'_i) = w(I_i, I'_i)$. This proves $w(B \cap B') = \min \{w(I_i, I'_i) : i = 2, \dots, d\} \geq \min \{w(I_i, I'_i) : i = 1, \dots, d\} = w(B, B')$. **Q.E.D.**

Lemma 6. *If B, B' are two boxes with $c(B) \in B'$, then*

$$w(B \cap B') \geq \min \{w(B)/2, w(B')\}.$$

Proof. It is sufficient to prove this for the case where B, B' are intervals. Let I, I' be intervals with $c(I) \in I'$. Say $I' = [-w'/2, w'/2]$ and $I = [c-w/2, c+w/2]$. Wlog, $-w'/2 \leq c \leq 0$. CASE $c+w/2 \leq w'/2$: Then $I \cap I'$ contains $[c, c+w/2]$ of width $w/2$. This proves the lemma. CASE $c+w/2 > w'/2$: Then $-c-w/2 \leq -w'/2$, and so $c-w/2 \leq -w'/2$. This proves that $I' \subseteq I$ and so $w(I \cap I') = w(I') = w'$, again proving the lemma. **Q.E.D.**

Lemma 7 (Canonical subpath for 2 overlapping boxes).

Let B, B' be free boxes in \mathbb{R}^d .

- (i) *If $B \subseteq B'$ then the straightline path from $c(B)$ to $c(B')$ has clearance at least $w(B)/2$.*
- (ii) *If $c(B) \in B'$, then the “canonical” 2-step path from $c(B)$ to $c(B \cap B')$ and then to $c(B')$ has clearance at least $\min \{w(B)/4, w(B')/2\}$.*

Proof. The proof is illustrated by Figure 1.2(i) and (ii).

- (i) The clearance of $c(B)$ is at least $w(B)/2$, and the clearance of $c(B')$ is at least $w(B')/2$. If $\mu : [0, 1] \rightarrow C_{free}$ is the straightline path from $c(B)$ to $c(B')$, then $\mu(t)$




Fig. 1.2 Paths from $c(B)$ to $c(B')$: (i) $B \subseteq B'$, (ii) $c(B) \in B'$, (iii) B, B' adjacent.

has clearance at least $((1-t)w(B) + tw(B'))/2$. Hence the path has clearance at least $\min\{w(B), w(B')\}/2 = w(B)/2$.

(ii) We apply part(i) twice: the straight path from $c(B)$ to $c(B \cap B')$ has width at least $w(B \cap B')/2$. A similar argument applies to the path from $c(B \cap B')$ to $c(B')$. By Lemma 6, $w(B \cap B') \geq \min\{w(B)/2, w(B')\}$. **Q.E.D.**

Lemma 8 (Canonical subpath for 2 adjacent boxes).

Let B, B' be two free aligned boxes that are also adjacent. Then there is a 2-step “canonical path” from $c(B)$ to $c(B')$ with clearance at least $w(B, B')/2$.

Proof. This proof is illustrated in Figure 1.2(iii). Let $F = B \cap B'$. We first show that there is a 2-step path μ from $c(B)$ to $c(F)$ with clearance at least $w(B, B')$.

With our usual notation, suppose $F = B \cap B' = \{a_1\} \times J_2 \times \cdots \times J_d$ where a_1 is an endpoint of I_1 . Consider the box $B'' = I_1 \times J_2 \times \cdots \times J_d$. Note that F is a face of B'' . We construct μ by concatenating two straightline paths: $\mu = \mu_1; \mu_2$ where μ_1 is the path from $c(B)$ to $c(B'')$, and μ_2 is the path from $c(B'')$ to $c(F)$. By Lemma 7(i), the clearance of μ_1 is at least $w(B'')/2$. But $w(B'') = \min\{w(I_1), w(F)\}$ but $w(F) \geq w(B, B')$ by Lemma 5. Thus μ_1 has clearance at least $w(B, B')/2$.

Next consider the clearance of μ_2 : it is not hard to see that the clearance of $\mu_2(t)$ is at least half of

$$w^* := \min\{w(F), w(I_1), w(I'_1)\}.$$

But $w(F) \geq w(B, B')$ by Lemma 5. Hence $w^* \geq w(B, B')$.

We are almost done: by repeating the above argument, we also have a 2-step path from $c(F)$ to $c(B')$. Concatenating, we have a 4-step path. But our lemma claimed a 2-step path: this 4-step path is actually equivalent to a 2-step path because of the properties of alignment. First of all, note that F is actually a face of either B or B' . If F is a face of B' then the 2-step path from $c(B')$ to $c(F)$ is actually a straightline path. Moreover, this straightline path is a continuation of the second half of the 2-step path from $c(B)$ to $c(F)$. This is illustrated in Figure 1.2(iii). **Q.E.D.**

Proof of Theorem 4A. Let $P = (B_1, \dots, B_m)$ be a free channel of aligned boxes. The canonical path from $\alpha \in B_1$ to $\beta \in B_m$ is a concatenation of the “canonical subpaths” given by Lemma 7(ii) and Lemma 8. The clearance of the

subpaths from Lemma 7(ii) is at least $\min\{\varepsilon/2, \delta/4\}$. The clearance of the subpaths from Lemma 8 is at least $\varepsilon/2$. This concludes our proof.

¶22. Exact Subdivision Algorithms. We now address exactness in the subdivision context.

Theorem 5. *If the inputs numbers describing $R_0, \Omega, \alpha, \beta$ are all algebraic numbers, there is an effectively computable number $\delta = \delta(R_0, \Omega, \alpha, \beta) > 0$ with this property: if there is a motion from μ from α to β , then the clearance of μ is $> \delta$.*

The truth of this theorem is not in question. We omit the tedious details of calculating δ in this version. The sketch in the text indicates some ways for doing this.

Theorem 6. *Suppose we have a resolution-exact planner with accuracy parameter $K \geq 1$. If we fix the resolution parameter ε to be $\leq \delta(R_0, \Omega, \alpha, \beta)/K$, then the planner is exact.*

Proof. If there exists a path, then there exists a path of clearance $> \delta \geq K\varepsilon$. By the correctness of our resolution-exact planner, if there is a solution path with clearance $> K\varepsilon$, then our algorithm will return a path. Conversely, if there is no path, then there is no path of clearance $< \varepsilon/K$. By the correctness of our resolution-exact planner, if there is a solution path with clearance $> K\varepsilon$, then our algorithm will return “No Path”. **Q.E.D.**

1.9 APPENDIX B: Justification of (A2) and (A3)

The theorems in this paper are proved in APPENDIX A, assuming axioms (A0-A4). Of these axioms, the reasonableness (A2) and (A3) is perhaps the least obvious. The current appendix shows why they hold in typical settings. These two axioms show a tight connection between the Euclidean norm $\|\cdot\|$ on X to the clearance function $Cl : X \rightarrow \mathbb{R}_{\geq 0}$. But mediating between these two concepts is the Hausdorff distance $d_H(A, B)$ between two closed sets of physical space, \mathbb{R}^k .

¶23. On Parametrizations of C_{space} In this paper we assumed C_{space} is parametrized by (and identified with) the set X which is a subset of \mathbb{R}^d , $X \subseteq \mathbb{R}^d$. Naturally, we must properly define “adjacencies” of cells in X in order to ensure that the topology of C_{space} is captured. For instance, if $C_{space} = SE(2)$, then $X = \mathbb{R}^2 \times [0, 2\pi]$ where we identify 0 and 2π to have the proper adjacency. For simplicity, we will assume *our cells are boxes*.

Observe that if X has only translational degrees of freedom, then (A2) and (A3) are immediate. To handle rotational degrees of freedom, we exploit the fact that these are compact groups, and can only have bounded effect on clearance. More precisely, suppose X can be written as the product $X = X_T \times X_R$ where X_T and X_R represent the “translational” and “rotational” components of parameter space.

Here $X_T = \mathbb{R}^t$ is unbounded but $X_R \subseteq \mathbb{R}^r$ is compact (for some $t \geq 0, r \geq 0, t + r = d$). For instance, if X is a parametrization of the configuration of $k \geq 1$ unconstrained rigid bodies, then $X_T = \mathbb{R}^{3k}$ and $X_R = (SO(3))^k \subseteq \mathbb{R}^{3k}$. Or, if X parametrizes the configuration space of a humanoid robot, then $X^T = \mathbb{R}^3$ while X_R is a very high degree rotational space.

Let $\gamma \in X$ and B be a box in X . Then we may write $\gamma = (\gamma_T, \gamma_R)$ where $\gamma_T \in X_T, \gamma_R \in X_R$. Likewise, $B = B_T \times B_R$ where $B_T \subseteq X_T$ and $B_R \subseteq X_R$. We will consider two kinds of balls centered at $c \in \mathbb{R}^d$ with radius δ :

$$\begin{aligned} D_\delta(c) &:= \{x \in \mathbb{R}^d : \|x - c\| \leq \delta\} \quad (\text{Euclidean ball}) \\ \text{Cube}_\delta(c) &:= \{x \in \mathbb{R}^d : \|x - c\|_\infty \leq \delta\} \quad (\text{Cube}) \end{aligned}$$

Thus the cube is just the ball under the ∞ -norm, $\|x\|_\infty = \max\{|x_i| : i = 1, \dots, d\}$ where $x = (x_1, \dots, x_d)$. These functions generalize naturally to sets: if $S \subseteq \mathbb{R}^d$, then we have the **generalized footprint** $Fp(S) := \cup_{c \in S} Fp(c)$, **generalized ball** $D_\delta(S) = \cup_{c \in S} D_\delta(c)$ and **generalized cube** $\text{Cube}_\delta(S) := \cup_{c \in S} \text{Cube}_\delta(c)$. The generalized ball (resp. cube) is just the Minkowski sum of S with $D_\delta(\mathbf{0})$ (resp., $\text{Cube}_\delta(\mathbf{0})$) where $\mathbf{0}$ is the origin. The following properties hold for typical parametrizations:

- (Translational component) If $c = (c_T, c_R) \in X_T \times X_R$, then the generalized footprint

$$Fp(\text{Cube}_\delta(c_T) \times c_R) = \text{Cube}_\delta(Fp(c)). \quad (1.2)$$

The import of (1.2) is that the left hand side involves $\text{Cube}_\delta(c_T)$ while the right hand side involves $\text{Cube}_\delta(Fp(c))$: the former is in translational space $X_T = \mathbb{R}^t$, the latter in physical space \mathbb{R}^k .

- (Rotational component) There is a constant L_1 that depends only on the robot such that if $c_T = c'_T$, then

$$d_H(Fp(c), Fp(c')) \leq L_1 \|c_R - c'_R\| \quad (1.3)$$

where $d_H(A, B)$ is the Hausdorff distance between closed sets.

¶24. Justification of (A3).

Lemma 9. *Property (1.2) implies (A3): if B is a free box, then*

$$C\ell(c(B)) \geq w(B)/2.$$

Proof. By definition of clearance,

$$C\ell(\gamma) = \sup_{\delta} \{\delta : D_\delta(Fp(\gamma)) \cap \Omega = \emptyset\}. \quad (1.4)$$

Thus, for any $\delta \geq 0$,

$$D_\delta(Fp(\gamma)) \cap \Omega = \emptyset \implies C\ell(\gamma) \geq \delta. \quad (1.5)$$

Suppose B is free, $c = c(B)$ and $\delta = w(B)/2$. Then $\text{Cube}_\delta(c_T) \times c_R$, as a subset of B , is free. But (1.2) implies $Fp(\text{Cube}_\delta(c_T) \times c_R) = \text{Cube}_\delta(Fp(c))$. Thus

$$\text{Cube}_\delta(Fp(c)) \cap \Omega = \emptyset.$$

Since $D_\delta(Fp(c)) \subseteq \text{Cube}_\delta(Fp(c))$, we infer $C\ell(c) \geq \delta = w(B)/2$ from (1.5). We have verified (A3). **Q.E.D.**

Counter example to (A3). To show that (A3) is a non-trivial property, we describe a situation where it fails. Suppose the center of a disc robot $R_0 \subseteq \mathbb{R}^2$ is constrained to lie on the x -axis. So the configuration space is $X = \mathbb{R}$ and consider an interval $B = [-\delta, \delta]$ with center $c(B) = 0$. Let Ω consists of a single point on the y axis whose separation from $Fp(0)$ is $\delta/2$. Then $C\ell(c(B)) = \delta/2 < w(B)/2 = \delta$. So (A3) does not hold.

¶25. Justification of (A2) . We first note a connection between Hausdorff distance and clearance:

Lemma 10. *Let $c \in X$.*

(a) *If c is free, then $C\ell(c)$ is equal to*

$$C\ell_1(c) := \inf_{c''} d_H(Fp(c), Fp(c'')) \quad (1.6)$$

where c'' ranges over semi-free configurations in X .

(b) *Then $C\ell(c)$ is equal to*

$$C\ell_2(c) := \inf_{c''} d_H(Fp(c), Fp(c'')) \quad (1.7)$$

where c'' ranges over non-free configurations in X .

Proof. (a) Suppose c'' is semi-free and let $\delta = d_H(Fp(c), Fp(c''))$. Then $Fp(c'')$ is contained in $D_\delta(Fp(c))$. The semi-freeness of c'' implies $D_\delta(Fp(c)) \cap \Omega$ is non-empty. This proves $C\ell(c) \leq \delta$ (cf. (1.4)). This proves

$$C\ell(c) \leq C\ell_1(c). \quad (1.8)$$

On the other hand, it is easy to see that there exists some semi-free c'' such that $C\ell(c) = d_H(Fp(c), Fp(c''))$. This proves $C\ell(c) \geq C\ell_1(c)$, and hence $C\ell(c) = C\ell_1(c)$.

(b) If c is non-free, then $C\ell(c) = 0$ and we get $C\ell(c) = C\ell_2(c)$ immediately (choose $c'' = c$ in (1.7)). Suppose c is free. The argument for (1.8) also shows $C\ell(c) \leq C\ell_2(c)$. But $C\ell_2(c) \leq C\ell_1(c)$ and $C\ell_1(c) = C\ell(c)$. This proves $C\ell(c) = C\ell_2(c)$. **Q.E.D.**

Lemma 11. *Assume Properties (1.2) and (1.3). There exists $L_0 > 0$ such that for all $c, c' \in X$:*

(a) *(Hausdorff distance on footprint is Lipschitz)*

$$d_H(Fp(c), Fp(c')) \leq L_0 \|c - c'\|.$$

(b) (Axiom (A2))

$$|Cl(c) - Cl(c')| \leq L_0 \|c - c'\|.$$

Proof. Let $c, c' \in X$. Here we use the **constant K_0** introduced for the proof of Theorem 4: we have $\|c - c'\| \leq K_0 \|c - c'\|_\infty$.

(a) Initially assume $c_R = c'_R$ and $\|c - c'\|_\infty = \delta$. Then $c'_T \in \text{Cube}_\delta(c_T)$ and hence

$$Fp(c') \subseteq Fp(\text{Cube}_\delta(c_T) \times c_R) = \text{Cube}_\delta(Fp(c)) \subseteq D_{K_0\delta}(Fp(c)).$$

Similarly, we can show $Fp(c) \subseteq D_{K_0\delta}(Fp(c'))$. This proves

$$d_H(Fp(c), Fp(c')) \leq K_0\delta. \quad (1.9)$$

In general, we may have $c_R \neq c'_R$ and:

$$\begin{aligned} d_H(Fp(c), Fp(c')) &\leq d_H(Fp(c), Fp(c'_T, c_R)) + d_H(Fp(c'_T, c_R), Fp(c')) \quad (\text{triangular inequality for } d_H) \\ &\leq K_0 \|c_T - c'_T\|_\infty + L_1 \|c - c'\| \quad (\text{by (1.3) and (1.9)}) \\ &\leq (K_0 + L_1) \|c - c'\|. \end{aligned}$$

This proves (a) if we let $L_0 = K_0 + L_1$.

(b) Note that this part is trivial if both c and c' are non-free. Suppose c' is non-free. Then we may assume c is free.

$$\begin{aligned} Cl(c) - Cl(c') &= Cl(c) \\ &\leq d_H(Fp(c), Fp(c')) \quad (\text{by Lemma 10(b)}) \\ &\leq L_0 \|c - c'\| \quad (\text{by Part(a)}). \end{aligned}$$

Suppose c' is free. Then choose a semi-free c'' such that $Cl(c') = d_H(Fp(c'), Fp(c''))$. Then

$$\begin{aligned} Cl(c) - Cl(c') &\leq d_H(Fp(c), Fp(c'')) - Cl(c') \quad (\text{by Lemma 10(a)}) \\ &\leq d_H(Fp(c), Fp(c')) + d_H(Fp(c'), Fp(c'')) - Cl(c') \quad (\text{triangular inequality}) \\ &\leq d_H(Fp(c), Fp(c')) \\ &\leq L_0 \|c - c'\| \quad (\text{by Part(a)}). \end{aligned}$$

Since the roles of c and c' can be interchanged in the above argument, we have shown $|Cl(c) - Cl(c')| \leq L_0 \|c - c'\|$. **Q.E.D.**