Motion Planning of a Ball Amid Segments in Three Dimensions*

Pankaj K. Agarwal†   Micha Sharir‡

Abstract

Let S be a set of n pairwise disjoint segments in $\mathbb{R}^3$, and let B be a ball of radius 1. The *free configuration space* $\mathcal{F}$ of B amid S is the set of all placements of B at which (the interior of) B does not intersect any segment of S. We show that the combinatorial complexity of $\mathcal{F}$ is $O(n^{5/2+\varepsilon})$, for any $\varepsilon > 0$, with the constant of proportionality depending on ε. This is the first subcubic bound on the complexity of the free configuration space even when S is a set of lines in $\mathbb{R}^3$. We also present a randomized algorithm that can compute the boundary of the free configuration space in $O(n^{5/2+\varepsilon})$ expected time.

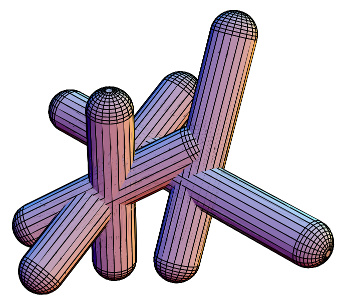

Figure 1: Union of the Minkowski sums of five segments with a ball.

1 Introduction

Problem statement. Let S be a collection of n pairwise disjoint segments in $\mathbb{R}^3$, and let B be a ball of radius 1. We regard S as a set of obstacles and consider the motion-planning problem in which B is allowed to move (translate) freely in $\mathbb{R}^3$ without intersecting any segment of S. The *free configuration space* $\mathcal{F}$ of B with respect to S is the set of all points $p \in \mathbb{R}^3$ so that if B is placed centered at p then it does not intersect any segment of S. We wish to bound the combinatorial complexity of $\mathcal{F}$ (defined below) and present an efficient algorithm for computing the boundary of $\mathcal{F}$.

Let B_0 be the placement of B with its center at the origin. For each $s \in S$, let K_s denote the *Minkowski sum* $s \oplus B_0 = \{x + y \mid x \in s, y \in B_0\}$. The set K_s, also referred to as the *expanded obstacle* of s, is the set of all centers of B at placements where it intersects s. Note that K_s is a bounded cylinder with two hemispheres attached to its bases; K_s is an infinite cylinder if s is a line. $\mathcal{F}$ is the complement of the union

$$U = U(S, B) = \bigcup_{s \in S} K_s$$

of the expanded obstacles; see Figure 1. The *combinatorial complexity* of U (and of $\mathcal{F}$) is the total number of faces of all dimensions (vertices, edges, and 2-faces) on the boundary of U (where a face is a (relatively open) maximal connected portion of ∂U that is contained in a fixed subset of boundaries of expanded obstacles and avoids all other boundaries). We want to bound the combinatorial complexity of U and describe an efficient algorithm for computing vertices, edges, and 2-faces of ∂U. Besides this motion-planning application, the problem of bounding the complexity of U is a precursor to the harder problem of obtaining a subcubic bound on the complexity of the Euclidean Voronoi diagram of S. Indeed, if the radius of B is r, then ∂U is the locus of all points whose Euclidean distance from their nearest segment in S is exactly r. In this sense, ∂U is a cross-section of the Voronoi diagram of S.

Previous results. Motivated by the motion-planning application, there has been much work on bounding the combinatorial complexity of the union of the Minkowski sums of a geometric object ('robot') with a family of geometric objects ('obstacles'), or, more generally, the complexity of the union of a set of geometric objects. See the book [17] and the survey paper [4] by the authors for a summary of known results on this topic. Boissonnat *et al.* [8] proved that the union of n axis-parallel hypercubes in $\mathbb{R}^d$ is $O(n^{\lceil d/2 \rceil})$; the bound improves to $O(n^{\lfloor d/2 \rfloor})$ if all hypercubes have the same size. Aronov *et al.* [7] proved that the complexity of the union

*Work by P.A. was supported by Army Research Office MURI grant DAAH04-96-1-0013, by a Sloan fellowship, by an NYI award, and by a grant from the U.S.-Israeli Binational Science Foundation. Work by M.S. has been supported by NSF Grants CCR-97-32101, CCR-94-24398 and CCR-93-11127, by grants from the U.S.-Israeli Binational Science Foundation, the G.I.F., the German-Israeli Foundation for Scientific Research and Development, and the ESPRIT IV LTR project No. 21957 (CGAL), and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

†Center for Geometric Computing, Department of Computer Science, Box 90129, Duke University, Durham, NC 27708-0129, USA. E-mail: `pankaj@cs.duke.edu`

‡School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel; and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA. E-mail: `sharir@math.tau.ac.il`

of n convex polyhedra in $\mathbb{R}^3$ with a total of s faces is $O(n^3 + sn \log n)$. Aronov and Sharir [6] proved that complexity of the union of the Minkowski sums of a convex polyhedron P with a set S of n convex polyhedra in $\mathbb{R}^3$ is $O(ns \log n)$, where s is the total number of faces in the set $\{P \oplus Q \mid Q \in S\}$. All these bounds are either optimal or near optimal in the worst case. These recent results concern unions in higher dimensions, and extend work on unions of objects in the plane. Among these, we mention the early result of Kedem *et al.* [14] that shows that the complexity of the union of n disks is $O(n)$, and the reults of Matoušek *et al.* [15] and of Efrat and Sharir [13] that prove near-linear bounds on the complexity of the union of "fat" triangles and general "fat" convex regions in the plane.

It is conjectured that Voronoi diagrams in three dimensions, under fairly general assumptions concerning the sites and the distance function, have near-quadratic complexity. A near-cubic bound on the complexity of such diagrams follows from the results in [16]. The maximum complexity of Voronoi diagrams of n point sites under the Euclidean distance is known to be $\Theta(n^2)$ [12]. The same bound has recently been established for point sites under the L_1 and L_∞ metrics, or under any simplicial distance function [8]. Near-quadratic bounds have also been recently established for the case of line sites and any polyhedral convex distance function [11], where the bound is $O(n^2\alpha(n) \log n)$, and for the case of point sites and any polyhedral convex distance function [18], where the bound is $O(n^2 \log n)$ (in both cases the distance function is induced by a convex polytope with a constant number of facets). No example with a substantially super-quadratic complexity (i.e. $\Omega(n^{2+c})$, for any fixed $c > 0$) is known. As noted above, any of these results also yields near-quadratic bounds on the complexity of the corresponding union of the Minkowski sums of the sites and of the unit ball under the given distance function.

Our results. If the conjecture on the complexity of the Voronoi diagram is true for the case of segment sites and Euclidean distance, then the complexity of U will be near-quadratic, and indeed this is what we continue to conjecture. The main result of this paper is an $O(n^{5/2+\varepsilon})$ upper bound on the complexity of U, for any $\varepsilon > 0$. Although this result falls short of proving the above conjecture, it is a significant result because it yields the first subcubic bound on this quantity (note that a cubic bound is trivial), even if S is a set of lines. Note that if S is a set of lines, then U is the union of n congruent cylinders, so, as a by-product, we obtain the same upper bound on the complexity of the union of congruent cylinders.

The proof technique does not seem to exploit too much the particular geometry of the cylindrical and spherical surfaces appearing on the boundary of expanded obstacles. It thus lends hope to extending the result for the union of more general objects in $\mathbb{R}^3$. For example, we can prove that the complexity of the union of n smooth compact convex objects of roughly the same size and with bounded curvature is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$.

For the sake of simplicity, we first prove the bound on the combinatorial complexity of U in the special case in which S is a set of n lines, i.e., each K_s is a cylinder, all being congruent (Section 2). We then extend this result to segments (Section 3). In Section 4 we present a randomized algorithm for computing the boundary of U. Finally, we discuss extensions of our technique.

2 Complexity of U: The Case of Lines

Preliminaries and overview. Let $S = \{s_1, \ldots, s_n\}$ be a set of n lines and B a ball in $\mathbb{R}^3$. Without loss of generality, we assume that the radius of B is 1. Let $K_i = K_{s_i} = s_i \oplus B$ and $c_i = \partial K_i$; K_i is a cylinder of radius 1. Set $\mathcal{K} = \{K_1, \ldots, K_n\}$ and $\mathcal{C} = \{c_1, \ldots, c_n\}$. We assume that the lines in S are in general position, which means that no pair of them are parallel or intersect, that no two K_i's are tangent to each other, that no curve of intersection of the boundaries of any two K_i's is tangent to a third one, and that no four boundaries meet at a point. Using an argument based on random perturbation, like the one given in [16], it follows that this assumption can be made with no loss of generality.

Let V denote the set of *vertices* of U, namely, intersection points of triples of boundaries of regions in $\mathcal{K}$ that lie on ∂U. By our general position assumption, each vertex lies on exactly three cylindrical surfaces.

The number of edges or 2-faces of ∂U that are not incident to any vertex is $O(n^2)$. By the general-position assumption, each vertex is incident to only a constant number of edges and faces, therefore the combinatorial complexity of U is $O(n^2+|V|)$. In the rest of this section, we prove that $|V| = O(n^{5/2+\varepsilon})$.

The proof consists of three main steps. In the first step, for technical reasons, we choose a subset of cylinders in $\mathcal{K}$ whose union boundary contains at least half of the vertices of V. We also choose the z-axis (by rotating the coordinate frame) carefully so that the acute angle between the z-axis and the axes of every chosen cylinder is at most $\cos^{-1}(1/6)$. In the second stage, we mark a subset $V' \subseteq V$ of vertices that have some useful properties that are needed in the proof, and reduce the problem to bounding the size of V'. Next, we partition $\mathbb{R}^3$ into a carefully chosen infinite grid of square prisms. We prove that if m cylinders intersect a prism Q then at most $O(m^{2+\varepsilon})$ vertices appear on the union boundary within Q. (Informally, the partition into prisms allows us to regard vertices of the union within a prism as vertices of the region enclosed between an upper envelope and a lower envelope of appropriate portions of the cylindrical surfaces, and such a region has a near-quadratic complexity by the results of [3].) Using this result and a double-counting argument, we bound the size of V'. Putting all these steps together we obtain the desired bound on V. We now describe each step in detail.

Choosing the z-direction. Let $\mathbb{S}^2$ denote the unit sphere of directions in $\mathbb{R}^3$. We choose a random point on $\mathbb{S}^2$ and regard it as the $(+z)$-axis. Then the following claim holds.

Lemma 2.1 *Let β_0 be the acute angle satisfying $\cos\beta_0 = 1/6$. Let v be a vertex in V incident to three cylindrical surfaces $a, b, c \in \mathcal{C}$. The probability that all three acute angles between the z-direction and the axes of a, b, c are $\le \beta_0$ is at least $1/2$.*

Proof: Indeed, for the acute angle between the z-axis and the axis of, say, a to be $> \beta_0$, the z-direction has to lie in the spherical band consisting of all directions at spherical distance $\le \frac{\pi}{2} - \beta_0$ from the great circle orthogonal to the axis of a. The area of this band is $4\pi\cos\beta_0$. Hence the probability that at least one of the above three acute angles is $> \beta_0$ is at most

$$\frac{12\pi\cos\beta_0}{4\pi} = \frac{1}{2}.$$

□

Hence, if we choose the $(+z)$-direction to be a random direction, the expected number of vertices in V for which all three corresponding acute angles are $\le \beta_0$ is at least $|V|/2$. In particular, there exists a z-direction so that at least half of the vertices in V have the property that the axes of all their incident cylinders form angles $\le \beta_0$ with the z-axis. We now remove from $\mathcal{C}$ all cylinders whose axes have an acute angle $> \beta_0$ with the z-direction. At least half of the original vertices in V still show up as vertices of the new union. Abusing the notation slightly, we will use $\mathcal{C}$ to denote the remaining set of cylinders. In summary, we assume that all cylinders of $\mathcal{C}$ are such that the acute angles between their axes and the z-direction are all $\le \beta_0$.

Divergent pairs and the selection lemma. For each $c \in \mathcal{C}$, let $\mathbf{n}_c \in \mathbb{S}^2$ denote the unit vector parallel to the axis of c and having positive z-component (with no loss of generality, we may assume that no line in S is horizontal). Let D be the resulting set of points on $\mathbb{S}^2$. Let v be a vertex of U incident to three surfaces $a, b, c \in \mathcal{C}$. We associate with v the triple $\Delta_v = \{\mathbf{n}_a, \mathbf{n}_b, \mathbf{n}_c\}$.

We call a pair of cylinders $a, b \in \mathcal{C}$ *divergent* with respect to a direction $\mathbf{u}$ if

$$\min\{\angle(\mathbf{n}_a, \mathbf{u}), \angle(\mathbf{n}_b, \mathbf{u})\} \le \angle(\mathbf{n}_a, \mathbf{n}_b).$$

Lemma 2.2 *There exist a direction $\mathbf{u}$, a subset V' of V and a subset Ξ of $\mathcal{C}$ of n' cylinders, so that*

(i) $\dfrac{|V'|}{n'} \ge \dfrac{|V|}{cn\log^2 n}$ *for some constant $c > 1$;*

(ii) for each vertex $v \in V'$, at least two of the cylinders incident to v belong to Ξ; and

(iii) for each $v \in V'$, all three pairs of cylinders incident to v are divergent.

Proof: Let v be a vertex of the union that is incident to three cylinders $a, b, c \in \mathcal{C}$. Among the pairs (a, b), (a, c), and (b, c), associate v with the (unordered) pair that has the smallest angle between the corresponding vectors in D (in case of a tie, we associate v with an arbitrary pair that attains this minimum).

For each pair (a, b) of surfaces, define the *weight* of (a, b), denoted by $w_{(a,b)}$, to be the number of vertices v associated with (a, b). Clearly, $\sum_{(a,b)} w_{(a,b)} = |V|$. Moreover, let $B_{(a,b)}$ denote the *spherical bounding box* spanned by a, b, defined as the smallest spherical quadrangle bounded by two latitudes and two meridians that contains both $\mathbf{n}_a$ and $\mathbf{n}_b$. We assign to $B_{(a,b)}$ the weight $w_{(a,b)}$. By extending the weighted selection lemma (Lemma 2.1) of [5], in the same manner as done in Theorem 2.2 of [9] for the unweighted case, we obtain that there exists a point $\mathbf{u} \in \mathbb{S}^2$ that lies in bounding boxes with a total weight W', spanned by a set $D' \subseteq D$ of n' directions, where

$$\frac{W'}{n'} \ge \frac{|V|}{cn\log^2 n}$$

for some constant $c > 1$. Let V' be the subset of vertices of V that contribute to the weights of the boxes pierced by $\mathbf{u}$, and let Ξ be the subset of cylinders corresponding to directions in D'. It is clear that properties (i) and (ii) are satisfied by V' and Ξ. To see that (iii) also holds, let $v \in V'$ be a vertex of the union, incident to three cylinders $a, b, c \in \mathcal{C}$, such that v is assigned to the pair (a, b). The pair (a, b) is divergent, because the spherical distance between $\mathbf{u}$ and either $\mathbf{n}_a$ or $\mathbf{n}_b$ is smaller than the spherical distance between $\mathbf{n}_a$ and $\mathbf{n}_b$. Since this latter distance is smaller than or equal to the spherical distances between $\mathbf{n}_a$ and $\mathbf{n}_c$ and between $\mathbf{n}_b$ and $\mathbf{n}_c$, and since $\min\{\angle(\mathbf{n}_a, \mathbf{u}), \angle(\mathbf{n}_c, \mathbf{u})\}$ and $\min\{\angle(\mathbf{n}_b, \mathbf{u}), \angle(\mathbf{n}_c, \mathbf{u})\}$ are both at most $\max\{\angle(\mathbf{n}_a, \mathbf{u}), \angle(\mathbf{n}_b, \mathbf{u})\}$, the pairs (a, c) and (b, c) are also divergent. □

We will prove below in Lemma 2.4 that

$$|V'| = O(n^{3/2+\varepsilon} n').$$

Hence, by property (i) of the lemma,

$$|V| \le \frac{cn \log^2 n \cdot |V'|}{n'} = O(n^{5/2+\varepsilon'}), \qquad (1)$$

for any $\varepsilon' > \varepsilon$. It thus suffices to bound the size of V' by $O(n^{3/2+\varepsilon} n')$.

Bounding the size of V'. We place in $\mathbb{R}^3$ a grid $\mathcal{Q}$ of infinite square prisms whose axes are parallel to the direction $\mathbf{u}$; see Figure 2. In particular, if we let $\mathbf{u}$ be the z-axis, then the prisms are of the form $Q_{ij} = [ti, t(i+1)] \times [tj, t(j+1)] \times \mathbb{R}$, for $i, j \in \mathbb{Z}$, where t is a sufficiently small constant as specified in the proof of Lemma 2.5 below.

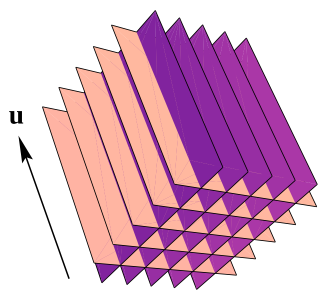

Figure 2: A system of prisms in direction $\mathbf{u}$.

For proving the bound on $|V'|$, we need the following lemma, which highlights the significance of divergent pairs.

Lemma 2.3 *If a and b are an intersecting pair of divergent cylindrical surfaces, then the intersection curve of a and b meets $O(1/t^2)$ prisms of $\mathcal{Q}$.*

Proof: Put $\gamma = \angle(\mathbf{n}_a, \mathbf{n}_b)$ and

$$\delta = \min\{\angle(\mathbf{n}_a, \mathbf{u}), \angle(\mathbf{n}_b, \mathbf{u})\};$$

by assumption, $\delta \le \gamma$, and our initial pruning step guarantees that $\gamma \le 2\beta_0$. Assume, with no loss of generality, that $\delta = \angle(\mathbf{n}_a, \mathbf{u})$.

Let us first analyze the intersection curve Γ between a and b, Let ℓ_a and ℓ_b denote the axes of a and b, respectively. Let λ denote the distance between ℓ_a and ℓ_b. Choose (temporarily) a coordinate frame in which ℓ_a is the z-axis, the xy-projection ℓ_b^* of ℓ_b is parallel to the y-axis, and the nearest distance between ℓ_a and ℓ_b materializes along the positive x-axis; see Figure 3 (i).

Thus ℓ_b passes through $\mathbf{p} = (\lambda, 0, 0)$ and its orientation is $\mathbf{n}_b = (0, \sin\gamma, \cos\gamma)$. Hence the equation of b is

$$\|\mathbf{x} - \mathbf{p}\|^2 - |(\mathbf{x} - \mathbf{p}) \cdot \mathbf{n}_b|^2 = 1,$$

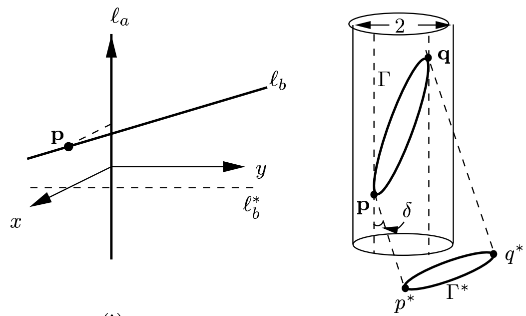

Figure 3: (i) Lines ℓ_a and ℓ_b. (ii) Intersection curve Γ and its projection Γ^*.

or

$$(x-\lambda)^2 + y^2 + z^2 - (y\sin\gamma + z\cos\gamma)^2 = 1,$$

or

$$(x-\lambda)^2 + (y\cos\gamma - z\sin\gamma)^2 = 1.$$

The curve Γ of intersection between a and b thus satisfies the equations

$$(x-\lambda)^2 + (y\cos\gamma - z\sin\gamma)^2 = 1, \; x^2 + y^2 = 1$$

or, alternatively,

$$z = \frac{y\cos\gamma \pm \sqrt{1-(x-\lambda)^2}}{\sin\gamma}, \quad x^2 + y^2 = 1.$$

Hence any point on Γ satisfies

$$|z| \le \frac{1 + |\cos\gamma|}{\sin\gamma} \le \frac{2}{\sin\gamma}.$$

We now project Γ orthogonally in the direction $\mathbf{u}$. Note that $\mathbf{u}$ makes an angle δ with the z-axis in our current transformed coordinate frame; see Figure 3 (ii). Clearly, any prism Q intersecting Γ also intersects its projection Γ^*. However, Γ^* is contained in the projection of the bounded cylinder

$$x^2 + y^2 = 1, \quad |z| \le 2/\sin\gamma. \qquad (2)$$

Suppose the projections of points $\mathbf{p} = (p_x, p_y, p_z)$ and $\mathbf{q} = (q_x, q_y, q_z)$ induce the diameter of Γ^*. Then

$$\begin{aligned} \mathrm{diam}(\Gamma^*) &\le \sqrt{(p_x - q_x)^2 + (p_y - q_y)^2} + (p_z - q_z)\sin\delta \\ &\le 2 + \frac{4}{\sin\gamma}\sin\delta. \end{aligned}$$

The last inequality follows from (2) and the fact that $\mathbf{u}$ makes an angle δ with the z-axis. We have

$$\frac{\sin\delta}{\sin\gamma} \le \begin{cases} 1 & \text{if } \gamma \le \pi/2, \\ \dfrac{1}{\sin(2\beta_0)} & \text{if } \pi/2 < \gamma \le 2\beta_0. \end{cases}$$

(This is the place where we need the property that γ is bounded away from π, and this is the reason for the initial pruning step of the analysis.) Hence,

$$\mathrm{diam}(\Gamma^*) \le 2 + \frac{72}{\sqrt{35}} \le 15.$$

Since the area of the projection of each prism is t^2, the number of prisms that intersect Γ^* is $O(1/t^2)$, as claimed. □

We are now in position to bound the size of V'.

Lemma 2.4 *$|V'| = O(n^{3/2+\varepsilon} n')$, for any $\varepsilon > 0$.*

Proof: Let $Q \in \mathcal{Q}$ be one of the prisms in the grid. Let $V_Q = V \cap Q$, $V'_Q = V' \cap Q$, $\mathcal{C}_Q \subseteq \mathcal{C}$ the set of cylinders that intersect Q, and $\Xi_Q = \Xi \cap \mathcal{C}_Q$. We say that a pair of cylindrical surfaces $a, b \in \mathcal{C}_Q$ is *good* in Q if the pair (a, b) is divergent and they intersect inside Q. For a cylinder $c \in \mathcal{C}_Q$, define $D(c)$ to be the set of cylinders $c' \in \mathcal{C}_Q$ such that (c, c') is a good pair in Q; define $D'(c) = D(c) \cap \Xi_Q$.

Put $q = \sqrt{n}$. Let

$$\mathcal{C}^1_Q = \{c \in \mathcal{C}_Q \mid |D(c)| \le q\} \quad \text{and} \quad \mathcal{C}^2_Q = \mathcal{C}_Q \setminus \mathcal{C}^1_Q.$$

Set $m_1 = |\mathcal{C}^1_Q|$ and $m_2 = |\mathcal{C}^2_Q|$. Let $\mu_1(Q)$ denote the number of good pairs in $\mathcal{C}^1_Q \times \Xi_Q$, and let $\mu_2(Q)$ denote the number of good pairs in $\mathcal{C}^2_Q \times \mathcal{C}_Q$. Lemma 2.3 implies that

$$\begin{aligned} \sum_{Q \in \mathcal{Q}} \mu_1(Q) &= \sum_{Q \in \mathcal{Q}} \sum_{c \in \mathcal{C}^1_Q} |D'(c)| \\ &\le \sum_{c \in \mathcal{C}, c' \in \Xi} O(1/t^2) = O(nn'). \end{aligned}$$

Similarly, we can prove that

$$\sum_{Q \in \mathcal{Q}} \mu_2(Q) = O(n^2).$$

Let v be a vertex in V'_Q, lying on three surfaces c_1, c_2, and c_3 (all belonging to $\mathcal{C}_Q$). We charge v to the cylinder c_i (for $i = 1, 2$, or 3) that has the smallest value of $|D(c_i)|$. Note that if v is charged to a surface in $\mathcal{C}^2_Q$, then all three surfaces c_i are in $\mathcal{C}^2_Q$. In any case, all three pairs of c_1, c_2, and c_3 are divergent and meet inside Q (at v), so all these pairs are good in Q. Moreover, by construction, at least two of the surfaces are in Ξ_Q. That is, if we fix one of the c_i's, say c_1, then both c_2 and c_3 are in $D(c_1)$ and at least one of them is in $D'(c_1)$.

We bound separately the number of vertices of V'_Q charged to each surface $c \in \mathcal{C}_Q$ and sum up over all surfaces in $\mathcal{C}_Q$.

Case (a): c is in $\mathcal{C}^1_Q$. By the above discussion, the number of vertices charged to c is at most proportional to $|D(c)| \cdot |D'(c)|$ (each triple of surfaces define only a constant number of vertices). Hence, the number of such vertices, over all cylinders in $\mathcal{C}^1_Q$, is at most

$$\sum_{c \in \mathcal{C}^1_Q} O(|D(c)| \cdot |D'(c)|) = O(q\mu_1(Q)).$$

Case (b): $c \in \mathcal{C}^2_Q$. If a vertex $v \in V'_Q$ is charged to a surface in $\mathcal{C}^2_Q$, then all three surfaces incident to v belong to $\mathcal{C}^2_Q$, by our charging scheme. Each such vertex lies on the boundary of the union of the regions bounded by the surfaces of $\mathcal{C}^2_Q$. Moreover, each vertex of V'_Q lies on two surfaces of Ξ_Q. Therefore, the total number vertices of V'_Q charged to any surface of $\mathcal{C}^2_Q$ is bounded by the number of those vertices in the union of $\mathcal{C}^2_Q$ that lie inside Q and that are incident on two surfaces of $\Xi_Q \cap \mathcal{C}^2_Q$. We will show below in Lemma 2.7 that the number of such vertices is $O(m_2|\mathcal{C}^2_Q \cap \Xi_Q|^{1+\varepsilon}) = O(m_2 n' n^\varepsilon)$, for any $\varepsilon > 0$. On the other hand, $\mu_2(Q) \ge m_2 q$, which implies that the number of vertices charged to cylinders in $\mathcal{C}^2_Q$ is at most $O(n' n^\varepsilon \mu_2(Q)/q)$.

Summing the two quantities, the number of vertices in V'_Q is at most

$$O\left(q\mu_1(Q) + \frac{n' n^\varepsilon \mu_2(Q)}{q}\right).$$

Summing this bound over all prisms $Q \in \mathcal{Q}$, we get, by the choice of q:

$$\begin{aligned} |V'| &= O\left(\sqrt{n} \sum_Q \mu_1(Q) + n^{-1/2+\varepsilon} n' \sum_Q \mu_2(Q)\right) \\ &= O(n^{3/2+\varepsilon} n'). \end{aligned}$$

This completes the proof of the lemma. □

Complexity of the union inside a prism. Next, we show that for any prism Q and for any collection $\mathcal{C}_Q$ of unit cylinders that intersect Q, we have $|V_Q| = O(|\mathcal{C}_Q|^{2+\varepsilon})$, for any $\varepsilon > 0$, where V_Q is the set of the vertices of the union of $\mathcal{C}_Q$ within Q. We assume that $\mathbf{u}$ is the $(+z)$-axis, i.e., the prisms are vertical. Let M be a sufficiently large constant, whose value will be chosen below. We partition each of the cylindrical surfaces in $\mathcal{C}$ into M canonical strips (parallel to the axis of the cylinder), each having an angular span of $2\pi/M$ (in the cylindrical coordinate frame induced by the cylinder). We say that a direction ρ is a *good direction* for a strip τ if the following two conditions hold.

(C1) $\angle(\rho, \mathbf{u}) \ge \pi/M$, and

(C2) each line tangent to (the relative interior of) τ forms an angle of at least π/M with ρ.

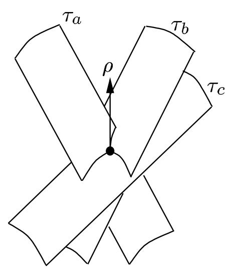

Figure 4: A vertex incident of the union on τ_a, τ_b, and τ_c along with a good direction.

We say that $\rho \in \mathbb{S}^2$ is a *good direction* for a vertex v incident to three canonical strips τ_a, τ_b, and τ_c if it is a good direction for all three strips; see Figure 4. It is easily checked that the set B_τ of bad directions for a fixed strip τ, contained in a cylindrical surface $c \in \mathcal{C}$ whose axis is assumed to be vertical, is the union $B_1 \cup B_2$, where we have:

- B_1 is the union of the two caps about the north and south poles of $\mathbb{S}^2$ at opening angles π/M.
- Let $\mathbf{n}_1$ and $\mathbf{n}_2$ be the normals to the planes passing through the axis of c and passing through the boundary of τ. By construction the angle between $\mathbf{n}_1$ and $\mathbf{n}_2$ is at most $2\pi/M$. The (thinner) "spherical double wedge" defined by the two great circles normal to $\mathbf{n}_1$ and $\mathbf{n}_2$ is the set of directions for which a line is tangent to τ. B_2 is the set of all points on $\mathbb{S}^2$ that lie at spherical distance at most π/M from this double wedge; see Figure 5. Thus B_2 is contained in a spherical band consisting of all points lying at spherical distance at most $2\pi/M$ from a great circle on $\mathbb{S}^2$ (namely, from the circle "bisecting" the double wedge); see Figure 5.

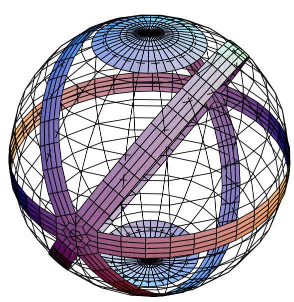

Figure 5: The set of bad directions for a vertex consists of two caps and three spherical bands.

It follows that the area of B_τ is at most

$$4\pi(1 - \cos\frac{\pi}{M} + \sin\frac{\pi}{M}).$$

This implies that the set of good directions for v has area at least

$$4\pi\left[1 - 3\left(1 - \cos\frac{\pi}{M} + \sin\frac{\pi}{M}\right)\right].$$

By choosing M sufficiently large, this area can be made close to the area of the entire sphere. Moreover, it is easy to verify that this set contains a spherical cap of opening angle $\delta \le \pi/8$, if M is sufficiently large.

Let Z be a set of $O(1/\delta^2)$ points on $\mathbb{S}^2$, with the property that any cap on $\mathbb{S}^2$ of opening angle δ contains at least one of these points. For each $\rho \in Z$ and a prism Q, we define $V_Q(\rho)$ to be the subset of all vertices in V_Q for which ρ is a good direction. The preceding analysis implies that each vertex of V has at least one good direction in Z.

Lemma 2.5 *Let $\rho \in Z$, and let v be any vertex in V_Q, incident to strips τ_a, τ_b, τ_c, for which ρ is a good direction. Then any line parallel to ρ intersects τ_a in at most one point. Moreover, if we go from any point $w \in \tau_a \cap Q$ into the cylinder K bounded by τ_a in the direction parallel to ρ, we reach ∂Q before exiting K. Similar properties hold for τ_b and τ_c.*

Proof: If τ_a were not monotone in the above sense, it would have to contain a point v so that a line parallel to ρ is tangent to τ_a at v, which is impossible by the definition of a good direction. As to the second assertion, let w be a point in $\tau_a \cap Q$, and let w' be the other intersection between ∂K and the line passing through w and parallel to ρ. It is easily verified that $|ww'|$ is minimized (relative to the constraints on good directions) when ww' is orthogonal to the axis of a and forms an angle π/M with the tangent plane to a at w. In this case $|ww'| = 2\sin\frac{\pi}{M}$. On the other hand, since ww' forms an angle of at least π/M with the z-direction, it follows that the horizontal distance between w and w' is at least $|ww'|\sin\frac{\pi}{M} \ge 2\sin^2\frac{\pi}{M}$. If t is chosen such that $t < \sqrt{2}\sin^2(\pi/M)$, then w' does not lie in Q, which completes the proof of the lemma. $\square$

For each prism Q, let $T_Q(\rho)$ denote the set of canonical strips τ such that τ crosses Q and τ contains at least one vertex in $V_Q(\rho)$. In particular, ρ is a good direction for τ. Let $n_Q(\rho) = |T_Q(\rho)|$.

Define the ρ-*upper* (resp. the ρ-*lower*) portion of a surface $c \in \mathcal{C}$ to be the portion of c consisting of points w with the property that if we move from w in the direction ρ we leave the region K bounded by c (resp. enter into K). Let τ be a strip in $T_Q(\rho)$, lying, say, on the ρ-upper portion of some surface c (by construction, τ cannot overlap both ρ-upper and ρ-lower portions of c). Lemma 2.5 implies that any line parallel to ρ that passes through a point in $\tau \cap Q$ meets the region K bounded by c in an interval whose other endpoint lies outside Q;

a symmetric property applies when τ lies on the ρ-lower portion of c. Let $T_Q^+(\rho)$ (resp. $T_Q^-(\rho)$) denote the set of surfaces $Q \cap \tau$, for all strips $\tau \in T_Q(\rho)$ that lie on the ρ-upper (resp. ρ-lower) portion of their containing surfaces.

Let v be a vertex in $V_Q(\rho)$. The preceding analysis implies that v is a vertex of the region R_Q enclosed between the ρ-upper envelope (i.e., the envelope in the positive ρ-direction) of the surfaces in $T_Q^+(\rho)$ and the ρ-lower envelope (the envelope in the negative ρ-direction) of the surfaces in $T_Q^-(\rho)$. By the result of [3], the number of such vertices is $O(n_Q(\rho)^{2+\varepsilon})$, for any $\varepsilon > 0$, with the constant of proportionality depending on ε. Repeating this step for all directions $\rho \in Z$, we obtain the following result.

Lemma 2.6 *Let Q be a prism, and let $\mathcal{C}_Q$ be a set of cylinders intersecting Q. Then the number of vertices of the union of $\mathcal{C}_Q$ lying inside Q is $O(|\mathcal{C}_Q|^{2+\varepsilon})$, for any $\varepsilon > 0$.*

Finally, let Ξ_Q be a subset of $\mathcal{C}_Q$. We wish to bound the number of those vertices of V_Q that are incident upon (at least) two surfaces of Ξ_Q. Set $\xi = |\Xi_Q|$ and $m = |\mathcal{C}_Q|$. Divide $\mathcal{C}_Q$ into $\beta = \lceil m/\xi \rceil$ subsets $\mathcal{C}_1, \ldots, \mathcal{C}_\beta$, each of size at most ξ. Each vertex of V_Q that is incident upon two surfaces of Ξ_Q is a vertex in the union of $\Xi_Q \cup \mathcal{C}_i$, for some $1 \le i \le \beta$. By Lemma 2.6, the number vertices in the union of $\Xi_Q \cup \mathcal{C}_i$ is $O(\xi^{2+\varepsilon})$. Hence, the total number of such vertices is $O((m/\xi) \cdot \xi^{2+\varepsilon}) = O(m\xi^{1+\varepsilon})$.

Lemma 2.7 *Let Q be a prism, let $\mathcal{C}_Q$ be a set of cylinders intersecting Q, and let Ξ_Q be a subset of $\mathcal{C}_Q$. Then the number of vertices of the union of $\mathcal{C}_Q$ that lie inside Q and that are incident upon at least two surfaces of Ξ_Q is $O(|\mathcal{C}_Q| \cdot |\Xi_Q|^{1+\varepsilon})$, for any $\varepsilon > 0$.*

Putting everything together, we thus conclude the following.

Theorem 2.8 *The complexity of the union of n congruent cylinders in $\mathbb{R}^3$ is $O(n^{5/2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε.*

3 The Complexity of U: The Case of Segments

We next extend Theorem 2.8 to the case of segments. Let $S = \{s_1, \ldots, s_n\}$ be a set of n segments in $\mathbb{R}^3$. Again, for each i, let $K_i = K_{s_i} = s_i \oplus B$. If s_i is a bounded segment, then ∂K_i consists of a bounded cylindrical surface, denoted by c_i, and two hemispheres σ_i^+, σ_i^- attached at the ends of c_i. Without loss of generality, we assume that σ_i^+, σ_i^- are full spheres, bounding two respective balls which are fully contained in K_i. Let $\mathcal{K} = \{K_1, \ldots, K_n\}$, $\mathcal{C} = \{c_1, \ldots, c_n\}$, and $\Sigma = \{\sigma_1^+, \sigma_1^-, \ldots, \sigma_n^+, \sigma_n^-\}$. Let V denote the set of vertices on the union of $\mathcal{K}$. Each vertex of V is an intersection of three cylindrical surfaces, of two cylindrical surfaces and one spherical surface, of a cylindrical surface and two spherical surfaces, or of three spherical surfaces. We will argue that a slight modification of the proof presented in the previous section can be extended to this case.

The intuition behind the extension is as follows. In the preceding argument, the complexity of the union of a set $\mathcal{C}_Q$ of cylinders that cross a fixed prism Q is $O(|\mathcal{C}_Q|^{2+\varepsilon})$ (Lemma 2.6). The main technical difficulty in the preceding proof is that we cannot bound $\sum_Q |\mathcal{C}_Q|^{2+\varepsilon}$, because each cylinder crosses an infinite number of prisms. We overcame this difficulty by (a) considering *pairs* of cylinders, and (b) managing to count only divergent pairs, whose intersection curve crosses only $O(1)$ prisms. Now making the cylinders bounded and adding spheres can only help us in this aspect of the proof, and the new proof simply follows the old one, with additional (and simpler) care given to the spherical surfaces. We only sketch the extended analysis.

(i) As in Lemma 2.1, we choose the z-axis and a subset $\mathcal{C}'$ of $\mathcal{C}$ so that at least half of the vertices of V are also vertices of the union of $\mathcal{C}' \cup \Sigma$, and so that the acute angles between the z-axis and the axes of the cylinders in $\mathcal{C}'$ are all at most β_0. Abusing the notation a little, as above, we set $\mathcal{C}$ to $\mathcal{C}'$. Note that even if we delete a cylindrical surface c_i, we retain the corresponding spheres σ_i^+, σ_i^- in the set. For each (hemi)sphere $\sigma \in \Sigma$, we assign an arbitrary unit vector $\mathbf{n}_\sigma \in \mathbb{S}^2$ so that the angle between the z-axis and n_σ is at most β_0. Let D be the set of points corresponding to the axis directions of cylinders in $\mathcal{C}$ and the vectors assigned to Σ.

(ii) We call every (unordered) pair in $\Sigma \times (\mathcal{C} \cup \Sigma)$ a divergent pair; the definition of divergent pairs for two cylindrical surfaces remains the same.

(iii) With these modifications, Lemma 2.2 still holds: There exist subsets $\Xi \subseteq \mathcal{C} \cup \Sigma$ and $V' \subseteq V$ so that

$$|V'|/|\Xi| \ge |V|/(3cn \log^2(3n))$$

and properties (ii) and (iii) of Lemma 2.2 continue to hold.

(iv) The intersection curve of any pair of surfaces in $\Sigma \times (\mathcal{C} \cup \Sigma)$ meets only $O(1)$ prisms, because any sphere in Σ intersects only $O(1)$ prisms. The proof of Lemma 2.3 continues to hold for pairs of cylindrical surfaces, and the lemma therefore holds for any pair of surfaces.

(v) Now Lemma 2.4 holds in the extended case as well, where the proof follows the preceding one almost verbatim.

(vi) Finally, concerning Lemmas 2.5 and 2.6, we divide each (hemi)sphere of Σ into $O(M^2)$ spherical caps, each of angular opening at most $2\pi/M$. We define a good direction for a spherical cap in the same manner as we did for a strip (see (C1) and (C2) above). The set of bad directions for such a spherical cap τ is again the union of $B_1 \cup B_2$, where B_1 is the same as earlier, and B_2 is the set of all points on $\mathbb{S}^2$ that lie at spherical distance at most π/M from the spherical band consisting of all points at spherical distance at most π/M from the great circle that is the locus of all directions orthogonal to the surface normal at the center of τ. In other words, B_2 is itself a spherical band consisting of all points at spherical distance at most $2\pi/M$ from that circle. Hence, we can again choose a set Z of $O(1)$ directions so that at least one direction in Z is good for every vertex of V_Q. It is now easy to check that both Lemmas 2.5 and 2.6 continue to hold in the extended case.

Putting everything together, we thus obtain the following extension.

Theorem 3.1 *The complexity of the union of the Minkowski sums of a ball and n line segments in $\mathbb{R}^3$ is $O(n^{5/2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε.*

4 Efficient Construction of the Union

In this section we sketch a randomized algorithm for computing the boundary of U in expected time $O(n^{5/2+\varepsilon})$. The algorithm is similar to the ones described in [1, 2]; see also [9, 17]. Our algorithm can also preprocess U in time $O(n^{5/2+\varepsilon})$ into a data structure of the same size, as in [2], so that a path between two query points in $\mathbb{R}^3 \setminus U$ that does not intersect the interior of U can be computed in $O(\log n + k)$ time, where k is the number of "edges" in the path.

Let $\mathcal{K} = \{K_1, \ldots, K_n\}$ be the set of expanded obstacles. Define $\mathcal{C}, \Sigma$ as in the previous section. For each surface $\sigma \in \mathcal{C} \cup \Sigma$, our algorithm computes $\partial U \cap \sigma$. We describe how to compute $\partial U \cap c_i$ for a cylindrical surface $c_i \in \mathcal{C}$; spherical surfaces can be handled in a similar manner. Without loss of generality, assume that the segment s_i is parallel to the z-axis.

So we now fix c_i. For each $j \neq i$, let $\gamma_j = K_j \cap c_i$; $\gamma_j \subseteq c_i$ is a two-dimensional surface patch. Our goal is to compute $\Pi = \bigcup_{j \neq i} \gamma_j$. Set $\Gamma^i = \{\gamma_j \mid j \neq i\}$. We compute Π by adding the patches γ_j one by one in a random order and by computing the union of the

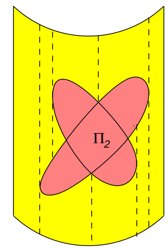

Figure 6: Vertical decomposition $\overline{\Pi}_j^{||}$ of the complement of the union of two patches.

patches added so far. Let $\langle \gamma_1, \gamma_2, \ldots \rangle$ be the (random) insertion sequence of patches, and let $\Pi_j = \bigcup_{k=1}^{j} \gamma_k$.

The algorithm maintains the vertical decomposition of $\overline{\Pi}_j = c_i \setminus \Pi_j$. The *vertical decomposition* $\overline{\Pi}_j^{||}$ is obtained by drawing segments in the $(+z)$- and $(-z)$-directions from every vertex of $\partial \Pi_j$ and every point of vertical tangency of $\partial \Pi_j$ (i.e., a point on $\partial \Pi_j$ at which the line parallel to the z-axis is tangent to Π_j), within $\overline{\Pi}_j$, until they hit $\partial \Pi_j$; see Figure 6. If any of the segments does not intersect $\partial \Pi_j$, it is extended to infinity (if c_i is unbounded) or to the appropriate circular base of c_i (if it is bounded). Each face of $\overline{\Pi}_j^{||}$ is a *pseudo-trapezoid* in the sense that it is bounded by at most four arcs, of which two are vertical segments and the other two are portions of edges of Π_j. Each pseudo-trapezoid is defined by at most four patches of Γ^i. Conversely, any four patches of Γ^i define a constant number of pseudo-trapezoids, namely, those which appear in the vertical decomposition of the complement of the union of these four patches. The algorithm actually maintains $\overline{\Pi}_j^{||}$. In the j-th step the algorithm inserts γ_j, finds all pseudo-trapezoids of $\overline{\Pi}_{j-1}^{||}$ that intersect γ_j, and modifies these trapezoids to compute $\Pi_j^{||}$, using the technique described in [1, 10]. We repeat this process for each surface in $\mathcal{C} \cup \Sigma$.

The analysis of the expected running time of the algorithm proceeds along the same lines as described in [1, 10, 17]. We define the weight, $w(\tau)$, of a pseudo-trapezoid τ, defined by the patches of Γ^i, to be the number of patches in Γ^i that intersect τ. As shown in the cited papers, the running time of the algorithm is proportional to the number of pseudo-trapezoids created by the algorithm, plus the sum of their weights, plus an $O(n^2)$ overhead term needed to construct Γ^i over all i. We omit from this abstract the details of estimating the expected values of these sums, and conclude with the following result.

Theorem 4.1 *Given a set S of n line-segments in $\mathbb{R}^3$, and a ball B, the boundary of the free configuration space of B amid S, namely the boundary of the union of the*

Minkowski sums $\{s \oplus B \mid s \in S\}$, can be computed in randomized expected $O(n^{5/2+\varepsilon})$ time, for any $\varepsilon > 0$.

5 Union of Objects with Bounded Curvature

Inspecting the preceding proofs, we see that they can be extended to several other interesting cases. Here is one such extension: Let $\mathcal{K}$ be a collection of n compact convex objects in $\mathbb{R}^3$ satisfying the following properties:

(i) The objects in $\mathcal{K}$ have *constant description complexity*, meaning that each object is a semialgebraic set defined by a constant number of polynomial equalities and inequalities of constant maximum degree.

(ii) The objects in $\mathcal{K}$ are all of roughly the same size, meaning that the ratio between the diameters of any pair of objects is at most some fixed constant α.

(iii) The objects in $\mathcal{K}$ are all smooth and the curvature of any $K \in \mathcal{K}$ at any point is at most some fixed contant κ_0.

In this case we have the following;

Theorem 5.1 *The complexity of the union of a collection $\mathcal{K}$ as above is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε, α, κ_0, and on the maximum algebraic complexity of an object in $\mathcal{K}$.*

Due to lack of space we omit the proof. Briefly, we partition $\mathbb{R}^3$ into a grid of cubes (rather than prisms) and show that the complexity of the union of $\mathcal{K}$ within a cube is near-quadratic in the number of objects intersecting the cube (the bounded curvature assumption allows us to reduce this problem to the study of a region enclosed between two envelopes, as above). Summing over all cubes, and using the "same size" assumption, the bound follows easily.

6 Conclusion

In this paper we proved the first nontrivial bound on the complexity of the free configuration space $\mathcal{F}$ of a set of pairwise-disjoint segments with respect to a ball in $\mathbb{R}^3$. We also presented an efficient randomized algorithm for computing the boundary of the configuration space. We are currently working on proving a near-quadratic bound on $\mathcal{F}$ and on extending the results to a set of pairwise-disjoint polyhedral obstacles in $\mathbb{R}^3$. We believe there is a good chance of achieving these goals. As mentioned in the introduction, the main open problem in this area is to prove a tight bound on the complexity of the Euclidean Voronoi diagram of a set of pairwise-disjoint polyhedral obstacles in $\mathbb{R}^3$.

Acknowledgments

The authors wish to thank Boris Aronov, Alon Efrat, Leo Guibas, Sariel Har-Peled, and János Pach for helpful discussions concerning this problem.

References

[1] P. K. Agarwal, B. Aronov, and M. Sharir, Computing envelopes in four dimensions with applications, *SIAM J. Comput.* 26 (1997), 1714–1732.

[2] P. K. Agarwal, B. Aronov, and M. Sharir, Motion planning for a convex polygon in a polygonal environment, to appear in *Discrete Comput. Geom.*.

[3] P.K. Agarwal, O. Schwarzkopf and M. Sharir, The overlay of lower envelopes in 3-space and its applications, *Discrete Comput. Geom.* 15 (1996), 1–13.

[4] P.K. Agarwal and M. Sharir, Arrangements and their applications, in *Handbook of Computational Geometry* (J. Sack, ed.), North-Holland, to appear.

[5] B. Aronov, B. Chazelle, H. Edelsbrunner, L. Guibas, M. Sharir and R. Wenger, Points and triangles in the plane and halving planes in space, *Discrete Comput. Geom.* 6 (1991), 435–442.

[6] B. Aronov and M. Sharir, On translational motion planning of a convex polyhedron in 3-space, *SIAM J. Comput.* 26 (1997), 1785–1803.

[7] B. Aronov, M. Sharir, and B. Tagansky, The union of convex polyhedra in three dimensions, *SIAM J. Comput.* 26 (1997), 1670–1688.

[8] J.D. Boissonnat, M. Sharir, B. Tagansky and M. Yvinec, Voronoi diagrams in higher dimensions under certain polyhedral distance functions, *Discrete Comput. Geom.* 19 (1998), 485–519.

[9] B. Chazelle, H. Edelsbrunner, L. Guibas and M. Sharir, Selecting heavily covered points, *SIAM J. Comput.* 23 (1994), 1138–1151.

[10] B. Chazelle, H. Edelsbrunner, L. J. Guibas, M. Sharir, and J. Snoeyink, Computing a face in an arrangement of line segments and related problems, *SIAM J. Comput.* 22 (1993), 1286–1302.

[11] L.P. Chew, K. Kedem, M. Sharir, B. Tagansky and E. Welzl, Voronoi diagrams of lines in three dimensions under polyhedral convex distance functions, *J. Algorithms* 29 (1998), 238–255.

[12] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.

[13] A. Efrat and M. Sharir, On the complexity of the union of fat objects in the plane, *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, 1997, pp. 104–112.

[14] K. Kedem, R. Livne, J. Pach and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.* 1 (1986), 59–71.

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

[16] M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* 12 (1994), 327–345.

[17] M. Sharir and P. K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

[18] B. Tagansky, *The Complexity of Substructures in Arrangements of Surfaces*, Ph.D. Dissertation, Computer Science Department, Tel Aviv University, 1996.