Models and Motion Planning

Mark de Berg[*]   Matthew J. Katz[†]   Mark H. Overmars[*]
A. Frank van der Stappen[*]   Jules Vleugels[*]

Abstract

We study the complexity of the motion planning problem for a bounded-reach robot in the situation where the n obstacles in its workspace satisfy two of the realistic models proposed in the literature, namely unclutteredness and small simple-cover complexity. We show that the maximum complexity of the free space of a robot with f degrees of freedom in the plane is $\Theta(n^{f/2} + n)$ for uncluttered environments as well as environments with small simple-cover complexity. The maximum complexity of the free space of a robot moving in a three-dimensional uncluttered environment is $\Theta(n^{2f/3} + n)$. All these bounds fit nicely between the $\Theta(n)$ bound for the maximum free-space complexity for low-density environments and the $\Theta(n^f)$ bound for unrestricted environments. Surprisingly—because contrary to the situation in the plane—the maximum free-space complexity is $\Theta(n^f)$ for a three-dimensional environment with small simple-cover complexity.

1   Introduction

It is well known that the maximum complexity of the free space of a robot with f degrees of freedom moving in a scene consisting of n disjoint obstacles of constant complexity can be $\Omega(n^f)$. Consequently, exact motion-planning algorithms often have a worst-case running time of at least the same order of magnitude. This is probably one of the reasons that most of the exact algorithms were never implemented. One exception is Bañon's implementation [3] of the $O(n^5)$ algorithm of Schwartz and Sharir [14] for a ladder moving in a two-dimensional workspace, which performs surprisingly well, and much better than the worst-case theoretical analysis predicts. The reason is that the running time of the algorithm is sensitive to the actual complexity of the free space, and this is in practice far less than the $\Theta(n^f)$ worst-case bound.

These observations inspired research [1, 2, 4, 7–11, 13, 15, 16, 19–21] where geometric problems are studied under certain (hopefully realistic) assumptions on the input—in the case of motion planning: the environment in which the robot is moving. The goal of this line of research is to be able to predict better the practical performance of algorithms. For instance, van der Stappen et al. [16] studied the free-space complexity for a *bounded-reach robot* moving in environments consisting of *fat obstacles*. (A robot has bounded reach if it is not too large compared to the obstacles in its workspace; an obstacle is fat if it has no long and skinny

parts.) They showed that in this restricted type of environments the worst-case free-space complexity is only $\Theta(n)$. Van der Stappen [17, 18] also proved that in such environments naive and slightly adapted versions of Schwartz and Sharir's ladder algorithm run in $O(n^2)$ and $O(n \log n)$ time, respectively, which is more in line with the experimental results of Bañon. Van der Stappen and Overmars [19] used the linear free-space complexity result to obtain an efficient general approach to robot motion planning amidst fat obstacles. These results were extended to the more general setting of low-density environments by van der Stappen et al. [20].

De Berg et al. [5] brought together various of the *realistic input models* that were proposed in the literature, namely *fatness*, *low density*, *unclutteredness*, and *small simple-cover complexity*—see Section 2 for formal definitions of these models. They showed that these models form a strict hierarchy in the sense that fatness implies low density, which in turn implies unclutteredness, which implies small simple-cover complexity, and that no other implications exist between the models. A natural question that arises is whether the results of van der Stappen et al. [20] remain valid when, instead of a low-density scene, we assume a more general setting, like an uncluttered scene or a scene with small simple-cover complexity. In other words, does the complexity of the free space of a bounded-reach robot with f degrees of freedom moving in an uncluttered scene (alternatively, in a scene with small simple-cover complexity) remain $O(n)$?

The main result of this paper is a negative answer to this question. We prove that the maximum complexity of the free space of a bounded-reach robot moving in either an uncluttered scene or a scene with small simple-cover complexity is $\Theta(n^{f/2} + n)$ when the workspace is two-dimensional. These bounds fit nicely between the $\Theta(n)$ bound for low-density scenes and the $\Theta(n^f)$ bound for general scenes. For three-dimensional uncluttered scenes the bound becomes $\Theta(n^{2f/3} + n)$. Contrary to the planar case, small simple-cover complexity does *not* result in a reduced maximum free-space complexity for three-dimensional workspaces: the maximum complexity is $\Theta(n^f)$.

Our upper-bound proofs use the concept of guarding sets [6]. A guarding set for a collection of objects—in our case the obstacles in the robot's workspace—is, informally speaking, a set of points (sometimes referred to as guards) that approximates the spatial distribution of these objects. Guarding sets allow us to define a simplifying generalization [6] of unclutteredness that implies small simple-cover complexity in 3D and is even equivalent to it in the plane.

Section 2 recalls the input models that play a role in this paper and briefly reviews the relations between these models and the concept of guarding sets. Section 3 establishes an upper bound on the number of large objects intersecting a hypercube given the number of guards in its vicinity. In Sections 4 and 5 we use the relations and the bound to obtain tight bounds on the maximum complexity of the free space for motion planning in uncluttered environments and environments with small simple-cover complexity in 2D and 3D, respectively. Section 6 concludes the paper.

2   Input models

Before we briefly describe the input models that play a role in this paper we list a few important assumptions and definitions. The dimension of the (work)space is denoted as d. We shall be dealing a lot with squares, cubes, rectangles, and so on. These are always assumed to be axis-aligned. All geometric objects we consider are d-dimensional and open; in

particular, if we talk about a point lying in a square or cube, we mean that the point lies in the interior of the square or cube. Furthermore, all objects we consider are assumed to have constant complexity. More precisely, each object is a compact connected set in $\mathbb{R}^d$, bounded by a constant number of algebraic surface patches of constant maximum degree.

The *size* of a square (more generally, of a hypercube) is defined to be its edge length, and the size of an object is the size of a smallest enclosing hypercube for the object.

An *L-shape* is the geometric difference of a hypercube σ with a hypercube $\sigma' \subseteq \sigma$ of less than half its size and sharing a vertex with it. An L-shape can be covered by $2^d - 1$ hypercubes contained in it: for each vertex v of σ not shared with σ', take the hypercube of maximal size with v as a vertex and contained in $\sigma \setminus \sigma'$.

Although this paper concentrates on motion planning in uncluttered environments and environments with small simple-cover complexity, we also briefly describe the model of low density for the sake of reference. It is the weakest model for which the free space of a bounded-reach robot is known to have linear complexity [20]. We leave fatness [5, 18, 19] out of our discussion as it imposes stronger constraints on the environment while leading to the same bound as low density.

The model of low density was introduced by van der Stappen et al. [20] and refined by Schwarzkopf and Vleugels [15]. It forbids any ball B to be intersected by many objects whose minimal-enclosing-ball radius is at least as large as the radius of B. In the definition, $\rho_{\mathrm{meb}}(\mathcal{P})$ denotes the radius of the minimal enclosing ball of an object $\mathcal{P}$.

Definition 2.1 *Let $\mathcal{O}$ be a set of objects in $\mathbb{R}^d$. We say that $\mathcal{O}$ has λ-low-density if for any ball B, the number of objects $\mathcal{P}_i \in \mathcal{O}$ with $\rho_{\mathrm{meb}}(\mathcal{P}_i) \geqslant \mathrm{radius}(B)$ that intersect B is at most λ.*

We say that a scene has *low density* if it has λ-low-density for a small constant λ.

Unclutteredness was introduced by de Berg [4]. The model is defined as follows.

Definition 2.2 *Let $\mathcal{O}$ be a set of objects in $\mathbb{R}^d$. We say that $\mathcal{O}$ is κ-cluttered if any hypercube whose interior does not contain a vertex of one of the bounding boxes of the objects in $\mathcal{O}$ is intersected by at most κ objects in $\mathcal{O}$.*

We call a scene *uncluttered* if it is κ-cluttered for a small constant κ.

The following definition of simple-cover complexity is a slight adaptation of the original definition by Mitchell et al. [12], as proposed by de Berg et al. [5]. Given a scene $\mathcal{O}$, we call a ball δ-*simple* if it intersects at most δ objects in $\mathcal{O}$.

Definition 2.3 *Let $\mathcal{O}$ be a set of objects in $\mathbb{R}^d$, and let $\delta > 0$ be a parameter. A δ-simple cover for $\mathcal{O}$ is a collection of δ-simple balls whose union covers the bounding box of $\mathcal{O}$. We say that $\mathcal{O}$ has (s, δ)-simple-cover complexity if there is a δ-simple cover for $\mathcal{O}$ of cardinality sn.*

We will say that a scene has *small simple-cover complexity* if there are small constants s and δ such that it has (s, δ)-simple-cover complexity.

Guarding sets [6] against hypercubes[1] provide a generalization of unclutteredness that turns out useful in our proofs. A guarding set for a collection of objects is, loosely speaking, a set of points that approximates the distribution of the objects. More precisely, guarding sets are defined as follows.

---

[1] In the paper by De Berg et al. [6] guarding sets are defined against an arbitrary family of ranges. It is sufficient for our purposes to concentrate on hypercubic ranges.

Definition 2.4 *Let $\mathcal{O}$ be a set of objects in $\mathbb{R}^d$, and let κ be a positive integer. A set G of points is called a κ-guarding set for $\mathcal{O}$ (against hypercubes) if any hypercube not containing a point from G intersects at most κ objects from $\mathcal{O}$.*

We will often call the points in G *guards.*

We are particularly interested in scenes that admit a small κ-guarding set for some small constant κ, that is, a guarding set of size linear in the number of objects in $\mathcal{O}$. Scenes that admit a linear-size guarding set fit nicely in the existing model hierarchy: a low-density scene is also an uncluttered scene, which, is a scene that admits a linear-size guarding set, which is a scene with small simple-cover complexity [5,6]. In the plane admitting a linear-size guarding set is even equivalent to having small simple-cover complexity, but this is not the case in higher dimensions. A consequence of these hierarchical relations is that upper bounds for scenes with linear-size guarding sets immediately transfer to planar and 3D uncluttered scenes as well as to planar scenes with small simple-cover complexity. This conclusion will come to our help in Sections 4 and 5.

3   Guards and vicinities

Guards provide information on the distribution of the objects in an environment. Let us assume we are given a κ-guarding set for a collection of objects. A hypercube without any guards is, by definition, intersected by at most κ objects. Moreover, a hypercube with exactly g guards in its interior is intersected by $O(\kappa g)$ objects [6]. Theorem 3.4 below states another, more surprising, relation between the distribution of the objects and the distribution of the guards. Again we look at hypercubes, but this time we only look at objects that are at least as large as the hypercube, and not only consider the guards inside the hypercube but also the ones in its vicinity.

We define the *vicinity* of a hypercube σ to be the hypercube obtained by scaling σ with a factor of 5/3 with respect to its center. Thus, if we partition the vicinity of σ into 5^d equal-sized subhypercubes, then σ consists of the 3^d middle subhypercubes. The planar case is illustrated in Figure 1.
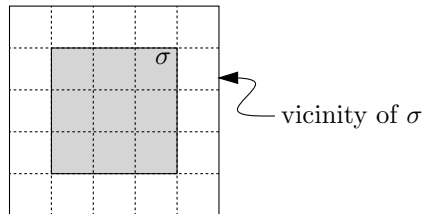

Figure 1: A square σ and its vicinity.

We will show that the number of objects intersecting a hypercube σ in $\mathbb{R}^d$ and at least as large as σ cannot be more than (roughly) $O(g^{1-1/d})$, where g is the number of guards in the vicinity of σ.

We first reduce the problem to a simpler problem on so-called 3-blocks. Define a 3-*block* to be the hyperrectangle obtained by scaling an axis-parallel hypercube by a factor of 1/3 along one of the coordinate axes. We say that an object *crosses* a given hyperrectangle if there exists a curve inside the intersection of the hyperrectangle and the object that connects

the two largest (and opposite) faces of the hyperrectangle. First we prove that if a hypercube is intersected by many larger objects, then there must be a 3-block in its vicinity that is crossed by many objects.

Lemma 3.1 *Let σ be a hypercube intersected by a collection $\mathcal{O}$ of m objects that are at least as large as σ. Then there is a 3-block contained in the vicinity of σ that is crossed by at least $m/(2d3^d)$ objects.*

Proof: Partition σ into 3^d equal-sized subhypercubes. One of the subhypercubes, say σ^*, intersects at least $m/3^d$ objects. Denote the set of objects intersecting σ^* by $\mathcal{O}^*$. Let Σ^* be the hypercube obtained by scaling σ^* with a factor three with respect to its center. Consider the $2d$ 3-blocks that are contained in Σ^* and have one of the $2d$ sides of Σ^* as a face—see Figure 2 for an illustration of the planar case. Note that Σ^* and, hence, all $2d$ 3-blocks, are contained in the vicinity of σ. We shall argue that one of the 3-blocks is crossed by at least $m/(2d3^d)$ objects.
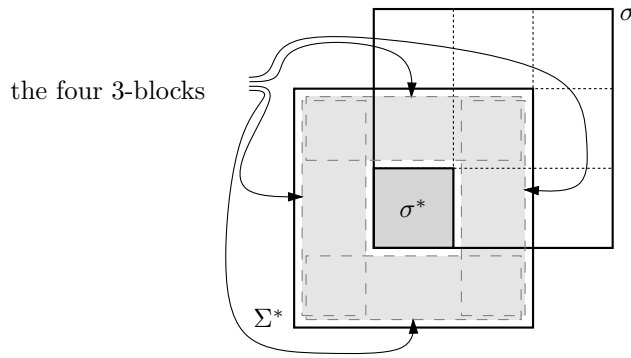

Figure 2: Illustration of the planar case in the proof of Lemma 3.1.

The hypercube Σ^* has the same size as the original hypercube σ. Hence, each object in $\mathcal{O}$ is at least as large as Σ^*. This implies that the objects cannot be fully contained in Σ^*. Thus, each object in $\mathcal{O}^*$ has a point inside σ^* and a point outside Σ^*. But this means that inside each such object we can find a curve connecting the two largest faces of one of the $2d$ 3-blocks. Hence, one of the 3-blocks must be crossed by at least $|\mathcal{O}^*|/2d = m/(2d3^d)$ objects. □

The next step is to prove a relation between the number of crossing objects and the number of guards in a 3-block B. The following auxiliary lemma will be called for in the proof of Lemma 3.3.

Lemma 3.2 *For any given hypercube σ intersected by a set $\mathcal{O}$ of m objects, and any constant b with $0 < b \leqslant m$, we can identify more than $\lfloor m/(2^{d+1} - 2)b \rfloor$ disjoint hypercubes inside σ each intersected by at least b objects from $\mathcal{O}$.*

Proof: We construct a tree on σ by recursively identifying between 2 and 2^d subhypercubes in the current hypercube if it is intersected by at least $2^d b$ objects. Assume that we have a hypercube σ^* intersected by at least $2^d b$ objects and consider its decomposition into 2^d equal-sized subhypercubes (by means of the d hyperplanes perpendicular to the coordinate

axes and cutting the hypercube into two equal halves). We call a (sub)hypercube crowded if it is intersected by at least b objects. Note that at least one of the subhypercubes of σ^* is crowded. Note also that each object intersecting σ^* intersects at least one of the 2^d subhypercubes (because the objects are d-dimensional and open).

- If the number of crowded subhypercubes of σ^* exceeds one, then each of these crowded subhypercubes becomes a child of σ^* in our tree. We charge the objects that do not intersect one of the crowded subhypercubes, to $\sigma*$. The number of such objects is at most $(2^d - 2)(b - 1)$.

- If the number of crowded subhypercubes of σ^* equals one, then this subhypercube σ' is shrunk towards the vertex it shares with σ^* until $\sigma^* \setminus \sigma'$ is intersected by at least $(2^d - 1)b$ objects but σ' is still intersected by at least b objects. (Note that this is always possible because the objects are d-dimensional.) Now consider the $2^d - 1$ hypercubes of maximal size contained in $\sigma^* \setminus \sigma'$ and sharing a vertex with σ^*. Since these hypercubes jointly cover $\sigma^* \setminus \sigma'$, at least one of them, say σ'', is crowded. The crowded hypercubes σ' and σ'' become the children of σ^* in our tree. We charge the at most $(2^d - 1)b - b = (2^d - 2)b$ objects that do not intersect σ' or σ'' to σ^*.

The leaves of the resulting tree correspond to disjoint hypercubes that are intersected by at least b objects. We will see that there are at least $m/(2^{d+1} - 2)b$ leaves.

Let L and I be the number of leaves and internal nodes of the resulting tree, respectively. As every internal node has at least two children the number of leaves is larger than the number of internal nodes, so $L > I$. We notice that the number of objects charged to an internal node is at most $(2^d - 2)b$ and the number of objects intersecting a hypercube corresponding to a leaf is at most $2^d b - 1$. Since the number of objects charged to all internal nodes plus the number of objects intersecting the hypercubes at the leaves should at least be equal to m we have that

$$I(2^d - 2)b + L(2^d b - 1) \geqslant m.$$

Using the inequality $L > I$ we obtain

$$L > \left\lfloor \frac{m}{(2^{d+1} - 2)b} \right\rfloor.$$

$\square$

Lemma 3.3 *Let G be a κ-guarding set for a collection $\mathcal{O}$ of objects in $\mathbb{R}^d$, with $d \geqslant 2$. Let B be a 3-block crossed by m objects from $\mathcal{O}$, and let $\kappa < m$. Then there must be at least*

$$(d-1)3^{d-1} \left\lfloor \frac{m}{d3^{d-1}(2^d - 2)(\kappa + 1)} \right\rfloor^{\frac{d}{d-1}}$$

guards from G inside B.

Proof: We assume withou loss of generality that the short side of B has unit length. We split B into

$$l := \left\lfloor \frac{m}{d3^{d-1}(2^d - 2)(\kappa + 1)} \right\rfloor^{\frac{1}{d-1}}$$

slices of height $1/l$ by means of hyperplanes parallel to its two largest faces, which are $(d-1)$-dimensional hypercubes of side length 3. The two largest (and opposite) faces of the resulting slices are again $(d-1)$-dimensional hypercubes of side length 3; see Figure 3 for a three-dimensional 3-block and the l slices. We observe that each of the objects from $\mathcal{O}$ crossing B
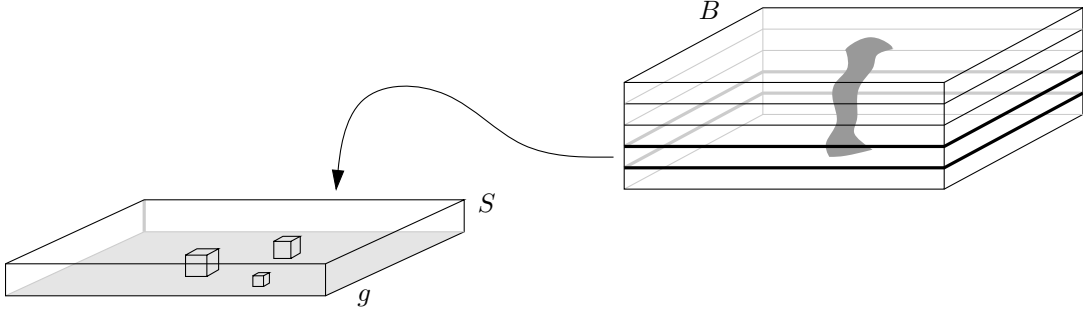

Figure 3: A 3-block B (with a crossing object shown shaded) is cut into l slices; each slice S contains a certain number of cubes (sharing a face with the face g of S) that must contain at least one guard.

also crosses each of the l slices. Consider a slice S and let g be one of its two largest faces. The intersections with g of the m objects crossing B are $(d-1)$-dimensional objects.

Lemma 3.2 states that we can identify at least $\lfloor m/(2^d - 2)(\kappa + 1)\rfloor$ disjoint $(d-1)$-dimensional hypercubes inside g, each of which is intersected by at least $\kappa + 1$ objects. Since the $(d-1)$-dimensional volume of g is 3^{d-1}, the number of $(d-1)$-dimensional hypercubes with a side length exceeding $1/l$ is less than $(3l)^{d-1}$, so at least $\lfloor m/(2^d - 2)(\kappa + 1)\rfloor - (3l)^{d-1}$ (which is positive) of such hypercubes have a side length of $1/l$ or less. For every such hypercube σ, take the d-dimensional hypercube σ' with σ as a face and contained in the slice S; this is possible because the side length of σ' is at most $1/l$ (see Figure 3). The hypercube σ' is intersected by at least $\kappa + 1$ objects because σ is intersected by $\kappa + 1$ objects, and hence it must contain a guard. It follows that we need at least $\lfloor m/(2^d - 2)(\kappa + 1)\rfloor - (3l)^{d-1}$ guards per slice, which sums up to a total of

$$l\left(\left\lfloor \frac{m}{(2^d - 2)(\kappa + 1)}\right\rfloor - (3l)^{d-1}\right) \geqslant (d-1)3^{d-1}\left\lfloor \frac{m}{d3^{d-1}(2^d - 2)(\kappa + 1)}\right\rfloor^{\frac{d}{d-1}}$$

guards for the entire 3-block B. $\qquad\qquad\square$

Combining the two lemmas above, we now prove that the number of relatively large objects intersecting a hypercube cannot exceed (roughly) the number of guards in its vicinity to the power $1 - 1/d$.

Theorem 3.4 *Let G be a κ-guarding set for a set $\mathcal{O}$ of objects in $\mathbb{R}^d$, with $d \geqslant 2$. Any hypercube whose vicinity contains exactly g guards from G is intersected by $O(\kappa(1+g^{1-1/d}))$ objects from $\mathcal{O}$ that are at least as large as σ.*

Proof: Let m denote the number of objects at least as large as σ intersecting σ. From Lemma 3.1 we know that there is a 3-block B in the vicinity of σ that is crossed by at least

$m/(2d3^d)$ object curves. Lemma 3.3 now implies that there must be at least

$$(d-1)3^{d-1} \left\lfloor \frac{m}{2d^2 3^{d(d-1)}(2^d-2)(\kappa+1)} \right\rfloor^{\frac{d}{d-1}}$$

guards in B. Since B is in the vicinity of σ, this number must be less than or equal to g, which (together with the fact that B can still be intersected by κ objects if it contains no guards) implies the theorem. $\qquad\square$

We now turn our attention to the complexity of motion planning in two-dimensional workspaces that are either uncluttered or have small simple-cover complexity, and then extend the obtained results to three dimensions in Section 5.

4 The complexity of motion planning in planar workspaces

Let $\mathcal{R}$ be a robot with f degrees of freedom, moving in a two-dimensional workspace amidst a set $\mathcal{O}$ of n obstacles. The robot $\mathcal{R}$ can be of any type: it can be a free-flying robot, a robotic arm, and so on. The only restriction is that it must have *bounded reach* [16], which is defined as follows. Let $p_{\mathcal{R}}$ be an arbitrary reference point inside $\mathcal{R}$. Then the *reach* of $\mathcal{R}$, denoted by $\mathrm{reach}(\mathcal{R})$, is defined as the maximum distance that any point of $\mathcal{R}$ can be from $p_{\mathcal{R}}$, taken over all possible configurations of $\mathcal{R}$. For instance, if $\mathcal{R}$ consists of two links of length 1 that are both attached to the origin, and the reference point is the tip of one of the links, then the reach of $\mathcal{R}$ is 2. (If the reference point would be the origin then the reach would be 1. For any two reference points, however, the two values $\mathrm{reach}(\mathcal{R})$ can be at most a factor of two apart.) A bounded-reach robot is now defined as a robot $\mathcal{R}$ with

$$\mathrm{reach}(\mathcal{R}) \leqslant c \cdot \min_{C \in \mathcal{O}}\{\mathrm{size}(C)\},$$

where c is a (small) constant.

In this section we study the complexity of the free space of a bounded-reach robot $\mathcal{R}$ under the assumption that the set of obstacles satisfies one of the models defined above. We prove an $\Omega(\kappa^f n^{f/2} + n)$ worst-case lower bound on the free-space complexity for the most restricted model, namely for κ-cluttered scenes. Because unclutteredness implies small simple-cover complexity in the hierarchy of input models [5], this bound carries over to scenes with small simple-cover complexity. Moreover, we prove an $O(\kappa^f((sn)^{f/2} + sn))$ upper bound for scenes with a κ-guarding set of size $s \cdot n$. By the conclusions from Section 2, the upper bound immediately carries over to uncluttered scenes and scenes with small simple-cover complexity. Hence, in both models we get a tight bound of $\Theta(n^{f/2} + n)$.

4.1 A lower bound for uncluttered scenes

The robot $\mathcal{R}$ in our lower bound example consists of f links, which are all attached to the origin. The links have length $1 + \varepsilon$, for a sufficiently small $\varepsilon > 0$. Obviously $\mathcal{R}$ has f degrees of freedom.

The set of n obstacles for the case of a 2-cluttered planar scene is defined as follows. (Later we adapt the construction to get the bound for κ-cluttered scenes for larger but still constant κ.) Recall that our obstacles are presumed to be two-dimensional. Fix an integer
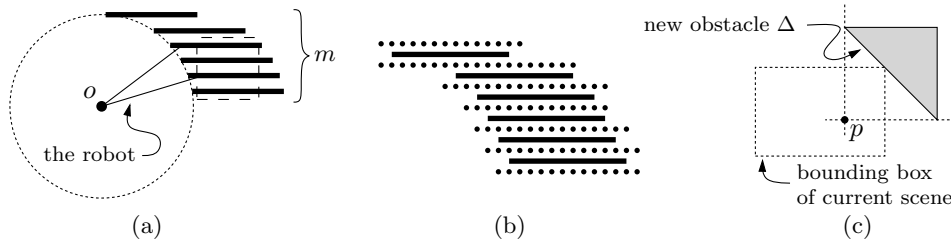
Figure 4: (a) Part of the lower bound construction. (b,c) Adding bounding-box vertices to make the scene uncluttered.

parameter m; it will turn out later that the appropriate value for m is roughly $\sqrt{n}$. For a given integer i, let C_i be the horizontal rectangle of length 1 and small height ϵ' whose lower left corner lies on the unit circle and has a y-coordinate equal to i/m—see Figure 4(a) for an example. Let $\mathcal{O}_1 := \{C_i \mid 1 \leqslant i \leqslant m\}$; this set forms a subset of the set of all obstacles. The remaining obstacles, which we describe later, are only needed to turn the environment into an uncluttered environment.

Consider any subset of f rectangles from $\mathcal{O}_1$. It is easy to see that there is a semi-free placement of $\mathcal{R}$ such that each rectangle in the subset is touched by a link of $\mathcal{R}$. Hence, the free-space complexity is $\Omega(m^f)$. When m is large, however, the set $\mathcal{O}_1$ does not form an uncluttered enviroment: the dashed square in Figure 4(a) for instance, intersects $\Omega(m)$ obstacles without having a bounding-box vertex of one of the rectangles in its interior. This problem would disappear if between every pair of adjacent horizontal rectangles there would be a collection of $\Theta(m)$ equal-spaced bounding-box vertices, as in Figure 4(b). If the distance between consecutive vertices is set to $1/2m$ then no square without a bounding-box vertex in its interior will intersect more than one obstacle from $\mathcal{O}_1$. Notice that in total we need $\Theta(m^2)$ bounding-box vertices for this. We cannot add tiny obstacles between the rectangles to achieve this, because such obstacles would be much smaller than the robot, so the robot would no longer have bounded reach. There is no need, however, to add obstacles between the rectangles; we can also create bounding-box vertices there by adding obstacles outside the current scene. Suppose that we wish to have a bounding-box vertex at a given point $p = (p_x, p_y)$, and suppose that the current set of obstacles is contained in the rectangle $[x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}]$. Then we add the right triangle Δ with vertices $(p_x, y_{\max}+x_{\max}-p_x)$, $(x_{\max}+y_{\max}-p_y, p_y)$, and $(x_{\max}+y_{\max}-p_y, y_{\max}+x_{\max}-p_x)$ as an obstacle—see Figure 4(c). The point p is a bounding-box vertex of Δ, and Δ is disjoint from the current set of obstacles. By iteratively adding obstacles that generate the necessary bounding-box vertices between the rectangles in $\mathcal{O}_1$ we transform the cluttered environment into an uncluttered one. The added obstacles are collected in a set $\mathcal{O}_2$; our final set of obstacles is $\mathcal{O} = \mathcal{O}_1 \cup \mathcal{O}_2$. It is not difficult to see that these obstacles form a 2-cluttered environment in this manner: any square without bounding-box vertices intersects at most one obstacle from $\mathcal{O}_1$ or two obstacles from $\mathcal{O}_2$.

We now have a collection of $\Theta(m^2)$ obstacles forming a 2-cluttered scene such that the free-space complexity is $\Omega(m^f)$. By choosing a suitable value for m (in the order of $\sqrt{n}$), we obtain a collection of n obstacles such that the free-space complexity is $\Omega(n^{f/2})$.

To get the general bound we replace each of the m rectangles in the set $\mathcal{O}_1$ by κ (even thinner) rectangles of length 1 that are quite close together. The lower left corners of these

9

rectangles still lie on the unit circle; the new scene is κ-cluttered. It is still possible to choose the value ε, which determines the length of the links of $\mathcal{R}$, small enough such that any f-tuple of rectangles in the new set $\mathcal{O}_1$ can be touched by a semi-free placement. Hence, the number of f-fold contacts has increased to $\Omega(\kappa^f m^f)$. By again choosing m to be roughly $\sqrt{n}$ we get a bound of $\Omega(\kappa^f n^{f/2})$. In the specific case that $f = 1$ the maximum complexity is clearly $\Omega(n)$.

Theorem 4.1 *The free-space complexity of a bounded-reach robot with f degrees of freedom moving in a two-dimensional κ-cluttered scene consisting of n obstacles can be $\Omega(\kappa^f n^{f/2} + n)$.*

4.2   An upper bound for scenes with linear-size guarding sets

We want to prove an upper bound on the complexity of the free space of a bounded-reach robot with f degrees of freedom moving in a scene with a linear-size κ-guarding set. The global structure of our proof will be as follows. We construct a decomposition of the workspace into cells that are not much smaller than the robot. The decomposition will have the property that none of its cells can have too many obstacles close to it. This means that the robot cannot have too many f-fold contacts when its reference point lies inside any given cell. Summing the number of f-fold contacts over all the cells using Theorem 3.4 yields the desired bound on the number of features of the free space.

The decomposition we use is obtained by adapting (the first stage of) the partitioning scheme described by de Berg [4]. First we describe the exact properties that we require, and then show how to obtain a decomposition with the desired properties.

Let $\rho := 2 \cdot \mathrm{reach}(\mathcal{R})$. Define the *expansion* $\hat{o}$ of an object o to be the Minkowski sum of o with a square of size 2ρ centered at the origin. Hence, $\hat{o}$ contains exactly those points that are at a L_∞-distance of less than ρ from o. Note that the expansion of a square σ is another square, whose edge length is 2ρ more than the edge length of σ. Let $\widehat{\mathcal{O}} := \{\widehat{C} \mid C \in \mathcal{O}\}$ denote the set of expanded obstacles.

Lemma 4.2 *Let $\mathcal{O}$ be a set of obstacles in $\mathbb{R}^2$ (or $\mathbb{R}^3$), and let G be a κ-guarding set for $\mathcal{O}$. Then there exists a set $\mathcal{S}$ of cells that are either squares (or cubes) or L-shapes with the following properties:*

(P1) the cells in $\mathcal{S}$ form a decomposition of a sufficiently large bounding square (or cube) of the set $\widehat{\mathcal{O}}$ of expanded obstacles;

(P2) the number of cells in $\mathcal{S}$ is $O(|G|)$;

(P3) every cell in $\mathcal{S}$ whose size is greater than 2ρ is intersected by $O(\kappa)$ expanded obstacles;

(P4) every cell in $\mathcal{S}$ whose size is less than or equal to 2ρ is a square (or cube) of size at least ρ.

Proof: We prove the lemma for the planar case; the generalization to three dimensions is straightforward.

Let $\widehat{G}$ denote the set of points obtained by adding to every guard $g \in G$ the four corner points of the square of size 2ρ centered at g. The set $\widehat{G}$ contains $5|G|$ points. We enclose $\widehat{\mathcal{O}}$ (and $\widehat{G}$) by a sufficiently large square and recursively decompose this square based on the points of $\widehat{G}$, as follows.

Let $\widehat{G}_\sigma$ denote the subset of points from $\widehat{G}$ contained in the interior of a square σ at some stage in the (quadtree-like) subdivision process. The square σ is handled according to the following set of rules.

10

1. If $\text{size}(\sigma) \leqslant 2\rho$ or $\widehat{G}_\sigma = \emptyset$ then σ is one of the cells in $\mathcal{S}$.

2. If $\text{size}(\sigma) > 2\rho$, $\widehat{G}_\sigma \neq \emptyset$, and not all points of $\widehat{G}_\sigma$ lie in the interior of a single quadrant of σ, then σ is subdivided into four quadrants, which are handled recursively.

3. If $\text{size}(\sigma) > 2\rho$, $\widehat{G}_\sigma \neq \emptyset$, and all points of $\widehat{G}_\sigma$ lie in the interior of a single quadrant of σ, then σ is subdivided as follows. Let σ' be the smallest square containing the points from $\widehat{G}_\sigma$ in its closure that shares a vertex with σ.

   (a) If $\text{size}(\sigma') > \rho$ then σ' is handled recursively, and the L-shape $\sigma \setminus \sigma'$ is a cell in $\mathcal{S}$.

   (b) If $\text{size}(\sigma') \leqslant \rho$, then let σ'' denote a square of size ρ contained in σ and containing σ'. The square σ'' and the L-shape $\sigma \setminus \sigma''$ are cells in $\mathcal{S}$.

It follows immediately from the construction that the cells in $\mathcal{S}$ satisfy Properties (P1) and (P4). A subdivision according to rule 2 splits the set $\widehat{G}_\sigma$ into two nonempty subsets; a subdivision according to rule 3 puts one of the points of $\widehat{G}_\sigma$ onto the boundary of a subcell. Both subdivisions can therefore be performed at most $|\widehat{G}|$ times. As a result, the number of cells will be $O(|\widehat{G}|) = O(|G|)$, which proves (P2).

It remains to prove Property (P3). By construction, any cell of size more than 2ρ contains no point from $\widehat{G}$. We now prove that any square σ of size at least ρ that does not contain any points from $\widehat{G}$ intersects at most κ expanded obstacles. If the cell under consideration is a square this immediately proves (P3), and if it is an L-shape then it also proves (P3) because an L-shape of size at least 2ρ can be covered by three squares of size at least ρ. So consider a square σ without points from $\widehat{G}$ and whose size is at least ρ. The fact that σ contains no points from $\widehat{G}$ implies that its expansion $\hat{\sigma}$ contains no guard from G—see Figure 5 for an illustration. This means that $\hat{\sigma}$ is intersected by at most κ original obstacles, which implies
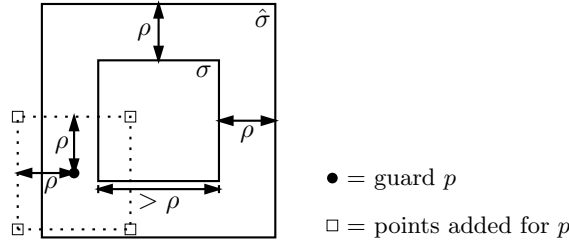

Figure 5: A guard from G inside $\hat{\sigma}$ implies a point of $\widehat{G}$ in σ.

that σ is intersected by at most κ expanded obstacles. $\qquad\square$

Now that we have a suitable decomposition of the workspace, we can use Theorem 3.4 to prove our main result.

Theorem 4.3 *Let $\mathcal{R}$ be a bounded-reach robot with f degrees of freedom, with f a constant, moving in a two-dimensional workspace containing a set $\mathcal{O}$ of n obstacles. If the set of obstacles has a κ-guarding set of size $s \cdot n$, then the complexity of the free space is $O(\kappa^f((sn)^{f/2} + sn))$.*

Proof: If $\mathcal{R}$ touches an obstacle C, its reference point must lie in the interior of $\widehat{C}$. (This is true because we defined ρ as *twice* the reach of $\mathcal{R}$.) Therefore we can bound, for any $1 \leqslant k \leqslant f$, the number of k-fold contacts of $\mathcal{R}$ by bounding the number of k-tuples of

expanded obstacles with a non-empty common intersection. The idea of the proof is to decompose the workspace according to Lemma 4.2 and then sum the number of k-tuples over all cells of the decomposition using Theorem 3.4.

Let G be a κ-guarding set of size sn for the obstacle set $\mathcal{O}$, and let $\mathcal{S}$ denote a decomposition having the properties stated in Lemma 4.2. To bound the free-space complexity we have to bound the number of simultaneous contacts involving k obstacles, for $k = 1, \ldots, f$. By Property (P1) this means that the free-space complexity is bounded by

$$\sum_{k=1}^{f} \sum_{\sigma \in \mathcal{S}} n_{\sigma}^{k},$$

where n_{σ} denotes the number of expanded obstacles intersecting the cell σ. The asymptotic value of this sum is dominated by the term where $k = f$, so we ignore the other terms from now on. Let $\mathcal{S}_1$ be the subset of $\mathcal{S}$ consisting of the cells of size larger than 2ρ, and let $\mathcal{S}_2$ be the subset of $\mathcal{S}$ consisting of the remaining cells. By Properties (P2) and (P3) we have

$$\sum_{\sigma \in \mathcal{S}_1} n_{\sigma}^{f} = O(sn\kappa^{f}).$$

Now consider the cells in $\mathcal{S}_2$. By Property (P4) these cells are squares whose size lies between ρ and 2ρ. Let σ be such a square. We claim that the number of expanded obstacles intersecting σ is $O(\kappa(1 + \sqrt{g_{\sigma}}))$, where g_{σ} is the number of guards from G in the vicinity of the expansion $\hat{\sigma}$. It is important to observe that $3\rho \leqslant \text{size}(\hat{\sigma}) \leqslant 4\rho$. Furthermore, any expanded obstacle intersecting σ corresponds to an original obstacle that intersects $\hat{\sigma}$. Because $\rho \leqslant 2c \cdot \min\{\text{size}(C) \mid C \in \mathcal{O}\}$ for a constant c, we can partition $\hat{\sigma}$ into $O(1)$ subsquares whose size is smaller than the size of the smallest obstacle. By Theorem 3.4, this means that the number of original obstacles intersecting $\hat{\sigma}$ is $O(\kappa(1 + \sqrt{g_{\sigma}}))$, where g_{σ} is the number of guards in the vicinity of $\hat{\sigma}$. Hence, the number of expanded obstacles intersecting σ is bounded by this quantity as well.

We conclude that the number of f-tuples of expanded obstacles with a non-empty common intersection in a cell of $\mathcal{S}_2$ is bounded by

$$\sum_{\sigma \in \mathcal{S}_2} n_{\sigma}^{f} = \sum_{\sigma \in \mathcal{S}_2} (\kappa(1 + \sqrt{g_{\sigma}}))^{f},$$

where g_{σ} is the number of guards in the vicinity of the expansion $\hat{\sigma}$. Since all squares have size at least ρ by Property (P4), the vicinity of an expanded square in $\mathcal{S}_2$ intersects $O(1)$ other vicinities of expanded squares. Hence, a guard from G lies in $O(1)$ vicinities, and we have

$$\sum_{\sigma \in \mathcal{S}_2} g_{\sigma} = O(sn),$$

which leads to

$$\sum_{\sigma \in \mathcal{S}_2} (\kappa(1 + \sqrt{g_{\sigma}}))^{f} = O(\kappa^{f}(sn)^{f/2}).$$

Therefore the total number of f-fold contacts of $\mathcal{R}$ is bounded by

$$\sum_{\sigma \in \mathcal{S}_1} n_{\sigma}^{f} + \sum_{\sigma \in \mathcal{S}_2} n_{\sigma}^{f} = O(\kappa^{f}((sn)^{f/2} + sn)).$$

$\square$

5 The complexity of motion planning in 3D workspaces

Having described the two-dimensional setting in the previous section, we now turn our attention to a robot R moving in a three-dimensional workspace amidst a set $\mathcal{O}$ of n obstacles. As in the two-dimensional case, the robot is allowed to be of any type—we only require that its reach is bounded. We prove an $\Omega(\kappa^{2f} n^{2f/3} + n)$ worst-case lower bound on the complexity of the free space for κ-cluttered scenes, and an $O(\kappa^f((sn)^{2f/3} + sn))$ upper bound for scenes with a κ-guarding of size $s \cdot n$. As before, this results in a tight bound of $\Theta(n^{2f/3} + n)$ for uncluttered scenes. We also prove an $\Omega(n^f)$ worst-case lower bound on the complexity of scenes with small simple-cover complexity. The maximum free-space complexity for such scenes is therefore $\Theta(n^f)$ and thus equivalent to the complexity for unrestricted scenes.

5.1 Lower bounds

A lower bound for scenes with small simple-cover complexity

We consider the scene consisting of n rings

$$C_i := \{(x, y, z) \mid \frac{i}{n} < x < \frac{i}{n} + \epsilon, \ 1 < y^2 + z^2 < 1 + \epsilon\}$$

with $0 \leqslant i < n$ and small ϵ shown in Figure 6. It was shown [6] that a similar scene consisting of unit circles has small simple-cover complexity but requires a κ-guarding set of size $\Omega(n^2)$ for any constant κ. It is clear that these properties carry over to our scene in which the circles are replaced by thin rings. Our robot $\mathcal{R}$ consists of f links, which are all attached to the point $(0, 2, 0)$. Each link has length 2 and rotates about the axis $h = \{(x, y, z) \mid x = 0, y = 2\}$, causing it to stay inside the xy-plane. Note that the size of the robot is comparable to the size of the obstacles so it has bounded reach. For any subset of f rings, there is a semi-
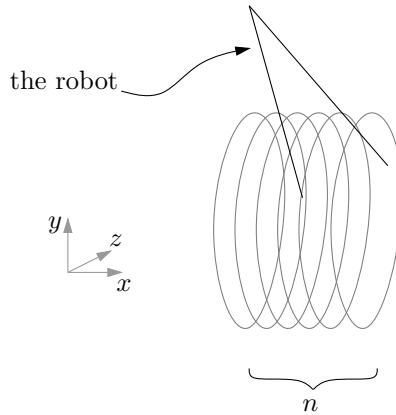

Figure 6: The lower-bound construction.

free placement of $\mathcal{R}$ such that each ring C_i in the subset is touched by a link of $\mathcal{R}$. As a consequence, the free-space complexity is $\Omega(n^f)$.

Theorem 5.1 *The free-space complexity of a bounded-reach robot with f degrees of freedom moving in a three-dimensional scene with small simple-cover complexity consisting of n obstacles can be $\Omega(n^f)$.*

13

A lower bound for uncluttered scenes

Our approach to obtaining a worst-case lower bound for 3D κ-cluttered scenes is similar to the planar case. We fix a parameter m and consider the set $\mathcal{O}_1$ of m^2 thin unit-length rectangloids

$$C_{i,j} = \{(x,y,z) \mid \frac{i}{m} < x < \frac{i}{m} + \epsilon, \; \frac{j}{m} < y < \frac{j}{m} + \epsilon, \; 0 < z < 1\}$$

for $0 \leqslant i, j < m$.

Consider the m^2 planes through pairs of obstacles and choose a point $q = (x_q, y_q, 1/2)$ that lies on none of these planes and satisfies $1/2 < x_q, y_q < 1/2 + 1/m$. Our robot $\mathcal{R}$ has f links, which are all anchored at q. Each link has length 1 and is able to rotate about the axis $h = \{(x,y,z) \mid x = x_q, \; y = y_q\}$, causing it to stay inside the plane $z = 1/2$; the choice of q allows it be placed in contact with each obstacle in $\mathcal{O}_1$. Figure 7(a) shows the intersection of the scene with the plane $z = 1/2$. The size of $\mathcal{R}$ is comparable to the size of the obstacles in
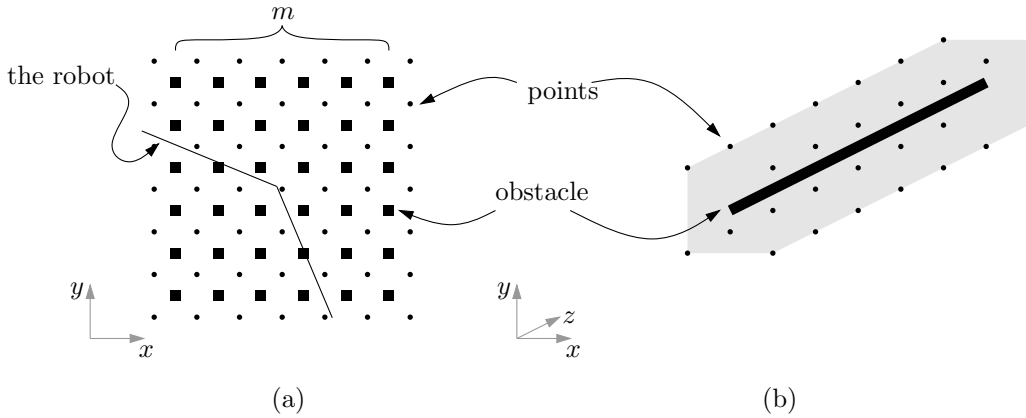

Figure 7: (a) Cross-section of the environment at $z = 1/2$. (b) Magnified portion of the scene showing a single obstacle and the four sequences of points (bounding-box vertices) immediately surrounding it.

$\mathcal{O}_1$. For any subset of f obstacles there is a semi-free placement of $\mathcal{R}$ such that each obstacle $C_{i,j}$ in the subset is touched by a link of $\mathcal{R}$. Hence, the complexity of the free space is $\Omega(m^{2f})$.

To prevent cubes from intersecting more than one obstacle from $\mathcal{O}_1$ we put sequences of points on each of the lines $l = \{(x,y,z) \mid x = (2i-1)/2m, \; y = (2j-1)/2m\}$ for $0 \leqslant i, j \leqslant m$, see Figure 7(b) for one obstacle and the points surrounding it. The distance between two consecutive points on a single line is again equal to $1/2m$. We turn the $\Theta(m^3)$ points into bounding-box vertices by iteratively adding tetrahedral obstacles in a way similar to the planar case. The resulting $\Theta(m^3)$ obstacles are collected in $\mathcal{O}_2$, and our final set of obstacles is $\mathcal{O} = \mathcal{O}_1 \cup \mathcal{O}_2$. Any cube without bounding box vertices intersects at most one obstacle from $\mathcal{O}_1$ or two obstacles from $\mathcal{O}_2$.

As a result, we now have a collection of $\Theta(m^3)$ obstacles forming a 2-cluttered scene with a free-space complexity of $\Omega(m^{2f})$. By choosing $m = n^{1/3}$, the free-space complexity becomes $\Omega(n^{2f/3})$ for the set of n obstacles.

As in the two-dimensional example, we now replace each of the m obstacles in $\mathcal{O}_1$ by κ obstacles that are close together and arranged such that each of them can be touched by the links of $\mathcal{R}$. The resulting scene is κ-cluttered, and the number of f-fold contacts increases

from $\Omega(m^{2f})$ to $\Omega(\kappa^f m^{2f})$. The theorem follows from again choosing $m = \sqrt[3]{n}$. By again choosing m to be roughly $n^{1/3}$ we get a bound of $\Omega(\kappa^f n^{2f/3})$. Again the maximum complexity is $\Omega(n)$ in the case that $f = 1$.

Theorem 5.2 *The free-space complexity of a bounded-reach robot with f degrees of freedom moving in a three-dimensional κ-cluttered scene of n obstacles can be $\Omega(\kappa^f n^{2f/3} + n)$.*

5.2 An upper bound for scenes with linear-size guarding sets

Proving an upper bound on the free-space complexity of a bounded-reach robot with f degrees of freedom moving in a scene with a linear-size guarding set is entirely analogous to the two-dimensional case.

Theorem 5.3 *Let $\mathcal{R}$ be a bounded-reach robot with f degrees of freedom, with f a constant, moving in a three-dimensional workspace containing a set $\mathcal{O}$ of n obstacles. If the set of obstacles has a κ-guarding set of size $s \cdot n$, then the complexity of the free space is $O(\kappa^f((sn)^{2f/3} + sn))$.*

Proof: Analogous to the proof of Theorem 4.3, except that

$$\sum_{\sigma \in \mathcal{S}_2} n_\sigma^f = \sum_{\sigma \in \mathcal{S}_2} (\kappa(1 + g_\sigma^{2/3}))^f = O(\kappa^f (sn)^{2f/3}),$$

which yields

$$\sum_{\sigma \in \mathcal{S}_1} n_\sigma^f + \sum_{\sigma \in \mathcal{S}_2} n_\sigma^f = O(\kappa^f((sn)^{2f/3} + sn)).$$

$\square$

6 Conclusion

We have established that the maximum complexity of the free-space of a bounded-reach robot with f degrees of freedom moving in an uncluttered scene is $\Theta(n^{f/2} + n)$ in $\mathbb{R}^2$ and $\Theta(n^{2f/3} + n)$ in $\mathbb{R}^3$; the planar bound also holds for scenes with small simple-cover complexity. These bounds fit nicely between the $\Theta(n)$ bound for low-density scenes—which are more restrictive—, and the $\Theta(n^f)$ bound for unrestricted scenes. Surprisingly, the maximum complexity of the free space for a robot moving in a 3D scene with small simple-cover complexity is equal to $\Theta(n^f)$—the bound for unrestricted scenes.

Motion planning in low-density environments can be solved in an amount of time that is almost equal to the maximum free-space complexity [20]. It is interesting to see if a similar result is possible for uncluttered scenes and scenes with small simple-cover complexity.

References

[1] P.K. Agarwal, M. Katz, and M. Sharir. Computing depth orders for fat objects and related problems. *Comput. Geom. Theory Appl.* **5** (1995), pp. 187–206.

[2] H. Alt, R. Fleischer, M. Kaufmann, K. Mehlhorn, S. Näher, S. Schirra, and C. Uhrig, Approximate motion planning and the complexity of the boundary of the union of simple geometric figures. *Algorithmica* **8** (1992), pp. 391–406.

[3] J.M. Bañon, Implementation and extension of the ladder algorithm, *Proc. IEEE Internat. Conf. Robot. Autom.* (1990), pp. 1548-1553.

[4] M. de Berg, Linear size binary space partitions for uncluttered scenes, *Algorithmica* (2000), to appear.

[5] M. de Berg, M.J. Katz, A.F. van der Stappen, and J. Vleugels, Realistic input models for geometric algorithms, *Proc. 13th Annu. ACM Sympos. Comput. Geom.* (1997), pp. 294-303.

[6] M. de Berg, H. David, M.J. Katz. M. Overmars, A.F. van der Stappen, and J. Vleugels, Guarding scenes against invasive hypercubes, Tech. Rep. UU-CS-2000-40, Inst. Inform. Comput. Sci., Utrecht Univ., Utrecht, the Netherlands (2000).

[7] A. Efrat and M.J. Katz, On the union of κ-curved objects, *Comput. Geom. Theory Appls.* **14** (1999), pp. 241-254.

[8] A. Efrat, M.J. Katz, F. Nielsen, and M. Sharir. Dynamic data structures for fat objects and their applications. *Comput. Geom. Theory Appls.* **15** (2000), pp. 215-227.

[9] M.J. Katz, 3-D Vertical ray shooting and 2-d point enclosure, range searching, and arc shooting amidst convex fat objects, *Comput. Geom. Theory Appl.* **8** (1997), pp. 299-316.

[10] M.J. Katz, M.H. Overmars, and M. Sharir, Efficient hidden surface removal for objects with small union size, *Comput. Geom. Theory Appl.* **2** (1992), pp. 223–234.

[11] J. Matoušek, J. Pach, M. Sharir, S. Sifrony, and E. Welzl, Fat triangles determine linearly many holes, *SIAM J. Comput.* **23** (1994), pp. 154-169.

[12] J.S.B. Mitchell, D.M. Mount, and S. Suri, Query-sensitive ray shooting, *Internat. J. Comput. Geom. Appl.* **7** (1997), pp. 317-347.

[13] M.H. Overmars and A.F. van der Stappen, Range searching and point location among fat objects, *J. Algorithms* **21** (1996), pp. 629-656.

[14] 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, *Commun. Pure Appl. Math.* **36** (1983), pp. 345-398.

[15] O. Schwarzkopf and J. Vleugels, Range searching in low-density environments, *Inform. Process. Lett.* **60** (1996), pp. 121-127.

[16] A.F. van der Stappen, D. Halperin, and M.H. Overmars, The complexity of the free space for a robot moving amidst fat obstacles, *Comput. Geom. Theory Appl.* **3** (1993), pp. 353–373.

[17] A.F. van der Stappen, The complexity of the free space for motion planning amidst fat obstacles, *J. Intell. Robotic Syst.* **11** (1994), pp. 21-44.

[18] A.F. van der Stappen, *Motion planning amidst fat obstacles*. Ph.D. thesis, Dept. Comput. Sci., Utrecht Univ., Utrecht, the Netherlands (1994).

[19] A.F. van der Stappen and M.H. Overmars, Motion planning amidst fat obstacles, *Proc. 10th Annu. ACM Sympos. Comput. Geom.* (1994), pp. 31–40.

[20] A. F. van der Stappen, M. H. Overmars, M. de Berg, and J. Vleugels. Motion planning in environments with low obstacle density. *Discr. Comput. Geom.* **20** (1998), pp. 561–587.

[21] J. Vleugels, *On fatness and fitness—Realistic input models for geometric algorithms*, Ph.D. thesis, Dept. Comput. Sci., Utrecht Univ., Utrecht, the Netherlands (1997).