

Pipes, Cigars, and Kreplach: the Union of Minkowski Sums in Three Dimensions*

P. K. Agarwal¹ and M. Sharir²

¹Center for Geometric Computing, Department of Computer Science,
Box 90129, Duke University,
Durham, NC 27708-0129, USA
pankaj@cs.duke.edu

²School of Mathematical Sciences, Tel Aviv University,
Tel Aviv 69978, Israel
sharir@math.tau.ac.il
and
Courant Institute of Mathematical Sciences, New York University,
New York, NY 10012, USA

Abstract. Let Ω be a set of pairwise-disjoint polyhedral obstacles in \mathbb{R}^3 with a total of n vertices, and let B be a ball in \mathbb{R}^3 . We show that the combinatorial complexity of the free configuration space \mathcal{F} of B amid Ω , i.e., (the closure of) the set of all placements of B at which B does not intersect any obstacle, is $O(n^{2+\epsilon})$, for any $\epsilon > 0$; the constant of proportionality depends on ϵ . This upper bound almost matches the known quadratic lower bound on the maximum possible complexity of \mathcal{F} . The special case in which Ω is a set of lines is studied separately. We also present a few extensions of this result, including a randomized algorithm for computing the boundary of \mathcal{F} whose expected running time is $O(n^{2+\epsilon})$, for any $\epsilon > 0$.

* Work by P.A. was supported by Army Research Office MURI Grant DAAH04-96-1-0013, by a Sloan fellowship, by NSF Grants EIA-9870724, EIA-997287, and CCR-9732787, 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. A preliminary version of this paper appeared in *Proc. 15th Annual ACM Symposium on Computational Geometry*, 1999, pp. 143–153.

1. Introduction

Problem Statement. Let Ω be a collection of pairwise-disjoint polyhedral obstacles in \mathbb{R}^3 with a total of n vertices, edges, and faces, and let B be a ball in \mathbb{R}^3 . With no loss of generality, we assume that the faces of Ω are triangles and that the radius of B is 1. We consider the motion-planning problem in which B is allowed to move (translate) freely in \mathbb{R}^3 without intersecting any obstacle. The *free configuration space* \mathcal{F} of B with respect to Ω is the space of all points $p \in \mathbb{R}^3$ so that if B is placed centered at p , then it does not intersect any obstacle. 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. \mathcal{F} can be expressed in the following standard manner (see, e.g., [20]). For each obstacle $\omega \in \Omega$, let K_ω denote the *Minkowski sum*¹

$$K_\omega = \omega \oplus B_0 = \{x + y \mid x \in \omega, y \in B_0\}.$$

The set K_ω , referred to as the *expanded obstacle* of ω , is the set of all centers of B at placements where it intersects ω . Hence $\mathcal{F} = \mathbb{R}^3 \setminus \bigcup_{\omega \in \Omega} K_\omega$. See Fig. 1.

Let S be the set of triangular faces of Ω . For each triangle $s \in S$, we can define $K_s = s \oplus B_0$. If s is bounded, then K_s is the (nondisjoint) union of (i) a triangular prism of height 2 with s as a middle cross section, (ii) three bounded cylinders of radius 1 whose axes are the edges of s , and (iii) three balls of radius 1 centered at the vertices of s . If s is unbounded, the structure of K_s changes accordingly. We refer to K_s as an *expanded triangle* (or a *krepl*).

A face of \mathcal{F} is a maximal connected closed portion of $\partial\mathcal{F}$ contained in a single triangular, cylindrical, or spherical portion of some ∂K_s . An edge of \mathcal{F} is a maximal




Fig. 1. The union of Minkowski sums of triangles and a ball.

¹ Strictly speaking, we should form the sum of ω with $-B_0$, the reflection of B_0 through the origin; of course, we obtain the same set since B_0 is symmetric.

connected portion of $\partial\mathcal{F}$ lying in the intersection of two distinct faces; the two faces may lie on the boundary of the same krepl or on the boundaries of different kreplach. A vertex of \mathcal{F} is the intersection of three distinct faces, not necessarily of distinct kreplach, that lies in $\partial\mathcal{F}$. The *combinatorial complexity* of \mathcal{F} , denoted by $|\mathcal{F}|$, is the number of vertices, edges, and two-dimensional faces of $\partial\mathcal{F}$.

Set $U = \bigcup_{s \in S} K_s$. Each connected component of \mathcal{F} is also a connected component of $\mathbb{R}^3 \setminus U$, but the latter may have some connected components that do not belong to \mathcal{F} . These components represent placements at which the ball moves inside an obstacle without touching its boundary. $|\mathcal{F}|$ is thus upper bounded by $|U|$, and \mathcal{F} can be computed by first constructing U and then discarding the connected components of $\mathbb{R}^3 \setminus U$ that do not belong to \mathcal{F} . The main problems we are concerned with are thus to estimate the combinatorial complexity of U and to compute efficiently its boundary ∂U .

Besides this motion-planning application, the problem of bounding the complexity of U is a precursor to the harder problem of obtaining a near-quadratic, or even just 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 triangle 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 [24] and the survey paper [5] by the authors for a summary of known results on this topic. Boissonnat et al. [12] proved that the maximum complexity of the union of n axis-parallel hypercubes in \mathbb{R}^d is $\Theta(n^{\lfloor d/2 \rfloor})$; the bound improves to $\Theta(n^{\lfloor d/2 \rfloor})$ if all hypercubes have the same size. Aronov et al. [9] proved that the complexity of the union of n convex polyhedra in \mathbb{R}^3 with a total of s faces is $O(n^3 + ns \log n)$. Aronov and Sharir [8] proved that the complexity of the union of the Minkowski sums of a convex polyhedron P with a collection S of n pairwise-disjoint convex polyhedra in \mathbb{R}^3 is $O(ns \log n)$, where s is the total number of faces of the polyhedra 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 the work on unions of objects in the plane. Among the two-dimensional results, we mention the early result of Kedem et al. [20] that shows that the complexity of the union of n disks (or “pseudodisks”) is $O(n)$, and the results of Matoušek et al. [22] and Efrat and Sharir [17] that prove near-linear bounds on the complexity of the union of “fat” triangles and general “fat” convex regions in the plane. See also [7], [16], and [21]. In a sense, our results are extensions of the analysis of [20] to three dimensions.

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 on lower envelopes [23]. The maximum complexity of Voronoi diagrams of n point sites under the Euclidean distance is known to be $\Theta(n^2)$ [15]. The same bound has recently been established for point sites under the L_1 and L_∞ metrics, or under any simplicial distance function [12]. Near-quadratic bounds have also been recently established for

the case of line sites and any polyhedral convex distance function [13], where the bound is $O(n^2\alpha(n)\log n)$, and for the case of point sites and any polyhedral convex distance function [26], 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 superquadratic 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 with 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 triangle sites and Euclidean distance, then the complexity of U will be near-quadratic. Although a subcubic bound on the complexity of the Voronoi diagram still remains elusive, we prove that the complexity of U is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Using this bound, we also derive a near-quadratic algorithm for constructing the complement of the union U , and thereby obtain a motion-planning algorithm for a ball amid polyhedral obstacles.

Our results extend and improve a previous initial attack on the problem by the authors [4], where we only managed to handle the cases in which S is a collection of lines or segments and to obtain a weaker bound of $O(n^{5/2+\varepsilon})$, for any $\varepsilon > 0$. The new analysis borrows ideas from the previous paper, but has many new ingredients.

The paper is organized as follows. In Section 2 we study the special case in which S is a set of lines, so U is the union of congruent cylinders (*pipes*). We extend the previous result to segments in Section 3; here U is the union of *cigars*. In Section 4 we prove the main result of the paper—a near-quadratic bound on the complexity of U for the case of pairwise-disjoint triangles, so U is the union of *kreplach*. In Section 5 we discuss two generalizations of our results. The first result proves a near-quadratic bound on the complexity of the union of convex objects of bounded curvature and of roughly the same size. The second result proves a near-quadratic bound on the number of changes in the combinatorial structure of a set of congruent disks in the plane, each moving with a fixed velocity. We also present a near-quadratic algorithm for constructing U and \mathcal{F} . We conclude the paper in Section 6 with a few open problems.

2. The Case of Pipes

Preliminaries and Overview. We first solve the problem, in which $S = \{s_1, \dots, s_n\}$ is a set of n lines in \mathbb{R}^3 . For $i = 1, \dots, n$, let $K_i = K_{s_i} = s_i \oplus B$ and $c_i = \partial K_i$; K_i is an infinite cylinder (or *pipe*) of radius 1. Set $\mathcal{K} = \{K_1, \dots, K_n\}$, $U = \bigcup_{i=1}^n K_i$, and $\mathcal{C} = \{c_1, \dots, c_n\}$. Whenever there is no ambiguity, we will also refer to the c_i 's as cylinders; otherwise we will refer to them as cylindrical surfaces. See Fig. 2. We assume that the lines in S are in general position, which means that every pair of lines is skew, 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. An argument based on random perturbation, similar to the one given in [23], shows that this assumption can be made with no loss of generality. The main result of this section is the following.




Fig. 2. Two arrangements of cylinders.

Theorem 2.1. *The combinatorial complexity of the union of n congruent cylinders in \mathbb{R}^3 is $O(n^{2+\epsilon})$, for any $\epsilon > 0$, where the constant of proportionality depends on ϵ .*

For a subset $R \subseteq \mathcal{C}$, let $U(R)$ denote the union of cylinders bounded by the cylindrical surfaces in R . Let $V(R)$ denote the set of *vertices* of $U(R)$, namely, intersection points of triples of boundaries of cylindrical surfaces in R that lie on $\partial U(R)$. By our general position assumption, each vertex lies on exactly three cylindrical surfaces, and is thus incident upon only a constant number of edges and faces. The number of edges or 2-faces of ∂U that are not incident upon any vertex is $O(n^2)$. Therefore the combinatorial complexity of U is $O(n^2 + |V(\mathcal{C})|)$. In the rest of this section we prove the following:

Proposition 2.2. *For any set \mathcal{C} of n congruent cylinders in \mathbb{R}^3 and for any $\epsilon > 0$,*

$$|V(\mathcal{C})| = O(n^{2+\epsilon}).$$

Overview of the Proof. The proof consists of several main steps, each presented in a separate subsection, and proceeds through a sequence of technical lemmas. To aid the reader in following the proof, we have written it from a certain point on in a “backward” manner: each step relies on a future key lemma and shows how Proposition 2.2 follows from the analysis so far and from that future lemma.

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 orientation of 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 $\cos^{-1}(\frac{1}{6})$.

In the second step we derive a recurrence relation to bound the number of vertices. The overhead term in the recurrence counts the number of vertices lying on cylinders whose axis directions are “well separated” in a certain sense.

In order to bound the overhead term, the third step introduces a key notion of “divergent pairs” of cylinders, relative to some direction \mathbf{u} , where the angle between the axes of such a pair is not much smaller than the angles that the axes form with \mathbf{u} . We show the existence of a direction \mathbf{u} so that many vertices $v \in V$ have the following property: all three pairs of the cylinders that are incident upon v are divergent with respect to \mathbf{u} .

In the fourth step we partition \mathbb{R}^3 into a carefully chosen infinite grid of square prisms whose infinite axes are in the direction \mathbf{u} , and count the number of vertices within each prism. We show that there are only $O(1)$ prisms Q that can be crossed by a fixed pair a, b of divergent cylinders so that the projections of $a \cap Q$ and $b \cap Q$ on a line in direction \mathbf{u} overlap.

In the fifth step we show that, within a prism Q , we can bound the number of vertices of U by regarding them as vertices of a “sandwich” region enclosed between an upper envelope of a collection of portions of the given cylinders and a lower envelope of another such collection. Using the results of [3] on the complexity of such a sandwich region, we get a near-quadratic bound for the number of vertices of U within a prism. We interpret this bound as counting the number of pairs of cylinders that cross the same prism.

Finally, in the sixth step, we sharpen the bound obtained in the fifth step so that it is proportional to the number of pairs of divergent cylinders that have “nearby” crossings with Q , in the sense of step 4. Hence, when we sum these improved bounds over all prisms we still get an overall near-quadratic bound. This is accomplished (a) by improving the bound of [3], and (b) by using a divide-and-conquer method that effectively decomposes a prism into a tree of boxes and counts the number of vertices within each box separately.

We now describe each step in detail.

2.1. Choosing the z -Direction

Let \mathbb{S}^2 denote the unit sphere of directions in \mathbb{R}^3 . For each $c \in \mathcal{C}$, let $\mathbf{n}_c \in \mathbb{S}^2$ denote a unit vector in the direction of the axis of c that points into the upper halfspace; if the axis of c is parallel to the xy -plane, we set \mathbf{n}_c to be any of the two unit vectors in the direction of the axis of c . There is a technical problem (e.g., in Lemma 2.7 below) with the definition of the directions \mathbf{n}_c , for $c \in \mathcal{C}$, which depend on the choice of the z -direction. Informally, we may have a pair a, b of cylinders whose directions $\mathbf{n}_a, \mathbf{n}_b$ are almost antipodal. In the foregoing analysis we treat this pair as having a large angle (close to π) between their axes, whereas the “real” angle between the axes is close to zero. We circumvent this problem by choosing a random point on \mathbb{S}^2 and by regarding it as the direction of the $(+z)$ -axis. The following claim holds.

Lemma 2.3. *Let β_0 be the acute angle satisfying $\cos \beta_0 = \frac{1}{6}$. Let v be a vertex in $V(\mathcal{C})$ incident upon three cylinders $a, b, c \in \mathcal{C}$. The probability that all three acute angles between the z -direction and the axes of a, b, c are at most β_0 is at least $\frac{1}{2}$.*

Proof. Indeed, for the acute angle between the z -axis and the axis of, say, a to be greater than β_0 , the z -direction has to lie in the spherical band consisting of all directions at spherical distance at most $\pi/2 - \beta_0$ 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 larger than β_0 is at most

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




Fig. 3. (a) A set of relevant directions, and (b) their projection on h .

We thus obtain the following:

Lemma 2.4. *We can choose a subset $C' \subseteq C$ and a direction $\mathbf{u}_0 \in \mathbb{S}^2$ so that the axes of all cylinders in C' form acute angles at most $\beta_0 = \cos^{-1}(\frac{1}{8})$ with \mathbf{u}_0 and $|V(C')| \geq |V(C)|/2$.*

We rotate the coordinate system so that \mathbf{u}_0 becomes the $(+z)$ -axis and remove from C all the cylinders whose axes have an acute angle larger than β_0 with the $(+z)$ -axis. At least half of the vertices of $V(C)$ still show up in the new union. Abusing the notation slightly, we use C to denote the set of remaining cylinders.

Let S be the spherical cap consisting of all points in \mathbb{S}^2 that form an angle of at most β_0 with the $(+z)$ -axis; see Fig. 3(a). We project S onto the horizontal plane $h : z = 1$ using the central projection. The resulting projection is a disk \mathcal{D} of radius $\tan \beta_0 = \sqrt{35}$ centered at $(0, 0, 1)$. For a point $\mathbf{u} \in \mathbb{S}^2$, we denote its projection on h by \mathbf{u}^* . For a cylinder a , we refer to \mathbf{n}_a^* as its *direction image*. For a vertex $v \in V(C)$, incident on three cylinders $a, b, c \in C$, we associate with v the triple $\Delta_v = \{\mathbf{n}_a^*, \mathbf{n}_b^*, \mathbf{n}_c^*\}$. Δ_v is referred to as the set of direction images of v .

2.2. Deriving the Main Recurrence

Let $\psi(n) = \max |V(C)|$, where the maximum is taken over all sets C of n cylinders of radius 1 whose axes make acute angles of at most β_0 with the z -axis. Fix a constant integer parameter $\xi > 2$ whose value depends on ε and will be specified later. Partition the plane h into a collection $\mathbf{W} = \{W_1, \dots, W_\xi\}$ of ξ horizontal strips by lines parallel to the x -axis, so that each strip contains direction images of at most n/ξ cylinders. See Fig. 3(b). For each pair of strips $W_i, W_j \in \mathbf{W}$, let C_{ij} denote the set of cylinders whose direction images lie in $W_i \cup W_j$. By definition,

$$|V(C_{ij})| \leq \psi\left(\frac{2n}{\xi}\right) \quad \text{for } 1 \leq i < j \leq \xi.$$

Next, we partition the plane into a collection $\mathbf{H} = \{H_1, \dots, H_\xi\}$ of ξ vertical strips by lines parallel to the y -axis, so that each strip contains at most n/ξ direction images. For each pair $H_k, H_l \in \mathbf{H}$, we also bound by $\psi(2n/\xi)$ the number of vertices in the union of cylinders whose direction images lie in $H_k \cup H_l$. These $2\binom{\xi}{2} < \xi^2$ subproblems

have accounted for all those vertices v of $V(\mathcal{C})$ whose direction images Δ_v lie in at most two horizontal or two vertical strips, and the number of these vertices is at most $\xi^2 \psi(2n/\xi)$. We thus have to count the number of vertices for which Δ_v lies in three different horizontal strips and in three different vertical strips.

The strips in \mathbf{H} and \mathbf{W} divide the plane h into a set $\mathcal{R} = \{R_1, \dots, R_{\xi^2}\}$ of ξ^2 rectangles. For a rectangle $R_i \in \mathcal{R}$ let \mathcal{C}_i be the set of cylinders whose direction images lie in R_i . For a triple i, j, k , let $V_{i,j,k} = V(\mathcal{C}_i, \mathcal{C}_j, \mathcal{C}_k) \subseteq V(\mathcal{C})$ denote the set of vertices v of $U(\mathcal{C}_i \cup \mathcal{C}_j \cup \mathcal{C}_k)$ lying on three cylinders a, b, c such that $a \in \mathcal{C}_i, b \in \mathcal{C}_j$, and $c \in \mathcal{C}_k$. Then $\psi(\mathcal{C}) = \sum_{i \leq j \leq k} |V_{i,j,k}|$. In view of the preceding discussion, it suffices to bound $V_{i,j,k}$ for each triple $i \neq j \neq k$ for which the rectangles R_i, R_j, R_k lie in different horizontal and vertical strips, i.e., their x - and y -projections are pairwise disjoint. We show below in Lemma 2.5 that for such a triple of rectangles $|V_{i,j,k}| = O(|\mathcal{C}_i \cup \mathcal{C}_j \cup \mathcal{C}_k|^2 \cdot n^\epsilon) = O(n^{2+\epsilon}/\xi^2)$, for any $\epsilon > 0$. Since there are $O(\xi^6)$ such triples of rectangles, we obtain the following recurrence:

$$\psi(n) \leq \xi^2 \psi\left(\frac{2n}{\xi}\right) + O(n^{2+\epsilon} \xi^4).$$

For any $\epsilon' > \epsilon$, by choosing $\xi = \xi(\epsilon')$ a sufficiently large constant, one can prove that the solution to the above recurrence is $\psi(n) = O(n^{2+\epsilon'})$ (see, e.g., [23]), thereby proving Theorem 2.1.

2.3. Bounding $|V_{1,2,3}|$ and Divergent Pairs

Let R_1, R_2, R_3 be three rectangles in \mathcal{R} whose x - and y -projections are pairwise disjoint; see Fig. 4. Let $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$ be the corresponding subsets of cylinders, as above. We want to bound the size of $V_{1,2,3} = V(\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3)$. We will prove the following:

Lemma 2.5. *Let R_1, R_2, R_3 be three rectangles as defined above, and let $\epsilon > 0$ be an arbitrarily small constant. Then*

$$|V_{1,2,3}| = O(|\mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3|^2 \cdot n^\epsilon).$$

Definition 2.6. We call a pair of cylinders $a, b \in \mathcal{C}$ λ -divergent with respect to a direction \mathbf{u} (assumed to lie in S) if

$$\min\{|\mathbf{n}_a^* \mathbf{u}^*|, |\mathbf{n}_b^* \mathbf{u}^*|\} \leq \lambda |\mathbf{n}_a^* \mathbf{n}_b^*|.$$

Roughly speaking, two cylinders a and b being divergent with respect to a direction \mathbf{u} means that the slopes of the projections of the axes of a and b on a plane normal to \mathbf{u} are not “very close” to each other. The significance of divergent pairs is illustrated in Lemma 2.8.

Lemma 2.7. *There exist a direction \mathbf{u} and three pairwise-disjoint subsets $\mathcal{C}'_1 \subseteq \mathcal{C}_1, \mathcal{C}'_2 \subseteq \mathcal{C}_2, \mathcal{C}'_3 \subseteq \mathcal{C}_3$ so that*

(i) $|V(\mathcal{C}'_1, \mathcal{C}'_2, \mathcal{C}'_3)| \geq |V_{1,2,3}|/2$, and




Fig. 4. The two cases in the proof of Lemma 2.7: (a) the R_i 's form a monotone sequence; (b) the R_i 's do not form a monotone sequence.

(ii) all pairs of cylinders in $C'_1 \times C'_2$, $C'_1 \times C'_3$, and $C'_2 \times C'_3$ are $\sqrt{17}$ -divergent relative to \mathbf{u} .

Proof. For $i = 1, 2$, we assume that the x -projection of R_i lies to the left of the x -projection of R_{i+1} . We say that a point p separates R_i and R_j if its x -coordinate separates the x -projections of R_i and R_j , and its y -coordinate separates the y -projection of R_i and R_j . There are two basic cases to consider (other cases can be reduced to them by reversing the direction of the $(+y)$ and/or the $(+x)$ -axis).

Case (a): The y -projection of R_i lies below that of R_{i+1} , for $i = 1, 2$. See Fig. 4(a). Let $w \in \mathcal{D}$ be a point that separates R_1 and R_2 , and let $z \in \mathcal{D}$ be a point that separates R_2 and R_3 . The perpendicular bisector of w and z splits R_2 into two subpolygons (one of which might be empty). Denote the one nearer to w by R'_2 and the one nearer to z by R''_2 . With no loss of generality, we may assume that at least half of the vertices in $V_{1,2,3}$ have one of their direction images in R'_2 . We set $C'_1 = C_1$, $C'_3 = C_3$, and C'_2 to be the set of cylinders whose direction images lie in R'_2 . By construction, $|V(C'_1, C'_2, C'_3)| \geq |V_{1,2,3}|/2$. We take the direction $\mathbf{u} \in \mathbb{S}^2$ to be the pre-image of w , i.e., the intersection point of \vec{ow} with \mathbb{S}^2 .

Property (ii) is proved as follows. Let a, b be cylinders such that $\mathbf{n}_a^* \in R_1$ and $\mathbf{n}_b^* \in R'_2$. Then clearly

$$|\mathbf{n}_a^* \mathbf{n}_b^*| \geq \max\{|\mathbf{n}_a^* w|, |\mathbf{n}_b^* w|\},$$

implying that (a, b) are 1-divergent with respect to \mathbf{u} . An identical argument implies that all pairs in $C_1 \times C_3$ are also 1-divergent. Let b, c be cylinders such that $\mathbf{n}_b^* \in R'_2$ and $\mathbf{n}_c^* \in R_3$. Then

$$|\mathbf{n}_b^* \mathbf{n}_c^*| > |\mathbf{n}_b^* z| \geq |\mathbf{n}_b^* w|,$$

implying that (b, c) are also 1-divergent with respect to \mathbf{u} . Hence the lemma holds for this case.

Case (b): The y -projection of R_2 lies above the y -projections of R_1 , which lies above the y -projection of R_3 . See Fig. 4(b). Let $w \in \mathcal{D}$ be a point that separates R_1 and R_2 , and let $z \in \mathcal{D}$ be a point that separates R_2 and R_3 . Let R_0 be the axis-parallel rectangle whose opposite vertices are w and z . Let d_x and d_y denote the lengths of the horizontal and

vertical edges of R_0 , respectively. Assume, without loss of generality, that $d_x \geq d_y$. If $d_y > d_x$, then we reverse the roles of R_1 and R_2 in the following analysis. Let s denote the third vertex of R_0 whose x -coordinate is that of z and whose y -coordinate is that of w .

The perpendicular bisector of w , which is parallel to the y -axis, and s splits R_2 into two subrectangles (one of which might be empty). Denote the one nearer to w by R'_2 and the one nearer to s by R''_2 . Clearly, one of the following two situations arise:

Case (b.i): At least half of the vertices in $V_{1,2,3}$ have a direction image in R'_2 . In this case we take V' to be this subset of vertices; the direction \mathbf{u} is the pre-image of w . The set C'_1 (resp. C'_2, C'_3) consists of those cylinders whose direction images lie in R_1 (resp. in R'_2, R_3).

Property (i) is obvious. Arguing as in case (a), all pairs of cylinders in $C_1 \times C_2$ are 1-divergent. Let a, c be cylinders such that $\mathbf{n}_a^* \in R_1$ and $\mathbf{n}_c^* \in R_3$. Then

$$|\mathbf{n}_a^* \mathbf{n}_c^*| > |\mathbf{n}_a^* z| \geq |\mathbf{n}_a^* w|,$$

where the last inequality follows from the easy observation that the perpendicular bisector of wz does not intersect R_1 , which in turn is a consequence of the assumption $d_x \geq d_y$. Hence, the pair (a, c) are 1-divergent with respect to \mathbf{u} . Similarly, let b, c be cylinders such that $\mathbf{n}_b^* \in R'_2$ and $\mathbf{n}_c^* \in R_3$. Then

$$|\mathbf{n}_b^* \mathbf{n}_c^*| > |\mathbf{n}_b^* s| \geq |\mathbf{n}_b^* w|,$$

implying that (b, c) are also 1-divergent. Hence, the lemma holds for this subcase too.

Case (b.ii): At least half of the vertices have one of their direction images in R''_2 . In this case we set $C'_1 = C_1, C'_3 = C_3$, and C'_2 to be the set of cylinders whose direction images lie in R''_2 , and set \mathbf{u} to be the pre-image of z . Again, property (i) is obvious. Arguing as above, all pairs of cylinders in $C_1 \times C_3$ and in $C'_2 \times C_3$ are 1-divergent with respect to \mathbf{u} . Let a, b be cylinders such that $\mathbf{n}_a^* \in R_1$ and $\mathbf{n}_b^* \in R''_2$. Let \tilde{R}_0 denote the reflection of R_0 about its edge ws . Suppose first that \mathbf{n}_b^* lies outside \tilde{R}_0 . Then

$$|\mathbf{n}_a^* \mathbf{n}_b^*| > |\mathbf{n}_b^* w| \geq |\mathbf{n}_b^* s|.$$

On the other hand,

$$2|\mathbf{n}_b^* s| > |\mathbf{n}_b^* s| + |sz| > |\mathbf{n}_b^* z|,$$

implying that (a, b) are 2-divergent with respect to \mathbf{u} .

Suppose next that \mathbf{n}_b^* lies in \tilde{R}_0 . Let t denote the midpoint of the edge of \tilde{R}_0 opposite to ws . Then

$$|\mathbf{n}_a^* \mathbf{n}_b^*| > |\mathbf{n}_b^* w| > \frac{d_x}{2}.$$

On the other hand,

$$|\mathbf{n}_b^* z| < |tz| = \sqrt{4d_y^2 + \frac{1}{4}d_x^2} \leq \frac{\sqrt{17}}{2}d_x.$$

Hence, we have $|\mathbf{n}_b^* z| < \sqrt{17} |\mathbf{n}_a^* \mathbf{n}_b^*|$, implying that (a, b) are $\sqrt{17}$ -divergent with respect to \mathbf{u} .

This completes the proof of the lemma. □

In view of Lemma 2.7, it suffices to bound the size of $V' = V(C'_1, C'_2, C'_3)$. Set $C' = C'_1 \cup C'_2 \cup C'_3$. All the vertices of V' appear on the boundary of $U(C')$.

2.4. Subdivision into Prisms and the Importance of Being Divergent

Let C' be the set of cylinders as above, and let \mathbf{u} be a direction such that all pairs of cylinders in C' are $\sqrt{17}$ -divergent with respect to \mathbf{u} . We place in \mathbb{R}^3 a grid \mathcal{Q} of infinite square prisms whose axes are parallel to the direction \mathbf{u} ; see Fig. 5. For simplicity of presentation, we rotate the coordinate system to make \mathbf{u} the positive z -direction. The prisms are thus 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. We bound the size of $V_Q = V' \cap Q$ for each $Q \in \mathcal{Q}$ separately and then sum these quantities. Let $\mathcal{C}_Q \subseteq C'$ be the set of cylinders in C' that intersect Q . We call a pair of cylinders $a, b \in \mathcal{C}_Q$ near inside Q if the z -projections of $a \cap Q$ and $b \cap Q$ overlap. Let μ_Q be the number of pairs of cylinders in $\mathcal{C}_Q \times \mathcal{C}_Q$ that are $\sqrt{17}$ -divergent with respect to \mathbf{u} and are near inside Q . We show below in Lemma 2.15 that $|V_Q| = O(\mu_Q \cdot n^\varepsilon)$, for any $\varepsilon > 0$. Hence,

$$|V'| = \sum_{Q \in \mathcal{Q}} |V_Q| = O\left(n^\varepsilon \sum_{Q \in \mathcal{Q}} \mu_Q\right). \tag{2.1}$$

Lemma 2.8. *If a and b are a pair of λ -divergent cylinders with respect to \mathbf{u} (which is assumed to be the same as the z -axis), then (a, b) is near inside at most $O(\lambda^2/t^2)$ prisms of \mathcal{Q} .*




Fig. 5. A system of prisms in direction \mathbf{u} .




Fig. 6. Illustration of the proof of Lemma 2.8

Proof. Suppose to the contrary that there are more than $\zeta^2 \lambda^2 / t^2$ prisms with the property in the lemma, where ζ is a sufficiently large constant that will be specified later. Then there are two prisms, Q and Q' , whose vertical center lines are at distance $d > (\zeta \lambda / t) \cdot t = \zeta \lambda$ apart and the pair (a, b) is near in both Q and Q' . Consequently, there exist four points $\mathbf{p}_a \in a \cap Q$, $\mathbf{p}'_a \in a \cap Q'$, $\mathbf{p}_b \in b \cap Q$, and $\mathbf{p}'_b \in b \cap Q'$, such that \mathbf{p}_a and \mathbf{p}_b have the same z -coordinate, say 0, and \mathbf{p}'_a and \mathbf{p}'_b also have the same z -coordinate, say $h > 0$; see Fig. 6.

We first claim that the angle $\delta = \angle(\mathbf{n}_a, \mathbf{n}_b)$ is small. Indeed, draw two balls B, B' of radius $r = 1 + t\sqrt{2}/2$ about the intersection of the center line of Q with $z = 0$ and about the intersection of the center line of Q' with $z = h$. Then the axes of a and b cross both balls. Translate b so that its axis touches the axis of a at some point $P \in B$, and so that it moves laterally no more than $2r$.

The distance between any point in B and any point in B' is at least

$$\sqrt{d^2 + h^2} - 2r \geq d - 2r > \zeta \lambda - 2r.$$

We obtain a triangle PZR , where Z lies on the axis of a inside B' and R lies on the axis of b at distance at most $3r$ from the center of B' . Hence we have $|PZ| \geq d - 2r \geq \zeta \lambda - 2r$ and $|ZR| \leq 4r$. Hence, by the sine theorem,

$$\frac{\sin \delta}{|ZR|} = \frac{\sin \angle PRZ}{|PZ|} \leq \frac{1}{|PZ|},$$

or

$$\sin \delta \leq \frac{|ZR|}{|PZ|} \leq \frac{4r}{d - 2r} \leq \frac{4r}{\zeta \lambda - 2r},$$

which can be made as small as we wish by choosing ζ large enough.

Next we estimate $|\mathbf{n}_a^* \mathbf{n}_b^*|$. Using the sine theorem once again, we have

$$\frac{|\mathbf{n}_a^* \mathbf{n}_b^*|}{\sin \delta} = \frac{|\mathbf{n}_a^*|}{\sin \theta},$$

where θ is the angle opposite to \mathbf{n}_a^* in the triangle formed by \mathbf{n}_a^* and \mathbf{n}_b^* . By the properties of \mathcal{D} , we have $|\mathbf{n}_a^*| \leq 6$ and $\pi/2 - \beta_0 \leq \theta \leq \pi/2 + \beta_0$. Hence $\sin \theta \geq \cos \beta_0 = \frac{1}{6}$. Thus

$$|\mathbf{n}_a^* \mathbf{n}_b^*| \leq 36 \sin \delta \leq \frac{144r}{d - 2r} \leq \frac{144r}{\zeta \lambda - 2r}.$$

Since the pair (a, b) are λ -divergent with respect to \mathbf{u} , we have, without loss of generality,

$$|\mathbf{n}_a^* \mathbf{u}^*| \leq \lambda |\mathbf{n}_a^* \mathbf{n}_b^*| \leq \frac{144r\lambda}{d - 2r} \leq \frac{144r\lambda}{\zeta \lambda - 2r},$$

which again can be made arbitrarily small if ζ is sufficiently large. This is easily seen to imply that the angle $\gamma = \angle(\mathbf{n}_a, \mathbf{u})$ is also small. Specifically, using the sine theorem yet another time, we have

$$\sin \gamma = \frac{|\mathbf{n}_a^* \mathbf{u}^*| \sin \varphi}{|\mathbf{n}_a^*|} \leq |\mathbf{n}_a^* \mathbf{u}^*| \leq \frac{144r\lambda}{d - 2r},$$

where φ is the angle opposite to \mathbf{n}_a^* in the triangle formed by \mathbf{n}_a^* and \mathbf{u}^* ; we use here the fact that $|\mathbf{n}_a^*| \geq 1$.

On the other hand, we have $\tan \gamma = H/V$, where H (resp. V) is the horizontal (resp. vertical) distance between P and Z . We have $H \geq \zeta \lambda - 2r$ and $V \leq h + 2r$, so that, for sufficiently small γ (that is, for sufficiently large ζ),

$$2 \sin \gamma > \tan \gamma \geq \frac{\zeta \lambda - 2r}{h + 2r}.$$

Note that since γ is small, h must be large, in fact much larger than $\zeta \lambda$, say.

Combining the last two inequalities, we obtain

$$\frac{\zeta \lambda - 2r}{h + 2r} \leq \frac{288r\lambda}{d - 2r} \leq \frac{288r\lambda}{h - 2r},$$

which is a contradiction if ζ is sufficiently large. □

Hence, a pair of cylinders in \mathcal{C}' that are $\sqrt{17}$ -divergent with respect to \mathbf{u} are near inside only $O(1)$ prisms. Putting Lemmas 2.7 and 2.8 together and using (2.1), we obtain that

$$|V_{1,2,3}| \leq 2|V'| = O(n^\epsilon) \cdot \sum_{a,b \in \mathcal{C}'} \# \text{ prisms in which } (a, b) \text{ is a near pair} = O(|\mathcal{C}'|^2 \cdot n^\epsilon).$$

This completes the proof of Lemma 2.5.

2.5. A Weaker Bound on $|V_Q|$

Let $Q = Q_{ij}$ be one of the prisms in \mathcal{Q} , and put, as above, $V_Q = V' \cap Q$. The next stretches of the analysis culminate in Lemma 2.15, which shows that $|V_Q| = O(\mu_Q \cdot n^\epsilon)$, where μ_Q is, as above, the number of pairs of cylinders in $\mathcal{C}' \times \mathcal{C}'$ that are $\sqrt{17}$ -divergent

with respect to \mathbf{u} and are near in Q . (Recall that we rotated the coordinate axis so that the orientation of the $(+z)$ -axis is \mathbf{u} .) This is achieved in two stages. First, in this subsection, we establish a weaker bound on $|V_Q|$ that does not exploit the nearness and divergence of cylinders. Then we sharpen the analysis to obtain the above improved bound.

The main idea in this subsection is to reduce the analysis to the problem of estimating the complexity of a region enclosed between a lower envelope of a collection of surfaces and an upper envelope of another collection, and then to apply the results of [3] that yield a near-quadratic bound on the complexity of such a region.

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) $\angle(\rho, \mathbf{u}) \geq \pi/M$, and
- (C2) each line tangent to (the relative interior of) τ forms an angle of at least π/M with ρ .

We say that $\rho \in \mathbb{S}^2$ is a *good direction* for a vertex v incident upon three canonical strips τ_a, τ_b , and τ_c if it is a good direction for all three strips; see Fig. 7. Recalling that \mathbf{u} is the positive z -direction, it is easily checked that the set B_τ of bad directions for a fixed strip τ , contained in a cylindrical surface $c \in \mathcal{C}$, 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 of opening angles π/M . The area of B_1 is $4\pi(1 - \cos(\pi/M))$.
- Let \mathbf{n}_1 and \mathbf{n}_2 be the normals to the planes tangent to c at the two lines delimiting 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 of the lines 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. 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). The area of B_2 is $4\pi \sin(2\pi/M)$.




Fig. 7. A vertex of the union incident upon three strips τ_a, τ_b , and τ_c , along with a good direction ρ .

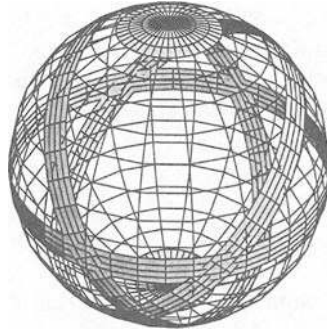


Fig. 8. The set of bad directions for a vertex is contained in the union of two caps and three spherical bands.

It follows that the area of B_τ is at most

$$4\pi \left(1 - \cos \frac{\pi}{M} + \sin \frac{2\pi}{M} \right).$$

This implies that the set of good directions for v contains the complement of the union of two caps with opening angles π/M and of three “great bands,” as above, each of width $4\pi/M$ (see Fig. 8). Hence, the area of this set is at least

$$4\pi \left[1 - \left(1 - \cos \frac{\pi}{M} \right) - 3 \sin \frac{2\pi}{M} \right].$$

By choosing M sufficiently large, the area of the set of good directions 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 some constant opening angle, say, δ , if M is sufficiently large (see Fig. 8).

Let Z be a set of $O(1/\delta^2)$ points on S^2 , with the property that any cap on 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 2.9. *Suppose the horizontal side-length t of a prism Q is less than $\sqrt{2} \sin^2(\pi/M)$. Let $\rho \in Z$, and let v be any vertex in V_Q , incident upon 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 $w \in \tau_a \cap Q$ inside the cylinder a bounded by τ_a in the direction parallel to ρ , we reach ∂Q before exiting a . 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 ∂a 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 $|ww'| = 2 \sin(\pi/M)$. On the other hand, since ww' forms an angle of at least π/M with the z -direction (that is, with \mathbf{u}), it follows that the horizontal distance between w and w' is at least $|ww'| \sin(\pi/M) \geq 2 \sin^2(\pi/M)$. If t , the horizontal side length of prisms in \mathcal{Q} , is chosen such that $t < \sqrt{2} \sin^2(\pi/M)$, then w' does not lie in \mathcal{Q} , which completes the proof of the lemma. \square

Remark 2.10.

- (i) The second part of the lemma crucially uses the fact that the cylinders are infinitely long. Otherwise we may exit a (through its base) before leaving the prism Q . See also Remark 2.16 below.
- (ii) The proof also uses the fact that the radius of the cylinder is 1. It, however, works as long as one can argue that the length of the segment ww' is bounded from below by a constant. For example, the lemma holds even if the radii of the cylinders are different but vary between α and 1, where $\alpha < 1$ is a constant; or if each cylinder is obtained by sweeping a smooth convex planar shape of diameter 1 and of bounded curvature normal to a line in \mathbb{R}^3 .

For a prism $Q \in \mathcal{Q}$ and a direction $\rho \in Z$, let $T_Q(\rho)$ denote the set of canonical strips τ that cross Q and contain at least one vertex in $V_Q(\rho)$. In particular, ρ is a good direction for any $\tau \in T_Q(\rho)$. Let $n_Q(\rho) = |T_Q(\rho)|$. We clip each strip in $T_Q(\rho)$ within Q . We partition $T_Q(\rho)$ into two subsets $T_Q^+(\rho)$ and $T_Q^-(\rho)$ as follows. A (clipped) strip τ contained in a cylinder c belongs to $T_Q^+(\rho)$ (resp. $T_Q^-(\rho)$) if for any point $w \in \tau$, the point $w + \alpha\rho$ lies in the exterior (resp. interior) of c for sufficiently small positive values of α . We define the ρ -upper envelope of $T_Q^+(\rho)$ to be the set of points w on the strips in $T_Q^+(\rho)$ so that a ray from w in the $(+\rho)$ -direction does not intersect any other clipped strip in $T_Q^+(\rho)$. Similarly, we define the ρ -lower envelope of $T_Q^-(\rho)$.

Let τ be a strip in $T_Q^+(\rho)$. Lemma 2.9 implies that any line parallel to ρ that passes through a point in $\tau \cap Q$ meets the interior of the cylinder c in an interval whose other endpoint lies outside Q ; the same property applies when $\tau \in T_Q^-(\rho)$. 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 of the surfaces in $T_Q^+(\rho)$ and the ρ -lower envelope of the surfaces in $T_Q^-(\rho)$. By the result of Agarwal et al. [3], the number of vertices in R_Q 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 2.11. *Let Q be a prism, and let \mathcal{C}_Q be any set of cylinders intersecting Q . Then the number of vertices of the union of (the interiors of the cylinders in) \mathcal{C}_Q lying inside Q is $O(|\mathcal{C}_Q|^{2+\varepsilon})$, for any $\varepsilon > 0$.*

In what follows we will need the following stronger version of the above lemma.

Lemma 2.12. *Let Q be a prism, let A_Q be any set of cylinders intersecting Q , and let B_Q be a subset of A_Q of size b . Then the number of vertices of the union of (the interiors*

of the cylinders in A_Q that lie inside Q and that are incident upon at least two surfaces of B_Q is $O(a \cdot b^{1+\varepsilon})$, for any $\varepsilon > 0$.

Proof. Partition A_Q into $\xi = \lceil a/b \rceil$ subsets A_1, \dots, A_ξ , each of size at most b . Each vertex of the union of A that lies in Q and is incident upon two surfaces of B_Q is a vertex of the union of $B_Q \cup A_i$, for some $1 \leq i \leq \xi$. By Lemma 2.11, the number of vertices in the union of $B_Q \cup A_i$ is $O(b^{2+\varepsilon})$. Hence, the total number of such vertices is $O((a/b) \cdot b^{2+\varepsilon}) = O(ab^{1+\varepsilon})$, for any $\varepsilon > 0$. \square

Remark 2.13. The technique used in the proof of the above lemma applies to the general setup in [3], which yields the following enhancement of the analysis of that paper: Let \mathcal{F} and \mathcal{G} be two sets of n bivariate functions, satisfying the assumptions stated in [3], let M be the “sandwich” region lying between the upper envelope of \mathcal{F} and the lower envelope of \mathcal{G} , and let $\mathcal{H} \subseteq \mathcal{F} \cup \mathcal{G}$ be a subset of size m . Then the number of vertices of M that are incident upon the graphs of at least two functions in \mathcal{H} is $O(nm^{1+\varepsilon})$, for any $\varepsilon > 0$.

2.6. A Stronger Bound on $|V_Q|$

One might interpret Lemma 2.11 as bounding the size of V_Q by $O(\tilde{\mu}_Q \cdot n^\varepsilon)$, where $\tilde{\mu}_Q$ is the number of pairs of cylinders in \mathcal{C} that both intersect Q . Unfortunately, $\tilde{\mu}_Q$ is too large, and $\sum_Q \tilde{\mu}_Q$ may be infinite. There are two “weaknesses” in using $\tilde{\mu}_Q$: it does not take into account divergence and nearness of pairs of cylinders. Both properties are essential for our analysis, as suggested by Lemma 2.8. The purpose of this subsection is to obtain an improved bound on $|V_Q|$ using these properties. This is achieved by combining Lemma 2.12 with a recursive divide-and-conquer analysis that allows us to consider only near (and divergent) pairs of cylinders. Recall that we are assuming that \mathbf{u} is the z -axis and that $Q = Q_{ij}$.

For a cylinder $a \in \mathcal{C}_Q$, let Z_a denote the z -projection of $a \cap Q$. Set $\mathcal{Z}_Q = \{Z_a \mid a \in \mathcal{C}_Q\}$. At each recursive step we have a box $\Pi = [ti, t(i + 1)] \times [tj, t(j + 1)] \times [z_1, z_2]$, for some $z_1, z_2 \in \mathbb{R}$ (a “slice” of Q). Let $\mathcal{C}_\Pi \subseteq \mathcal{C}_Q$ be the set of cylinders that intersect Π . A cylinder $a \in \mathcal{C}_\Pi$ is called *long* in Π if a intersects both the top and bottom faces of Π (i.e., $[z_1, z_2] \subseteq Z_a$), otherwise it is called *short* in Π . Let $L_\Pi, S_\Pi \subseteq \mathcal{C}_\Pi$ denote the sets of long and short cylinders in Π , respectively. Let E_Π be the set of those endpoints of intervals in \mathcal{Z}_Q which lie in the open interval (z_1, z_2) . By the general position assumption and by shifting slightly the grid of prisms, we may assume that all endpoints in E_Π are distinct. We have $|S_\Pi| \leq |E_\Pi| \leq 2|S_\Pi|$. Let $V(L_\Pi, S_\Pi) \subseteq V_Q$ denote the subset of vertices of V_Q that lie in Π and that are incident upon at least two (short) cylinders of S_Π . Initially, $\Pi = Q, L_Q = \emptyset, S_Q = \mathcal{C}_Q, V(L_Q, S_Q) = V_Q$. The recursive process will bound the sizes of the subsets $V(L_\Pi, S_\Pi)$.

If $S_\Pi = \emptyset, |V(L_\Pi, S_\Pi)| = 0$. Otherwise, we partition Π into two subprisms Π_1, Π_2 by a horizontal plane so that the relative interior of the z -projections of each of the two subprisms contains at most half of the endpoints of E_Π . Set $L_1 = L_{\Pi_1}, L_2 = L_{\Pi_2}, S_1 = S_{\Pi_1}$, and $S_2 = S_{\Pi_2}$. For $i = 1, 2$, let $S'_i \subseteq L_i$ be the set of cylinders that are long in Π_i but short in Π . Note that $S_i \cup S'_i$ is the set of all cylinders of S_Π that meet Π_i . Let

$v \in V(L_\Pi, S_\Pi)$ be a vertex lying in Π_1 . If v is incident upon at least two cylinders of S_1 , then $v \in V(L_1, S_1)$. Otherwise, it is incident upon at most one cylinder of S_1 , at most one cylinder of $L_1 \setminus S'_1$, and at least one cylinder of S'_1 . Let V'_1 denote the set of such vertices; V'_2 is defined analogously for Π_2 . It suffices to bound the sizes of V'_1, V'_2 . We define $\Sigma_\Pi \subseteq E_\Pi \times C_\Pi$ to be a set of pairs as follows. A pair $(p, b) \in \Sigma_\Pi$, where p is an endpoint of an interval $Z_a \in \mathcal{Z}_Q$, if the cylinders a and b satisfy the following conditions: (i) $a, b \in C_\Pi$, (ii) they are $\sqrt{17}$ -divergent relative to \mathbf{u} , and (iii) they are near in Q . Set $\sigma_\Pi = |\Sigma_\Pi|$. Since each $\sqrt{17}$ -divergent pair of cylinders that is near inside Q contributes at most two pairs to Σ_Q , we have $\sigma_Q \leq 2\mu_Q$, where μ_Q is, as above, the number of $\sqrt{17}$ -divergent pairs in C_Q that are near in Q .

Lemma 2.14. $|V'_1| + |V'_2| = O(\sigma_\Pi \cdot n^\epsilon)$.

Proof. Let $v \in V'_1$ be a vertex lying on the boundary of three cylinders a, b, c . By definition, up to a permutation of $\{a, b, c\}$, we have $a \in L_1, b \in S'_1$, and $c \in S_1 \cup S'_1$. On the other hand, by definition of V_Q , some permutation of $\{a, b, c\}$ appears in $C'_1 \times C'_2 \times C'_3$ (where C'_1, C'_2 , and C'_3 are as in Lemma 2.7). For specificity, we bound the size of $V(X_1, X_2, X_3) \cap V'_1$, where $X_1 = L_1 \cap C'_1, X_2 = S'_1 \cap C'_2$, and $X_3 = (S_1 \cup S'_1) \cap C'_3$. The other vertices of V'_1 can be counted in a similar manner.

Suppose, without loss of generality, that $|X_1| \geq |X_2| \geq |X_3|$. Then, applying Lemma 2.12 with $A_Q = X_1 \cup X_2 \cup X_3$ and $B_Q = X_2 \cup X_3$, we obtain that the number of vertices of $V(X_1, X_2, X_3)$ that lie in Π_1 is at most $O(n^\epsilon |X_1| \cdot |X_2|)$. Hence, in general, the number of such vertices is at most $O(n^\epsilon \cdot \sum_{i \neq j} |X_i| |X_j|)$. Let $a \in X_i, b \in X_j$, for $i \neq j$. Then, by Lemma 2.7, (a, b) is $\sqrt{17}$ -divergent pair. We charge (a, b) to a pair in Σ_Π . By examining all possible combinations, it suffices to consider only two cases: (i) $a \in L_1 \setminus S'_1$ and $b \in S'_1 \cup S_1$; (ii) $a \in S'_1$ and $b \in S'_1 \cup S_1$. In case (i), one of the endpoints p of Z_b lies in (z_1, z_2) and $p \in Z_a$ (since a is long in Π), so (a, b) is a $\sqrt{17}$ -divergent pair that is near in Q . In case (ii), since a is long in $\Pi_1, Z_a \cap Z_b \neq \emptyset$. Moreover, a and b are both short in Π , so at least one of the endpoints, say p , of $Z_a \cap Z_b$ lies in (z_1, z_2) . If p is an endpoint of Z_a , then $(p, b) \in \Sigma_\Pi$; otherwise, $(p, a) \in \Sigma_\Pi$. Hence, in both cases (a, b) can be charged to a unique pair of Σ_Π , thereby implying that $\sum_{i \neq j} |X_i| |X_j| = O(\sigma_\Pi)$. This completes the proof of Lemma 2.14. \square

Let $\varphi(m, \sigma) = \max |V(L_\Pi, S_\Pi)|$, where the maximum is taken over all pairs L_Π, S_Π such that $|E_\Pi| = m$ and $\sigma_\Pi = \sigma$. Then we obtain the following recurrence:

$$\varphi(m, \sigma) = \begin{cases} 0 & \text{if } \sigma = 0, \\ \varphi(m/2, \sigma_1) + \varphi(m/2, \sigma_2) + O(\sigma n^\epsilon) & \text{if } \sigma > 0. \end{cases}$$

Since $E_{\Pi_1} \cap E_{\Pi_2} = \emptyset, \sigma_1 + \sigma_2 \leq \sigma$. The solution to the above recurrence is

$$\varphi(m, \sigma) = O(\sigma \log m \cdot n^\epsilon) = O(\sigma \cdot n^{\epsilon'})$$

for any $\epsilon' > \epsilon$. Hence, we obtain the following.

Lemma 2.15. *Let Q be a prism in \mathcal{Q} . If there are μ_Q pairs of cylinders in C_Q that are*

$\sqrt{17}$ -divergent with respect to \mathbf{u} and are near inside Q , then $|V_Q| = O(\mu_Q \cdot n^\varepsilon)$, for any $\varepsilon > 0$.

This completes the proof of Theorem 2.1.

Remark 2.16.

- (i) The only place where we need the fact that the cylinders are infinitely long is in Lemma 2.9. The rest of the proof works for bounded cylinders as well. However, if we take a set of n bounded cylinders, each of radius 1 and of sufficiently small height, the complexity of their union can be $\Omega(n^3)$.
- (ii) The current proof does not extend to cylinders with different radii because, as noted in Remark 2.10, Lemma 2.9 uses the fact that the radius of each of the cylinders is 1. However, the above proof, combined with the limited flexibility of Lemma 2.9 (as noted in Remark 2.10) gives an $O(n^{2+\varepsilon})$ bound on the complexity of the union of n cylinders if the ratio of the largest to the smallest radii is bounded by a constant. See also Section 5.

3. The Case of Cigars

We now extend Theorem 2.1 to the case of segments. Let $S = \{s_1, \dots, s_n\}$ now denote a set of n segments in \mathbb{R}^3 . For each i , put $K_i = K_{s_i}$; each K_i is referred to as a *cigar*; see Fig. 9. Let c_i denote the cylindrical portion of ∂K_i , and let σ_i^+, σ_i^- denote the two hemispherical portions of ∂K_i ; the whole boundary is thus $c_i \cup \sigma_i^+ \cup \sigma_i^-$. Let $\mathcal{K} = \{K_1, \dots, K_n\}$ and $U = \bigcup_{i=1}^n K_i$. Let $\mathcal{C} = \{c_1, \dots, c_n\}$ denote the collection of the ∂K_i 's, let $\Sigma = \{\sigma_1^+, \sigma_1^-, \dots, \sigma_n^+, \sigma_n^-\}$ denote the collection of the corresponding hemispherical portions, and let \mathcal{B} denote the set of $2n$ balls whose boundaries contain the hemispheres in Σ .

Again 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 . We assume general position of the segments in




Fig. 9. The union of cigars, the Minkowski sums of line segments and a ball.




Fig. 10. Different types of vertices: (a) *ccc*-vertex, (b) *ccs*-vertex, (c) *css*-vertex, and (d) *sss*-vertex.

S , which now means that every pair of them is skew, 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, that no triple intersection of the boundaries of the K_i 's lie on any circle separating the cylindrical and spherical portions of one of them, and that no four boundaries meet at a point. Each vertex of V is an intersection point 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; see Fig. 10. We denote these vertices mnemonically as *ccc*-, *ccs*-, *css*-, and *sss*-vertices, respectively. We denote the corresponding subsets of V as V_{ccc} , V_{ccs} , V_{css} , and V_{sss} . We bound each of them separately.

3.1. Handling Easy Cases

Any *sss*-vertex v of the union is also a vertex of the union of the $2n$ balls in \mathcal{B} . It is well known that the complexity of the union of m balls in \mathbb{R}^3 is $O(m^2)$ (this follows trivially from [20]), so the number of *sss*-vertices of U is $O(n^2)$.

Lemma 3.1. *The number of *css*-vertices of U is $O(n^{2+\epsilon})$, for any $\epsilon > 0$.*

Proof. We place in \mathbb{R}^3 the same grid \mathcal{Q} of infinite square prisms, as in the previous section, whose axes are parallel to the z -axis. That is,

$$\mathcal{Q} = \{[ti, t(i + 1)] \times [tj, t(j + 1)] \times \mathbb{R} \mid i, j \in \mathbb{Z}\},$$

where t is a sufficiently small constant, as above. For $Q \in \mathcal{Q}$, let $\mathcal{C}_Q \subseteq \mathcal{C}$, $\Sigma_Q \subseteq \Sigma$ be the set of cylindrical and spherical surfaces that intersect Q . Put $m_Q = |\Sigma_Q|$ and $n_Q = |\mathcal{C}_Q|$. Let M be the same constant as in Section 2. We partition each of the cylindrical surfaces in \mathcal{C}_Q into M canonical strips as before, and we cover each sphere in Σ_Q by $O(M^2)$ spherical caps, each of opening angles at most π/M , so that no point lies in more than a constant number of caps. We define a good direction for a spherical cap in the same manner as we did for a strip (see (C1) and (C2) in Section 2.5). 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 defined as follows. Let β_τ be the great circle on S^2 parallel to the tangent plane of the cap τ at its center. We define B_2 to be the spherical band consisting of all points at spherical distance at most $2\pi/M$ from β_τ .

Following the same argument as in Section 2.5, 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 $V_Q = V_{css} \cap Q$. It is now easy to check that both Lemmas 2.9 and 2.11 continue to hold in the extended case. That is, we can decompose the set of cylindrical strips and spherical caps into $u = O(1)$ pairs of subsets $(A_1, B_1), \dots, (A_u, B_u)$, where each A_i, B_i is a subset of strips and/or caps, so that each vertex of V_Q appears in the sandwich region lying between the upper envelope of A_i and the lower envelope of B_i , for some $i \leq u$. This implies that $|V_Q| = O((m_Q + n_Q)^{2+\epsilon})$. However, we want to count the number of *css*-vertices. The argument in the proof of Lemma 2.12 implies that the number of *css*-vertices in Q is $O(m_Q^{1+\epsilon}(m_Q + n_Q))$. Summing over all prisms, the total number of *css*-vertices is $\sum_Q O(m_Q^{1+\epsilon}(m_Q + n_Q))$. Since each hemisphere in Σ intersects $O(1/t^2) = O(1)$ prisms, the total number of *css*-vertices is $O(m^{1+\epsilon}(m + n)) = O(n^{2+\epsilon})$, as claimed. \square

It thus suffices to bound the number of *ccc*- and *ccs*-vertices of U . Using the same argument as in Lemma 2.3, we can again prove that we can choose a subset $C' \subseteq C$ and a direction ρ_0 so that the axes of cylinders in C' form an acute angle of at most $\beta_0 = \cos^{-1}(\frac{1}{6})$ with ρ_0 and the number of *ccc*- and *ccs*-vertices in the union of $C' \cup \Sigma$ is at least half of the number of such vertices in U . We rotate the coordinate system so that ρ_0 becomes the $(+z)$ -axis and remove from C all the cylinders whose axes have an acute angle larger than β_0 with the $(+z)$ -axis. When such a cylinder C_i is removed, we retain the two corresponding balls σ_i^+, σ_i^- . We use \mathcal{C} to denote the remaining set of cylindrical surfaces.

As mentioned in Remark 2.16, only Lemma 2.9 uses the fact that the cylinders in C are unbounded. Nevertheless, the lemma still holds because of the half-balls attached at the endpoints of the segments in S . In other words, a line parallel to a good direction, as in the proof of Lemma 2.9, will exit the whole cigar after exiting Q . Hence, the number of *ccc*-vertices in U is $O(n^{2+\epsilon})$, for any $\epsilon > 0$.

3.2. Bounding the Number of *ccs*-Vertices

We next prove that the number of *ccs*-vertices is also $O(n^{2+\epsilon})$. The proof is very similar to the one described in the previous section, but is considerably simpler, so we mainly focus on the modifications needed to make the proof work for this case.

Let \mathcal{C} be a set of n bounded cylinders of unit radius and let Σ be a set of m unit-radius spheres such that the axes of \mathcal{C} make an acute angle of at most β_0 with the z -axis and the unit spheres centered at the endpoints of the axis of any cylinder in \mathcal{C} are contained in Σ . Let $V = V(\mathcal{C}, \Sigma)$ denote the set of *ccs*-vertices on the boundary of the union of $\mathcal{C} \cup \Sigma$. Set $\varphi(n, m) = \max|V(\mathcal{C}, \Sigma)|$, where the maximum is taken over all sets of n bounded cylinders and over all sets of m spheres that satisfy the axes and containment conditions. We will derive a recurrence for $\varphi(n, m)$ similar to the one in the previous section.

Fix a constant integer parameter $\xi > 2$, whose value depends on ϵ and will be specified later. Partition the plane h into a collection $\mathbf{W} = \{W_1, \dots, W_\xi\}$ of ξ horizontal strips by lines parallel to the x -axis so that each strip contains direction images of at most n/ξ cylinders. For each strip $W_i \in \mathbf{W}$, let \mathcal{C}_i denote the set of cylinders whose direction

images lie in W_i . By construction, $|V(C_i, \Sigma)| \leq \varphi(n/\xi, m)$. Next, we partition the plane into a collection $\mathbf{H} = \{H_1, \dots, H_\xi\}$ of ξ vertical strips by lines parallel to the y -axis, so that each strip contains at most n/ξ direction images. For each strip $H_k \in \mathbf{H}$, we also bound by $\varphi(n/\xi, m)$ the number of ccs -vertices v so that the direction images of the two cylindrical surfaces containing v lie in H_k . These 2ξ subproblems account for all those vertices v of $V(\mathcal{C}, \Sigma)$ that lie on two cylinders whose direction images lie in at most one horizontal or one vertical strip. Let \mathcal{R} be the set of ξ^2 rectangles induced by \mathbf{H} and \mathbf{W} . For a rectangle $R_i \in \mathcal{R}$, let \mathcal{C}_i be the set of cylinders whose direction images lie in R_i . For a pair $i \neq j$, let $V_{i,j} = V(\mathcal{C}_i, \mathcal{C}_j, \Sigma) \subseteq V(\mathcal{C}, \Sigma)$ denote the set of vertices v of $U(\mathcal{C}_i \cup \mathcal{C}_j \cup \Sigma)$ lying on two cylinders a, b such that $a \in \mathcal{C}_i$ and $b \in \mathcal{C}_j$.

Lemma 3.2. *Let R_1, R_2 be two rectangles in \mathcal{R} whose x - and y -projections are disjoint, then $|V_{1,2}| = O((n/\xi)^{1+\varepsilon}m)$.*

Before proving this lemma, we bound the number of the ccs -vertices in U using the lemma. Since there are $O(\xi^4)$ such pairs of rectangles, we obtain the following recurrence:

$$\varphi(n, m) \leq 2\xi \cdot \varphi\left(\frac{n}{\xi}, m\right) + O(n^{1+\varepsilon}\xi^3m).$$

For any $\varepsilon' > \varepsilon$, by choosing $\xi = \xi(\varepsilon')$ a sufficiently large constant, one can prove that the solution to the above recurrence is $\varphi(n, m) = O(n^{1+\varepsilon'}m)$ (see, e.g., [23]). This implies that the number of ccs -vertices in U is $O(n^{2+\varepsilon})$.

We now prove Lemma 3.2. Let $w \in h$ be a point whose x - and y -coordinates separate, respectively, the x -ranges and the y -ranges of R_1 and R_2 , and let \mathbf{u} be the pre-image of w . Then, arguing as in Case (a) of the proof of Lemma 2.7, it follows that all pairs of cylinders in $\mathcal{C}_1 \times \mathcal{C}_2$ are 1-divergent with respect to \mathbf{u} .

In order to bound the size of $V(\mathcal{C}_1, \mathcal{C}_2, \Sigma)$, we place in \mathbb{R}^3 the grid \mathcal{Q} of infinite square prisms, as defined above. We bound the size of $V_Q = V(\mathcal{C}_1, \mathcal{C}_2, \Sigma) \cap Q$ for each $Q \in \mathcal{Q}$ separately and then sum these quantities over all prisms Q . Let $Q \in \mathcal{Q}$ be fixed, and let $\mathcal{C}_Q^{(1)} \subseteq \mathcal{C}_1, \mathcal{C}_Q^{(2)} \subseteq \mathcal{C}_2$, be the subsets of these sets of cylinders that intersect Q , and let $\Sigma_Q \subseteq \Sigma$ be the set of spheres that intersect Q ; set $n_Q = |\mathcal{C}_Q^{(1)} \cup \mathcal{C}_Q^{(2)}|$ and $m_Q = |\Sigma_Q|$. Let ν_Q be the number of pairs of cylinders in $\mathcal{C}_Q^{(1)} \times \mathcal{C}_Q^{(2)}$ that are near inside Q , where nearness is defined as in Section 2 (all these pairs are also 1-divergent with respect to \mathbf{u}). The proof of Lemma 2.11 implies that $|V_Q| = O((n_Q + m_Q)^{2+\varepsilon})$ for any $\varepsilon > 0$. Since we are counting only the number of ccs -vertices, Lemma 2.12 implies that $|V_Q| = O(n_Q^{1+\varepsilon}(n_Q + m_Q))$. Finally, using the same recursive argument as in Section 2.6, we can show that $|V_Q| = O((\nu_Q + n_Q m_Q) \cdot n_Q^\varepsilon)$, for any $\varepsilon > 0$. Hence,

$$|V_{1,2}| \leq \sum_{Q \in \mathcal{Q}} |V_Q| = \sum_Q O((\nu_Q + n_Q m_Q) \cdot n_Q^\varepsilon).$$

By Lemma 2.8, $\sum_Q \nu_Q = O((n/\xi)^{2+\varepsilon})$. Since a sphere in Σ intersects only $O(t^2) = O(1)$ prisms, we have $\sum_Q m_Q = O(m)$. Finally, $m \geq n \geq n/\xi$, therefore $|V_{1,2}| = O((n/\xi)^{1+\varepsilon}m)$. This completes the proof of Lemma 3.2.

Putting everything together we conclude the following.

Theorem 3.3. *Let S be a set of n segments in \mathbb{R}^3 and let B be a ball. The complexity of the union of the Minkowski sums of B and the segments in S is $O(n^{2+\epsilon})$, for any $\epsilon > 0$.*

A result by Clarkson and Shor [14] implies the following corollary which will be useful in the analysis of the next section.

Corollary 3.4. *Let S be a set of n segments in \mathbb{R}^3 and let B be a ball. Set $\mathcal{K} = \{s \oplus B \mid s \in S\}$. The number of vertices of the arrangement of \mathcal{K} that lie in the interior of at most k regions of \mathcal{K} is $O(n^{2+\epsilon} k^{1-\epsilon})$.*

4. The Case of Kreplach

Armed with the bound in Theorem 3.3, we now turn to the general case in which S consists of n pairwise disjoint triangles. For each $s \in S$, let $K_s = s \oplus B_0$. Let $\mathcal{K} = \{K_s \mid s \in S\}$ and $U = \bigcup_{s \in S} K_s$. We also define $\mathcal{K}^{(0)} = \{K_e \mid e \text{ is an edge of a triangle in } S\}$. Let T denote the set of triangular faces of the kreplach in \mathcal{K} , let \mathcal{C} be the set of cylindrical surfaces of cigars in \mathcal{K}_e , and let \mathcal{B} be the set of balls bounding the spherical surfaces of \mathcal{K}_e . A point lying in k regions of \mathcal{K} lies in at most $3k$ regions of $\mathcal{K}^{(0)}$. Let $\mathcal{A}(\mathcal{K})$ (resp. $\mathcal{A}(\mathcal{K}^{(0)})$) be the arrangement defined by the boundary surfaces of the regions of \mathcal{K} (resp. $\mathcal{K}^{(0)}$), and define the *level* of a point p in \mathbb{R}^3 in either arrangement to be the number of regions K_s of the arrangement that contain p in their interior. The closure of the complement of U is the set of points of level 0 in $\mathcal{A}(\mathcal{K})$. The main result of this section is the following.

Theorem 4.1. *Let S be a set of n pairwise disjoint triangles in \mathbb{R}^3 , and let B be a ball. The combinatorial complexity of the union of the Minkowski sums of B with the triangles of S is $O(n^{2+\epsilon})$, for any $\epsilon > 0$.*

As in the previous sections, it suffices to bound the number of vertices of U . Moreover, we can assume general position of the triangles in S , which now means that no pair of triangles in S are parallel or intersect; that no two edges of distinct triangles in S are parallel or coplanar; 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; that no triple intersection of the boundaries of the K_i 's lie on any circle or segment separating the triangular, cylindrical and spherical portions of one of them; and that no four boundaries meet at a point. Using a standard argument based on a slight perturbation of the triangles (as in [24]), one can show that this assumption involves no loss of generality.

4.1. Preliminaries and Overview

We use the shorthand notation of referring to a triangular, cylindrical, or spherical surface as a t -surface, c -surface, and s -surface, respectively. We also use the notation n -surface to

refer to a (“nontriangle”) surface that is either a cylinder or a sphere. As in the preceding section, we call a vertex of $\mathcal{A}(\mathcal{K})$ an xyz -vertex, for $x, y, z \in \{t, c, s, n\}$, if it is incident upon an x -surface, a y -surface, and a z -surface.

Our analysis relies crucially on the following two lemmas. The first lemma, known as the *pseudosphere property*, is an extension of a two-dimensional result by Kedem et al. [20].

Lemma 4.2 (Pseudosphere Property). *Let A_1, A_2 be two disjoint compact, convex bodies in \mathbb{R}^3 , and let B be another compact, convex body with nonempty interior. Let $K_1 = A_1 \oplus B, K_2 = A_2 \oplus B$ be the Minkowski sums of A_1 and A_2 with B . Then the intersection $\partial K_1 \cap \partial K_2$ is connected.*

This lemma was originally obtained by János Pach in the early 1980s. Since this result has never been published, we present in an appendix the proof for the special case in which A_1 and A_2 are triangles and B is a ball (i.e., K_1 and K_2 are kreplach). Recently, another proof, for the polyhedral case, has been given by Hernandez-Barrera et al. [19].

Next, we prove a simple property of kreplach that is used repeatedly in our analysis. We note that this is the only place where the disjointness of the triangles of S is used in the analysis.

Lemma 4.3. *Let s be a triangle in S , and let a, a' be the two triangular portions of ∂K_s . Let t be another triangle in S , and let γ be an arc along ∂K_t that is contained in K_s and connects a point $v \in a$ to a point $v' \in a'$. Then γ must intersect a cylinder or a sphere induced by an edge or a vertex of s ; in other words, the distance of γ from ∂s is smaller than 1. (See Fig. 11.)*

Proof. (We are indebted to Boris Aronov for the following simplification of an earlier more complicated proof.) Suppose to the contrary that this is not the case. For simplicity, assume that s lies in the xy -plane, and that a, a' lie in the planes $z = 1$ and $z = -1$, respectively. For each point $u \in \gamma$, let $\psi(u)$ denote the point in t closest to u (obviously, $\|u - \psi(u)\| = 1$). It is easily seen that ψ is continuous. Let $\delta = \{\psi(u) \mid u \in \gamma\} \subset t$




Fig. 11. Illustration to the proof of Lemma 4.3: (a) triangles s and t and the image δ of γ on t ; (b) K_t and the path γ lying on ∂K_t ; and (c) the cylindrical and spherical surfaces corresponding to the edges and vertices of s and an intersection of γ with a cylindrical surface of K_s .

denote the (connected) image of γ ; see Fig. 11(a). Put $w = \psi(v)$ and $w' = \psi(v')$. Clearly, w lies in the halfspace $z \geq 0$ and w' in the halfspace $z \leq 0$. Since $w, w' \in \delta$ and δ is connected, δ must intersect the plane $z = 0$. Hence, there exists $u \in \gamma$ such that $p = \psi(u) \in t \cap \{z = 0\}$. Since s and t are disjoint, $p \notin s$. Let q be the vertical projection of u on the plane $z = 0$. If $q \notin s$, i.e., u does not lie vertically above s , then u lies inside K_e , for one of the edges $e \in s$, which implies that the path $\gamma[v, u]$ intersects ∂K_e , as claimed. So assume that $q \in \text{int}(s)$. Then pq must cross the boundary of s at some point b . Since $|uq| < 1$ and $|up| = 1$, it follows that $|ub| < 1$, which establishes the lemma. \square

Remark 4.4. The above proof relies on the fact that s is planar, but it does not use the fact that s is polygonal. The proof works as long as S is a family of pairwise-disjoint convex planar objects, e.g., a family of pairwise-disjoint disks.

We derive a recurrence similar to the ones used in the analysis of the complexity of lower envelopes and other substructures in arrangements (see, e.g., [24] for details), but we use a simple enhancement of it, as follows. Let χ denote the (constant) maximum possible number of intersections between any three boundary surfaces of regions in \mathcal{K} . For three triangles $a, b, c \in S$, let v be a vertex incident upon the boundaries of three regions K_a, K_b, K_c . Let N denote the network formed by the vertices and edges (i.e., 1-skeleton) of $K_{abc} = K_a \cap K_b \cap K_c$. By Lemma 4.2, any pair of boundaries ∂K_a and ∂K_b intersect in a connected curve, which implies that N is connected. Let $m \leq \chi$ be the number of vertices in N . Let N_v be the set of vertices in N , including v , that do not lie in the interior of any krepl and that can be reached from v along the edges of N without intersecting any other krepl. We define the *index* of v , denoted $\text{ind}(v)$, to be $m - |N_v|$. $\text{ind}(v) = m$ is equivalent to $v \notin \partial U$; $\text{ind}(v) = m - 1$ is equivalent to $v \in \partial U$ but each of the three edges of K_{abc} adjacent to v is intersected by at least one other region; $\text{ind}(v) = 0$ is equivalent to the entire network N not being intersected by any other krepl. For $0 < j < m$, we call a vertex v of index j a *frontier vertex* if an edge of N adjacent to v crosses the boundary of a krepl; by definition, if v is a vertex of index $0 < j < m$, then N_v contains at least one frontier vertex. If we remove some of the triangles from S , excluding the three whose expansion boundaries are incident upon v , the index of v can only decrease or remain unchanged. Note that the notion of an index used here is different from the one used in the previous works (as presented in [24]).

Let $F^{(j)}(S)$ denote the number of vertices of \mathcal{A} of index at most j that lie on ∂U , and let

$$F^{(j)}(n) = \max_{|S|=n} F^{(j)}(S).$$

Let $F(S) = F^{(\chi-1)}(S)$ denote the overall number of vertices of ∂U ; set

$$F(n) = \max_{|S|=n} F(S).$$

We now derive the recurrence for $F^{(j)}(n)$, for $j > 0$. For $j = 0$ we need a special analysis of the structure of the sets K_{abc} , which lies at the heart of our proof.

We choose some threshold parameter $\xi = \xi_j$ that we will fix later. Let v be a vertex of index $j > 0$, lying on the boundaries of K_a, K_b , and K_c . If v is not a frontier vertex, we

charge v to a frontier vertex of N_v ; each frontier vertex is charged at most χ times. If v is a frontier vertex, then let e be an edge of K_{abc} adjacent to v that is crossed by the boundary of another krepl K_d . If e is crossed by at least ξ other boundary surfaces, we charge v (and the nonfrontier vertices charged to v) to the first ξ vertices of \mathcal{A} encountered along e . These vertices are at level at most ξ in $\mathcal{A}(\mathcal{K})$ and each can be charged this way only $O(1)$ times. Hence, applying the Clarkson–Shor probabilistic analysis technique [14] and arguing as in earlier proofs (see [24]), the number of vertices v at level at most ξ is $O(\xi^2 F(n/\xi))$. Otherwise, if we remove the at most ξ triangles whose expansion boundaries meet e (but retain a, b , and c), then the index of v decreases by at least one. Hence, applying again the Clarkson–Shor technique, the number of vertices v of this kind is $O(\xi^3 F^{(j-1)}(n/\xi))$.

We thus obtain the following recurrences, for $j = 1, \dots, \chi - 1$:

$$F^{(j)}(n) = O\left(\xi_j^2 F\left(\frac{n}{\xi_j}\right) + \xi_j^3 F^{(j-1)}\left(\frac{n}{\xi_j}\right)\right). \tag{4.1}$$

We next derive a recurrence for $F^{(0)}(n)$.

Lemma 4.5. *Let $F^{(tnn)}(n)$ be the maximum number of tnn -vertices of index 0 on the union, maximized over all sets of n pairwise-disjoint triangles. Then, for any parameters ξ_0, ζ_0 , and $\varepsilon > 0$, we have*

$$F^{(0)}(n) = O(\xi_0^{3-\varepsilon} n^{2+\varepsilon}) + O\left(\xi_0^2 F\left(\frac{n}{\xi_0}\right) + \xi_0^3 F^{(tnn)}\left(\frac{n}{\xi_0}\right)\right), \tag{4.2}$$

$$F^{(tnn)}(n) = O(\zeta_0^{3-\varepsilon} n^{2+\varepsilon}) + O\left(\zeta_0^2 F\left(\frac{n}{\zeta_0}\right)\right).$$

Following an argument similar to the one in [23], one can show that the combined solution of the recurrences (4.1) and (4.2) satisfies $F(n) = O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. In the remainder of the section we prove the above lemma.

4.2. Bounding $F^{(0)}(n)$

Let v be a vertex of index 0 lying on three kreplach K_a, K_b, K_c . Then all vertices of K_{abc} lie on ∂U and none of the edges of K_{abc} meets any other kreplach. We refer to such a vertex v as a *free* vertex and to such a K_{abc} as a *free triple intersection*. We charge all free vertices of K_{abc} to some specific representative vertex on K_{abc} and count the number of representative vertices. This counting is done in several stages, depending on the type of representative vertices. The overall analysis will lead to the recurrences (4.2).

Handling Easy Cases. The definition of a free triple intersection K_{abc} only implies that its edges do not intersect the other kreplach, but it still allows the 2-faces of K_{abc} to meet other regions. If a 2-face f of K_{abc} lying, say, on ∂K_a intersects another krepl K_d but no edge of K_{abc} intersects K_d , then a whole connected component γ of the intersection curve $\partial K_a \cap \partial K_d$ lies entirely in f . Lemma 4.2 implies that $\partial K_a \cap \partial K_d$

is connected, therefore $\gamma = \partial K_a \cap \partial K_d$. We charge f to γ . Hence, the number of free triple intersections that intersect other kreplach is only $O(n^2)$. We can thus assume that the entire K_{abc} does not intersect any krepl. (It is easily seen that no K_d can be fully contained in the interior of K_{abc} .) If any of the edges of K_{abc} has a *transition point*, namely, a point on a seam of a t -, c - or s -surface of a krepl, then we can charge K_{abc} to that transition point. Since there are only $O(n^2)$ transition points, the number of such free triple intersections is also $O(n^2)$. This also implies that there are $O(n^2)$ free triple intersections K_{abc} that contain ttt -vertices. Indeed, let v be a ttt -vertex incident upon three triangular faces of K_{abc} . Since v is the only intersection point of the corresponding three triangles, at least one of the edges of K_{abc} adjacent to v contains a transition point, thereby implying that there are $O(n^2)$ such triple intersections.

Next, if K_{abc} has an nnn -vertex v (a vertex that does not lie on any displaced triangle), we choose v as its representative vertex; v is also a vertex of the union of $\mathcal{K}^{(0)}$. By Theorem 3.3, the number of such vertices is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Suppose next that K_{abc} has a tss -vertex. Let a' be a triangular face of K_a . Since every tss -vertex of U lying on a' is a vertex of the union of a set of at most $3n$ disks, within the plane containing a' , the number of such vertices is $O(n)$ [20]. Hence, there are $O(n^{2+\varepsilon})$ free triple intersections that contain an nnn - or a tss -vertex.

In view of the above discussion, we can thus assume that each vertex of K_{abc} lies on at least one triangular face, that K_{abc} has no ttt - or tss -vertex, that K_{abc} is disjoint from any other krepl, and that none of the edges of K_{abc} contains a transition point. Then all vertices of K_{abc} are tcn - or ttn -vertices. We call such triple intersections *interesting*. We call a vertex *interesting* if it is a vertex of an interesting triple intersection.

The rest of the proof, which bounds the number of interesting free triple intersections, consists of two parts. The first part bounds the number of interesting triple intersections that contain at least one tcn -vertex. We show that the number of interesting tcs -vertices is proportional to the number of certain degree-2 faces, called *bubbles*, in the arrangement of \mathcal{K} and $\mathcal{K}^{(0)}$. Following an approach similar to the one used in [17], we obtain a recurrence that bounds the number of these bubbles. The same recurrence can be derived to bound the number of interesting tcc -vertices. The second part of the proof bounds the number of interesting triple intersections that contain only ttn -vertices. Roughly speaking, we choose a parameter ξ and charge each ttn -vertex either to ξ ttn -vertices of level at most ξ or to one tcn -vertex of level at most ξ .

4.3. Bounding the Number of Interesting tcn -Vertices

We derive a recurrence for the number of interesting tcs -vertices. Let v be a tcs -vertex lying on some K_{abc} . Suppose v lies on the t -surface of the triangle a , and let e and p be the original edge (say of b) and vertex (of c) whose expanded cylinder and ball, respectively, contain v on their boundaries. We replace b by e and c by p and consider the triple intersection K_{aep} . This set is contained in K_{abc} but is otherwise free of intersections with any other region K_d (because K_{abc} avoids all these regions). We call v a *regular* tcs -vertex if all vertices of K_{aep} lie on one of the triangular faces a' of K_a and on the cylindrical surface of K_e . Otherwise, it is called *irregular*.

Lemma 4.6. *There are $O(n^{2+\epsilon})$ irregular vertices on interesting triple intersections.*

Proof. If an interesting K_{aep} contains an irregular vertex, then either it contains an nnn -vertex, or one of the edges of K_{aep} contains a transition point, or the vertices of K_{aep} lie on two distinct triangular faces of one of the kreplach. By the previous discussion, there are $O(n^{2+\epsilon})$ triple intersections of the first two types.

Suppose there exist two vertices of K_{aep} that lie on two distinct triangular faces of K_a . Since $\gamma = \partial K_e \cap \partial K_p$ is connected (as already noted, this is a consequence of Lemma 4.2, but can also be verified explicitly), it follows that there is a portion of γ that lies on ∂K_{aep} and connects between two points that lie in the two displaced copies of a . By Lemma 4.3, this portion of γ must intersect one of the expanded edges $K_{e'}$ of a , at a ccs - or csc -vertex that lies on the union of $\mathcal{K}^{(0)}$. The number of such vertices is $O(n^{2+\epsilon})$. Hence, there are $O(n^{2+\epsilon})$ irregular vertices. \square

It thus suffices to bound the number of regular tcs -vertices.

Bounding the Number of Regular tcs -Vertices. Let v be a regular vertex on K_{aep} , i.e., all vertices of K_{aep} lie on a displaced copy a' of a , on the cylindrical surface C_e of K_e , and on the sphere ∂K_p . As we follow the boundary of $R = a' \cap C_e \cap \partial K_p$ from v , we encounter only those vertices at which the intersection ellipse of a' and the cylindrical surface C_e crosses the intersection circle of a' and ∂K_p , implying that R has either two or four vertices. Since all vertices of K_{aep} lie on a' , it follows that K_{aep} has only two or four vertices. We first consider the case in which K_{aep} has exactly four vertices, all lying on the triangle a' . We consider ∂K_{aep} as a spherical map, and apply to it Euler's formula, as follows. The map has $V = 4$ vertices and each vertex is of degree 3. Moreover, as is easily seen, each face of the map has even degree, namely, either 2 or 4. Suppose there are E edges, F_2 faces of degree 2, and F_4 faces of degree 4. Since each vertex has degree 3, we have $E = 6$. Then Euler's formula yields

$$V + F_2 + F_4 = E + 2 \quad \text{or} \quad F_2 + F_4 = 4.$$

We also have $2E = 2F_2 + 4F_4$, or $F_2 + 2F_4 = 6$, thereby implying that $F_2 = F_4 = 2$. It is easily verified that K_a contributes to ∂K_{aep} one face of degree 4 (on a'), that another surface contributes another face of degree 4, and that the third surface contributes two faces of degree 2. See Fig. 12. On the other hand, if $V = 2$, then $E = 3$, and $F_2 + F_4 = 3$. Moreover, $2F_2 + 4F_4 = 6$, which implies that $F_4 = 0$ and $F_2 = 3$. That is, each of a' , K_e , and K_p contributes a 2-face to K_{aep} .




Fig. 12. An example of a regular tcs -vertex.

Lemma 4.7. *For any parameter $\xi \geq 1$, there are $O(\xi^2 F(n/\xi) + \xi^{3-\varepsilon} n^{2+\varepsilon})$ regular tcs -vertices v on free interesting K_{aep} 's, such that*

- (i) either K_{aep} has two vertices, or
- (ii) K_{aep} has four vertices and its two degree-2 faces lie on ∂K_e .

Proof. As defined earlier, let \mathcal{C} be the set of cylindrical surfaces of the ∂K_e 's, where e is an edge of a triangle in S . We bound the number of desired vertices that lie on each surface $C \in \mathcal{C}$ and sum these bounds up over all surfaces in \mathcal{C} . Assume that the axis of C is parallel to the z -axis. Let $\mathcal{K}_C = \{K_a \cap C \mid a \in S\}$ and $\mathcal{K}_C^{(0)} = \{K_g \cap C \mid K_g \in \mathcal{K}^{(0)}\}$. Let $\mathcal{A} = \mathcal{A}(\mathcal{K}_C)$, $\mathcal{A}^{(0)} = \mathcal{A}(\mathcal{K}_C^{(0)})$, $U_C = \bigcup \mathcal{K}_C$, and $U_C^{(0)} = \bigcup \mathcal{K}_C^{(0)}$. Clearly, $U_C^{(0)} \subseteq U_C$. Let ν_C denote the combinatorial complexity of $U_C^{(0)}$. By Theorem 3.3, we have $\sum_C \nu_C = O(n^{2+\varepsilon})$, where the sum ranges over all surfaces in \mathcal{C} . The level of a point $q \in C$ with respect to \mathcal{A} (resp. $\mathcal{A}^{(0)}$) is the number of regions in \mathcal{K}_C (resp. $\mathcal{K}_C^{(0)}$) that contain q in their interior. The closure of the complement of U_C (resp. $U_C^{(0)}$) is the set of points at level 0 with respect to \mathcal{A} (resp. $\mathcal{A}^{(0)}$).

The intersection of C with a triangular face a' of K_a , for $a \in S$, is an elliptic arc. Any pair $C \cap a'$, $C \cap b'$ of these elliptic arcs intersect in at most two points because $a' \cap b'$ is a line segment and it intersects C in at most two points. Moreover, any generator line on C , a line parallel to its axis, intersects any of these elliptic arcs in at most one point, which is the intersection of the generator with the respective displaced triangle. Finally, an endpoint of any elliptic arc is a transition point that lies on the boundary of the corresponding displaced triangle a' . Let v be a regular tcs -vertex of one of the degree-2 faces induced on C by K_{aep} (i.e., a vertex of K_{aep}). Note that v lies on an elliptic arc γ (a portion of the intersection of a' with C) and on a portion of a sphere-cylinder intersection curve δ . Since K_{aep} is free, by definition, the degree-2 faces of ∂K_{aep} on C appear as faces (which we refer to as *bubbles*) of the arrangement \mathcal{A} . Moreover, γ and δ appear in a fixed vertical order along C outside these bubbles (i.e., any generator that crosses both curves crosses them in the same order); see Fig. 13. We call the bubble *upward* if the elliptic arc γ is the top edge of the bubble; otherwise we call it *downward*. If γ and δ form upward bubbles, then δ lies *above* γ outside these bubbles. We bound the number of upward bubbles that do not lie inside any region of \mathcal{K}_C . By reversing the direction, we




Fig. 13. Bubbles of \mathcal{A} and quasi-regular vertices on C .

obtain a similar bound on the number of downward “free” bubbles. Together, this yields a bound on the overall number of desired vertices.

Let ξ be a parameter. Let $Q_C^{(\xi)}$ denote the closure of the points on C whose level is at most ξ with respect to $\mathcal{K}_C^{(0)}$. We take the θ -extremal points w (in the cylindrical coordinate system attached to C) of every edge of $\mathcal{A}^{(0)}$ that lies inside $Q_C^{(\xi)}$ and draw through w a maximal vertical segment contained in $Q_C^{(\xi)}$ and crossing at most ξ arcs of $\mathcal{A}^{(0)}$ on either side of w . The edges of $\mathcal{A}^{(0)}$ and these segments decompose $Q_C^{(\xi)}$ into “pseudotrapezoidal” cells, and we denote the collection of these cells by \mathcal{V} . Using the Clarkson–Shor analysis technique, the total number of vertices of $\mathcal{A}^{(0)}$ within $Q_C^{(\xi)}$, summed over all cylinders C , is $O(\xi^3(n/\xi)^{2+\varepsilon}) = O(\xi^{1-\varepsilon}n^{2+\varepsilon})$. Hence the total number of cells and edges in the decompositions \mathcal{V} , summed over all cylinders C , is $O(\xi^{2-\varepsilon}n^{2+\varepsilon})$, for any $\varepsilon > 0$.

Note that the new vertical segments may split some of the “bubbles” into two faces, but the number of such bubbles, summed over all cylinders, is only $O(\xi^{2-\varepsilon}n^{2+\varepsilon})$ since each vertical segment splits at most ξ bubbles. We thus count only those upward bubbles whose vertices both lie on the same edge of a cell of \mathcal{V} .

For each edge δ in \mathcal{V} whose level is 0 with respect to $\mathcal{K}_C^{(0)}$, we count the number of upward bubbles formed by δ that were not split by the vertical segments and sum this quantity over all such edges. Let \mathcal{E}_δ be the set of elliptic arcs that form upward bubbles with δ ; set $m_\delta = |\mathcal{E}_\delta|$. Each arc in \mathcal{E}_δ intersects δ in either two or four points and all these intersection points are the vertices of the bubbles; otherwise the corresponding bubble is not a face of a regular triple intersection. If $m_\delta \leq \xi$, the number of upward bubbles that lie on δ is at most 2ξ . We charge them to δ . The total number of such bubbles charged to arcs of \mathcal{V} , summed over all cylinders C , is $O(\xi^{3-\varepsilon}n^{2+\varepsilon})$.

Suppose next that $m_\delta > \xi$. Let γ be an elliptic arc that forms a bubble f with δ . Let v_L, v_R be the left and right vertices of f . First assume that γ intersects δ at two points. We trace γ from v_L (resp. v_R) leftward (resp. rightward) until we reach a point w_L (resp. w_R) for which one of the following conditions holds:

- (C1) we have reached an endpoint of γ ;
- (C2) we have encountered ξ vertices of \mathcal{A} ;
- (C3) we have reached a point that lies below an endpoint of δ .

By construction, $\gamma[w_L, v_L]$ and $\gamma[v_R, w_R]$ lie below δ .

Claim 4.8. *The relative interiors of the traced arcs $\gamma[w_L, v_L]$ and $\gamma[v_R, w_R]$ do not contain a vertex of an upward bubble.*

Proof. Let w'_L be the point on δ lying vertically above w_L , and let ρ_L be the region bounded by the arcs $\delta[w'_L, v_L]$, $\gamma[w_L, v_L]$ and the vertical segment $w_L w'_L$ (e.g., the left shaded region in Fig. 14). Similarly we define the region ρ_R lying between $\gamma[v_R, w_R]$ and δ . Suppose the relative interior of $\gamma[w_L, v_L]$ contains a vertex $v' \in \delta' \cap \gamma$ of an upward bubble formed by some curve δ' and γ . We assume that v' is the rightmost such vertex. We claim that the right endpoint of δ' lies in ρ_L . First, we observe that the bubble β containing v' lies to the left of v' . Indeed, if v' were the left vertex of β , then the right vertex of β would have to lie to the right of v_L because $\gamma(v', v_L)$ does not contain the




Fig. 14. Tracing an elliptic arc γ .

vertex of any upward bubble. But then β contains the arc $\delta[v_L, v_R]$, implying that β is not a face of \mathcal{A} , a contradiction. Hence v' is the right vertex of β . Since β lies to the left of v' , the arc δ' lies above γ to the right of v' . All intersection points of γ and δ' are vertices of the upward bubbles formed by them, so δ' cannot intersect $\gamma[v', v_L]$. Since δ is an edge of \mathcal{V} , δ' does not intersect δ . This implies that the right endpoint of δ' has to lie in the region ρ_L , as claimed above.

Let σ be the rightmost endpoint of an arc in $\mathcal{K}_C^{(0)}$ that lies inside ρ_L , and let $\sigma' \in \delta$ be the point lying vertically above σ . Any arc of $\mathcal{K}_C^{(0)}$ intersecting the segment $\sigma\sigma'$ has to intersect $\gamma[w_L, v_L]$ because it can neither intersect δ nor end inside ρ_L . Since $\gamma[w_L, v_L]$ contains at most ξ vertices of \mathcal{A} , the vertical segment $\sigma\sigma'$ intersects at most ξ arcs of $\mathcal{K}_C^{(0)}$. However, then the vertical segment erected through σ would have to intersect δ , thereby implying that σ' is the left endpoint of δ , a contradiction. Hence, v' does not exist. The same argument applies to ρ_R . \square

Actually, the preceding argument shows that no arc in $\mathcal{K}_C^{(0)}$ (or in \mathcal{K}_C) has an endpoint inside ρ_L or ρ_R .

If w_L or w_R is an endpoint of the elliptic arc γ , we charge f to γ . Since no other upward bubble can be charged to the same endpoint of γ , each elliptic arc is charged at most twice. Hence, the total number of such bubbles over all cylindrical surfaces in \mathcal{C} is $O(n^2)$. If the traced portion of γ (i.e., $\gamma[w_L, v_L] \cup \gamma[v_R, w_R]$) contains ξ vertices of \mathcal{A} , we charge f to ξ of these vertices whose levels are at most ξ . Each such intersection point is charged by $O(1)$ upward bubbles, over all cylinders C .

If we are not able to charge f to an endpoint of γ or to the vertices of \mathcal{A} , then w_L lies below the left endpoint of δ and w_R lies below the right endpoint of δ . Since ρ_L and ρ_R do not contain the endpoints of any elliptic arc $\gamma' \in \mathcal{E}_\delta$ and γ' does not intersect $\gamma[v_L, v_R]$, γ' has to intersect the traced portion of γ . Repeating this argument for all arcs of \mathcal{E}_δ and recalling that we have assumed $m_\delta > m$, we conclude that the traced portion of γ contains at least ξ vertices of \mathcal{A} , a contradiction. Hence, we are always able to charge an upward bubble.

Next, if γ and δ form two upward bubbles (as in Fig. 13), then lqt $v_1 = v_L, v_2, v_3, v_4 = v_R$ be the four intersection points of δ and γ , sorted from left to right. We trace γ from v_L and v_R as earlier and stop as soon as one of conditions (C1)–(C3) holds. In addition, we also trace γ from v_2 rightward until we either collect ξ vertices of \mathcal{A} or we reach v_3 . If $\gamma[v_2, v_3]$ contains less than ξ vertices of \mathcal{A} , the above argument implies that the region formed by $\gamma[v_2, v_3]$ and $\delta[v_2, v_3]$ does not contain the endpoints of any arc in

\mathcal{E}_δ . Hence, even in this case each arc of \mathcal{E}_δ intersects the traced portion of γ and we can charge both upward bubbles to ξ vertices of \mathcal{A} whose levels are at most ξ .

Repeating the same argument for downward bubbles and summing over all arcs δ of level 0 in \mathcal{V} and over all cylinders C , we conclude that the number of quasi-regular vertices incident upon upward or downward bubbles formed by those edges of \mathcal{V} (of level 0) for which $m_\delta > \xi$ is $O(n^2 + F_{\leq \xi}(n)/\xi)$, where $F_{\leq \xi}(n)$ is the number of vertices of level at most ξ in an arrangement of n kreplach. By a result of Clarkson and Shor [14], $F_{\leq \xi}(n) = O(\xi^3 F(n/\xi))$. Adding the number of bubbles that lie on edges δ of \mathcal{V} for which $m_\delta \leq \xi$, we conclude that the total number of bubbles is $O(\xi^2 F(n/\xi) + \xi^{3-\epsilon} n^{2+\epsilon})$. This completes the proof of the lemma. \square

Next, the case in which the degree-2 faces of K_{aep} lie on the sphere ∂K_p can be handled in a similar manner. We take ∂K_p and draw on it the arrangements \mathcal{A} , formed by its intersections with the regions $K_s \in \mathcal{K}$, and $\mathcal{A}^{(0)}$, formed by its intersection with the regions $K_e \in \mathcal{K}^{(0)}$. The degree-2 faces of a regular triple intersection K_{aep} appear as two faces of \mathcal{A} . We draw a (θ, φ) -coordinate system of longitudes and latitudes on ∂K_p and regard the longitudes of ∂K_p as the generator lines. If a circular arc γ is not θ -monotone, then we split γ at the points that are tangent to longitudes. We now proceed exactly as in the previous case. A similar argument shows that the overall number of regular tcs -vertices that lie on free interesting K_{aep} 's in which the two degree-2 faces lie on ∂K_p is also $O(\xi^2 F(n/\xi) + \xi^{3-\epsilon} n^{2+\epsilon})$. We leave it to the reader to verify the details.

To conclude, we have shown the following.

Lemma 4.9. *The number of free triple intersections K_{abc} that contain a regular tcs -vertex is at most*

$$O\left(\xi^{3-\epsilon} n^{2+\epsilon} + \xi^2 F\left(\frac{n}{\xi}\right)\right). \tag{4.3}$$

Bounding the Number of tcc -Vertices. Next suppose that K_{abc} has no tss -vertex and no tcs -vertex but has a tcc -vertex v . The analysis of this case is very similar to that of a tcs -vertex, with the following modification. In full analogy, we consider the intersection $K_{aee'}$, where e and e' are edges of b and c , respectively, on whose expanded cylinders v lies. We may assume that $K_{aee'}$ does not have a tcs -vertex, tss -vertex, nnn -vertex, or a transition point on any intersection curve because then we can apply the same analysis as above to conclude that (4.3) bounds the number of such free triple intersections. We define $\gamma = a' \cap \partial K_e$, $\delta = a' \cap \partial K_{e'}$, and $R = a' \cap K_e \cap K_{e'}$. Arguing as above, the preceding assumptions imply that all vertices of $K_{aee'}$ lie on $E \cap E'$, where E and E' are the elliptic intersection curves of a' with the cylindrical portions of ∂K_e and $\partial K_{e'}$, respectively; in particular, there are only two or four such vertices. The rest of the analysis proceeds exactly as above, and implies that the overall number of free triple intersections K_{abc} of the above type is bounded by the bound in (4.3).

4.4. Bounding the Number of ttn -Vertices

We now bound the number of interesting triple intersections that contain only ttn -vertices (Fig. 15). Let K_{abc} be such a triple intersection. The preceding analysis implies that all




Fig. 15. An example of a ttn -vertex.

vertices of K_{abc} lie on a single displaced copy a' of a and on a single displaced copy b' of b . However, then all vertices lie on the line segment $a' \cap b'$, and the convexity of K_{abc} implies that it has only two vertices v, v' .

We first assume that one of them, say v , is a ttc -vertex, lying on the cylindrical portion C of ∂K_e , for some edge e of c . Since we assume the edges of K_{abc} do not contain any transition point, the edges adjacent to v and lying on K_e lie fully in C . Therefore v' also lies on C and ∂K_{abc} has three (free) edges, one of which is the straight segment vv' and the other two are elliptic arcs contained in $a' \cap C$ and $b' \cap C$, respectively. We only study this case; the case in which v is a ttS -vertex is treated in essentially the same manner, replacing C by an appropriate sphere and the elliptic arcs by circular arcs along that sphere.

As in the proof of Lemma 4.7, we assume that the axis of C is vertical, and we form two arrangements on C . Let $\mathcal{K}_C, \mathcal{K}_C^{(0)}, \mathcal{A}, \mathcal{A}^{(0)}, U_C, U_C^{(0)}$, and v_C be the same as defined in that proof; recall that $\mathcal{A}, \mathcal{A}^{(0)}$ are the arrangements of \mathcal{K}_C and $\mathcal{K}_C^{(0)}$, respectively. Let \mathcal{E} be the set of at most $2n$ elliptic arcs in \mathcal{K}_C , formed by the intersection of C with the triangular faces of kreplach in \mathcal{K} . We take the complement of $U_C^{(0)}$ within C and decompose it into pseudotrapezoidal cells, by extending a vertical segment from each vertex or θ -extreme point on $\partial U_C^{(0)}$ until it hits this boundary again. The total number of cells, over all cylindrical surfaces in \mathcal{C} , is $O(n^{2+\epsilon})$. Let \mathcal{V} denote the resulting vertical decomposition.

Fix a cell τ of this vertical decomposition, and consider the set $\mathcal{E}_\tau \subseteq \mathcal{E}$ of all elliptic arcs that cross τ and that contain at least one ttc -vertex; set $m_\tau = |\mathcal{E}_\tau|$. Any ttc -vertex v that lies in τ is an intersection of two elliptic arcs in \mathcal{E}_τ . Since each endpoint of an elliptic arc lies on the boundary of a region in $\mathcal{K}_C^{(0)}$, none of the arcs in \mathcal{E}_τ can have an endpoint inside τ . Let $E \in \mathcal{E}_\tau$ be an elliptic arc and let ℓ be a generator line on C that intersects E . If we follow ℓ from $E \cap \ell$ (recall that there is a unique such point) into the region K_a bounded by E and apply Lemma 4.3, we conclude that we will meet some cylindrical surface in \mathcal{C} or some sphere in \mathcal{B} before exiting K_a , and therefore we will exit τ before exiting K_a . Let \mathcal{E}_τ^+ (resp. \mathcal{E}_τ^-) be the set of elliptic arcs $E \in \mathcal{E}_\tau$ so that a ray emanating from a point on the arc (within τ) in the $(+z)$ -direction (resp. $(-z)$ -direction) enters the corresponding K_a .

It follows that any ttc -vertex v under consideration is a vertex of the region lying between the lower envelope of \mathcal{E}_τ^+ and the upper envelope of \mathcal{E}_τ^- . Since any pair of arcs in \mathcal{E}_τ intersect in at most two points, it follows that the complexity of this sandwich region, and thus also the number of ttc -vertices under consideration within τ , is $O(m_\tau)$. It thus suffices to bound the value of $\sum_{C \in \mathcal{C}} \sum_{\tau \in \mathcal{V}} m_\tau$.




Fig. 16. Counting the number of arcs in \mathcal{E}_τ .

We fix a threshold parameter $\xi > 0$. The overall number of *ttc*-vertices under consideration that lie in cells τ with $m_\tau \leq \xi$, over all cylinders C , is $O(\xi n^{2+\varepsilon})$, so assume that $m_\tau > \xi$.

Since each arc $E \in \mathcal{E}_\tau$ intersects the boundary of τ , let w_E be any such intersection point; see Fig. 16. If the level of w_E with respect to \mathcal{A} is at most ξ , we charge E to w_E . There are at most ξ such intersection points lying on each of the vertical edges of τ . Summed over all cells in \mathcal{V}_C and over all cylindrical surfaces in C , the number of such intersection points is $O(\xi n^{2+\varepsilon})$. If w_E lies on the top or bottom boundary of τ , w_E is a *tcc*- or a *tcs*-vertex of \mathcal{A} . Using the Clarkson–Shor analysis technique, we conclude that the number of *tcc*- or *tcs*-vertices at level at most ξ in \mathcal{A} , summed over all cells $\tau \in U_C^{(0)}$ and over all cylindrical surfaces C , is $O(\xi^3 F^{(ttn)}(n/\xi))$, where $F^{(ttn)}(m)$ is the maximum possible number of free *ttn*-vertices on the boundary of the union of the expansions K_s , for s in a set of at most m pairwise-disjoint triangles.

Next, suppose that the level of w_E with respect to \mathcal{A} is greater than ξ . This means that as we walk from a free *ttc*-vertex v_E on E within τ to w_E along E , we visit at least ξ vertices of \mathcal{A} , each of which has level at most ξ . We charge E to these ξ vertices of \mathcal{A} . Since each such vertex is charged only $O(1)$ times in this manner (because we only want to count m_τ), the total number of such elliptic arcs E is $O(\xi^2 F(n/\xi))$.

We have thus proved that the number of *ttc*-vertices that appear on interesting free triple intersections is

$$O\left(\xi n^{2+\varepsilon} + \xi^3 F^{(ttn)}\left(\frac{n}{\xi}\right) + \xi^2 F\left(\frac{n}{\xi}\right)\right). \tag{4.4}$$

A similar analysis proves the same bound on the number of *tts*-vertices that appear on interesting free triple intersections. We thus conclude the following.

Lemma 4.10. *For any parameter $\xi > 1$ and any $\varepsilon > 0$,*

$$O\left(\xi n^{2+\varepsilon} + \xi^3 F^{(ttn)}\left(\frac{n}{\xi}\right) + \xi^2 F\left(\frac{n}{\xi}\right)\right)$$

ttn-vertices appear on free interesting triple intersections.

Wrapping Up. Putting Lemmas 4.6, 4.7, and 4.10 together, we obtain the following recurrence, which is the same as in (4.2):

$$F^{(0)}(n) = O\left(\xi^{3-\varepsilon}n^{2+\varepsilon} + \xi^2 F\left(\frac{n}{\xi}\right) + \xi^3 F^{(tnn)}\left(\frac{n}{\xi}\right)\right), \tag{4.5}$$

$$F^{(tnn)}(n) = O\left(\xi^{3-\varepsilon}n^{2+\varepsilon} + \xi^2 F\left(\frac{n}{\xi}\right)\right). \tag{4.6}$$

As argued above, this completes the proof of Theorem 4.1.

5. Extensions

In this section we extend Theorem 2.1 to prove a near-quadratic bound on the complexity of the union of objects with bounded curvature in \mathbb{R}^3 and on the number of combinatorial changes in the union of moving congruent disks in the plane. We also discuss algorithms for computing the union of cylinders.

5.1. Objects with Bounded Curvature

Let $\mathcal{K} = \{K_1, \dots, K_n\}$ 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 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 C^2 -continuous and the *mean* curvature of any object at all points is at most some fixed constant κ .

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 $\varepsilon, \alpha, \kappa$, and on the maximum algebraic complexity of an object in \mathcal{K} .*

Proof (Sketch). We assume that the diameter of each object K_i is between 1 and α . Let V be the set of vertices on the union of \mathcal{K} . Choose a sufficiently small constant δ whose value will be specified later. We partition \mathbb{R}^3 into a grid \mathcal{C} of cubes, each of size δ (see Fig. 17), i.e.,

$$\mathcal{C} = \{[i\delta, (i + 1)\delta] \times [j\delta, (j + 1)\delta] \times [k\delta, (k + 1)\delta] \mid i, j, k \in \mathbb{Z}\}.$$

For each cube $C \in \mathcal{C}$, let $\mathcal{K}_C \subseteq \mathcal{K}$ be the set of objects that intersect C . Each K_i intersects $O(\alpha^3/\delta^3)$ cubes of \mathcal{C} , so $\sum_C \mathcal{K}_C = O(n)$. It is easily seen that only $O(n^2)$ vertices appear

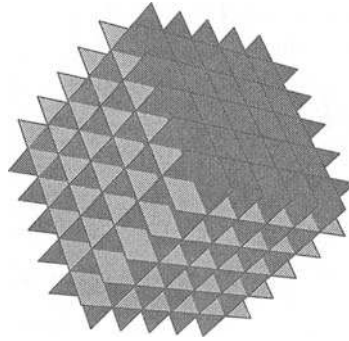


Fig. 17. Partitioning \mathbb{R}^3 into a grid of cubes.

on the boundary of any cube in \mathcal{C} , therefore it suffices to bound the number of vertices that lie in the interior of a cube.

Fix a cube $C \in \mathcal{C}$. Let $V_C \subseteq V$ be the set of vertices that lie in the interior of C . Let $\Sigma_C = \{\sigma_1, \dots, \sigma_u\}$, where $u = O(n)$, be the set of connected components of $(\partial K_i) \cap C$, for $i = 1, \dots, n$; each σ_i is a two-dimensional surface patch lying in the interior of C . Every vertex of V_C lies on three surface patches. Since each σ_i is C^2 -continuous and its curvature is bounded by κ , the normals of σ_i vary continuously and their directions lie inside a spherical cap of \mathbb{S}^2 of radius $c\kappa\delta$, for some constant $c > 0$.

We say that a direction $\rho \in \mathbb{S}^2$ is *good* for σ_i if each tangent line to σ_i makes an angle of at least $\beta\kappa\delta$ with ρ , for some constant $\beta > 0$; ρ is *bad* for a vertex $v \in V_C$ if it is bad for any of the three surfaces containing v . Since the normals of σ_i lie inside a spherical cap of radius $c\delta\kappa$, the bad directions for σ_i lie inside a spherical band consisting of all points in \mathbb{S}^2 that lie within distance $(\beta + c)\kappa\delta$ from a great circle. Hence, if we choose δ such that $\kappa\delta \ll 1$, then we can show, as in Section 2.5, that there exists a set $Z \subseteq \mathbb{S}^2$ of $O(1)$ points with the property that, for any vertex $v \in V_C$, there exists a direction $\rho \in Z$ that is good for v .

Let w, w' be two points on ∂K_i . Since K_i is convex and its mean curvature is at most κ , it follows that the sphere B_w of radius $1/\kappa$ and tangent to K_i at w from the inside is contained in K_i . If the direction $\overrightarrow{ww'}$ is good for σ_i , then $\overrightarrow{ww'}$ makes an angle of at least $\beta\kappa\delta$ with any line tangent to B_w at w . Since w' does not lie in the interior of B_w , $|ww'| \geq (2/\kappa) \sin(\beta\delta\kappa/2) \geq \beta\delta/2$, assuming that δ is a sufficiently small. If $|ww'| > \sqrt{3}\delta$, then both w and w' cannot lie in the same cube of \mathcal{C} . By choosing $\beta > 2\sqrt{3}$ we can guarantee that, for any point $w \in C \cap K_i$, the other intersection of the ray in a good direction from w does not lie in C . Now, following the same argument as in Section 2.5, one can reduce the problem of bounding $|V_C|$ to that of counting the number of vertices in the region lying between the ρ -upper and the ρ -lower envelopes of two respective subsets of Σ_C , summed over all $\rho \in Z$. Hence, $|V_C| = O(n^{2+\epsilon})$. This completes the proof of the theorem. \square

Remark 5.2. We can relax condition (iii) on C^2 -continuity. What we really need is that each object in \mathcal{K} intersects $O(1)$ cubes of \mathcal{C} and that, for each pair $C \in \mathcal{C}$, $K_i \in \mathcal{K}$, the normals of $C \cap K_i$ lie in a sufficiently small cap of \mathbb{S}^2 . For example, we

can obtain a quadratic bound on the union of convex polytopes that satisfy these two conditions.

5.2. Union of Moving Disks

Let $\mathcal{D} = \{D_1, \dots, D_n\}$ be a set of n unit-radius disks in the plane, each moving with a fixed velocity. That is, the position of the center \mathbf{c}_i of D_i is a linear function $\mathbf{c}_i(t) = \mathbf{a}_i + t\mathbf{b}_i$ of the time t , for some pair $\mathbf{a}_i, \mathbf{b}_i \in \mathbb{R}^2$. Let $U(t) = \bigcup_i D_i(t)$ denote the union \mathcal{D} at time t . We want to bound the number of changes in the combinatorial structure of $U(t)$ as t varies from $-\infty$ to $+\infty$.

For each $1 \leq i \leq n$, let K_i denote the *slanted cylinder*

$$K_i = \{(\mathbf{x}, t) \mid \mathbf{x} \in \mathbb{R}^2 \text{ and } d(\mathbf{x}, \mathbf{c}_i(t)) \leq 1\}.$$

(See Fig. 18.) The intersection of K_i with a plane normal to the axis of K_i , i.e., normal to the line $(\mathbf{a}_i + t\mathbf{b}_i, t)$ is an ellipse E_i whose major and minor semiaxes are 1 and $1/\sqrt{1 + \|\mathbf{b}_i\|^2}$, respectively. Set $U = \bigcup_{i=1}^n K_i$. $U(t)$ is the cross section of U at the plane $z = t$. The number of changes in the combinatorial structure of $U(t)$ is proportional to the combinatorial complexity of U .

Note that $U(t)$ is the cross section of the Euclidean Voronoi diagram of the point set $\{\mathbf{c}_i(t) \mid 1 \leq i \leq n\}$ in the sense discussed in the Introduction. Hence, the number of changes in $U(t)$ bounds the number of changes in the combinatorial structure of a cross section of the Voronoi diagram as the points move. The best known bound on the number of changes in the entire Voronoi diagram of a set of n points, each moving with fixed velocity, is near-cubic [6], [18], [24]. De Berg et al. [11] showed that if each K_i is a convex polygonal pseudodisk (i.e., each K_i is a convex polygon such that the boundaries of any pair always intersect in at most two points) moving with a fixed velocity, then the number of changes in their union is $O(n^2\alpha(n))$.

Without loss of generality, we can assume that the speed of all disks is at most 1. Then the minor semiaxis of each ellipse E_i is at least $1/\sqrt{2}$, and therefore the diameter and curvature of E_i are at most 2. By Remark 2.10, we can extend the proof of Theorem 2.1 to show that the combinatorial complexity of U is $O(n^{2+\epsilon})$. Hence, we can conclude the following.




Fig. 18. The “space-time” tracing of a set of moving disks.

Theorem 5.3. *Let \mathcal{D} be a set of n congruent disks in the plane, each moving with a fixed velocity. Then the number of combinatorial changes in their union is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$.*

5.3. Computing the Union

Let S be a set of n triangles in \mathbb{R}^3 with pairwise-disjoint interiors, and let B be a ball. The vertices, edges, and two-dimensional faces of the union of $\{s \oplus B \mid s \in S\}$ can be computed using the randomized incremental algorithm described by Agarwal et al. [1]. Basically, for each krepl K_i , their algorithm will compute the vertices, edges, and faces of U that lie on ∂K_i , by a straightforward incremental construction that inserts all the other K_j 's in a random order. Omitting all the details, which can be found in [1] (see also [2]), we conclude the following.

Theorem 5.4. *Let S be a set of n triangles in \mathbb{R}^3 with pairwise-disjoint interiors, and let B be a ball. The boundary of the union of the Minkowski sums $\{s \oplus B \mid s \in S\}$, can be computed in randomized expected $O(n^{2+\varepsilon})$ time, for any $\varepsilon > 0$.*

As mentioned in the Introduction, once the boundary ∂U is available, we can also compute the boundary of the free configuration space \mathcal{F} of B . We can then add artificial edges and vertices into \mathcal{F} so that all connected components of the boundary of any connected component of \mathcal{F} are connected. This can be done, using, for example, the technique by Sifrony and Sharir [25]. This step adds $O(n)$ additional vertices and edges. Then, given any two free placements Z_1, Z_2 of B , we can compute in $O(n)$ time, the placements W_1, W_2 that lie on ∂U immediately below (in the z -direction) Z_1 and Z_2 , respectively; here we are assuming that all connected components of \mathcal{F} are bounded. By locating W_1 and W_2 in the appropriate faces of $\partial \mathcal{F}$, we can then determine whether Z_1 and Z_2 lie in the same connected component of \mathcal{F} . That is, we can determine in $O(n)$ time whether B can be moved from Z_1 to Z_2 without intersecting any obstacle. If Z_1 and Z_2 lie in the same connected component, we can also compute a path from Z_1 to Z_2 that lies within \mathcal{F} . We do not know whether such a motion-planning query can be answered more efficiently, e.g., in polylogarithmic time.

6. Conclusion

In this paper we proved near-optimal (i.e., near-quadratic) bounds on the complexity of the free configuration space \mathcal{F} of a ball moving amid a set of polyhedral obstacles in \mathbb{R}^3 . We conclude by mentioning a few open combinatorial problems in this area. In each case the best known bound is cubic, and we conjecture the right bound to be near-quadratic.

- (i) What is the complexity of the Euclidean Voronoi diagram of a set of pairwise-disjoint polyhedral sites in \mathbb{R}^3 ? Even the case of line sites is still open.
- (ii) What is the complexity of the union of n cylinders of different radii?
- (iii) What is the complexity of the union of n congruent cubes in \mathbb{R}^3 ? What about n arbitrary cubes?

- (iv) What is the complexity of the union of n donuts, each being the Minkowski sum of a fixed ball with a circle, where the disks bounding the circles are assumed to be pairwise disjoint? (This problem was raised by Emo Welzl.)
- (v) In general, what is the complexity of the union of the Minkowski sums of a compact convex set B with n pairwise-disjoint compact convex sets A_1, \dots, A_n , under the assumption that the sets A_1, \dots, A_n, B all have “constant description complexity” (as defined, e.g., in [24])?
- (vi) What is the complexity of the dynamic Voronoi diagram of n moving points in the plane, where each point is moving at some fixed velocity?
- (vii) What is the complexity of the union of n “fat” tetrahedra? A tetrahedron is fat if the maximum aspect ratio of a face is a constant and the minimum dihedral angle is a constant.

Acknowledgments

The authors thank Boris Aronov, Alon Efrat, Leo Guibas, Dan Halperin, Sariel Har-Peled, and János Pach for helpful discussions concerning this problem. The authors thank an anonymous referee for reading the paper thoroughly and for the detailed comments that improved the quality of the paper.

Appendix. Proof of Lemma 4.2

We prove the lemma for the special case in which A_1 and A_2 are triangles and B is a ball. We assume that A_1 and A_2 are in general position as described in the beginning of Section 4.

Pick arbitrary points p_1 and p_2 in the relative interiors of triangles A_1 and A_2 , respectively, and assume, without loss of generality, that the center of B lies at the origin. For a parameter $t \in [0, 1]$, put

$$A_1(t) = tA_1 + (1-t)p_1, \quad A_2(t) = tA_2 + (1-t)p_2, \quad B(t) = tB,$$

and

$$K_1(t) = A_1(t) \oplus B(t), \quad K_2(t) = A_2(t) \oplus B(t).$$

Note that $K_1(t)$ and $K_2(t)$ are smooth for each $t > 0$.

We vary t from 0 to 1, and watch for topological changes in $C(t) = \partial K_1(t) \cap \partial K_2(t)$. Initially, $C(t) = C(0)$ is empty. $C(t)$ changes continuously as we vary t , so the number of connected components of $C(t)$ can change only when $K_1(t)$ and $K_2(t)$ are tangent to each other at some point. When this happens, either some component of $C(t)$ is a singleton point (when a new component has just appeared or an old component is about to vanish), or some component of $C(t)$ is not a simple closed curve (when two components of $C(t)$ are about to split or have just merged). It can be checked that if a component of $C(t)$ is not a singleton, then it is a simple closed curve. Therefore a connected component of $C(t)$ cannot split, or two components cannot merge. Hence, only a new component may appear or an existing component may disappear, as t varies.

Since $C(0)$ is empty, let t_0 be the minimum value of t at which $C(t)$ becomes a singleton. As is easily checked, $K_1(t_0)$ and $K_2(t_0)$ lie on the opposite sides of the (unique) plane supporting them at $C(t_0)$, and thus they have disjoint interiors. For any $t > t_0$, the interiors of $K_1(t)$ and $K_2(t)$ intersect. Suppose that $C(t')$, for some $t' > t_0$, has a new singleton component, call this point w . Let π be the common tangent plane to $K_1(t')$ and $K_2(t')$ at w . Without loss of generality, assume that π is parallel to the xy -plane. Since the interiors of $K_1(t')$ and $K_2(t')$ intersect, both of them lie on the same side of π , say below π . Then we can write w as

$$w = t'a_1 + (1 - t')p_1 + t'b,$$

where a_1 is a point on A_1 with the maximum z -coordinate, and b is the unique point on B with the maximum z -coordinate, and also as

$$w = t'a_2 + (1 - t')p_2 + t'b,$$

where a_2 is a point on A_2 with the maximum z -coordinate. We thus obtain

$$t'a_1 + (1 - t')p_1 = t'a_2 + (1 - t')p_2.$$

This however is impossible since $t'a_1 + (1 - t')p_1$ lies in A_1 and $t'a_2 + (1 - t')p_2$ lies in A_2 , and they are disjoint. This contradiction completes the proof of the lemma.

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, *Discrete Comput. Geom.* **22** (1999), 201–221.
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, Motion planning of a ball amid segments in three dimensions, *Proc. 10th ACM-SIAM Symp. Discrete Algorithms*, 1999, pp. 21–30.
5. P. K. Agarwal and M. Sharir, Arrangements and their applications, in *Handbook of Computational Geometry* (J. Sack and J. Urrutia, eds.), North-Holland, Amsterdam, 2000, pp. 49–119.
6. G. Albers, L. J. Guibas, J. S. B. Mitchell, and T. Roos, Voronoi diagrams of moving points, *Internat. J. Comput. Geom. Appl.* **8** (1998), 365–380.
7. 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), 391–406.
8. B. Aronov and M. Sharir, On translational motion planning of a convex polyhedron in 3-space, *SIAM J. Comput.* **26** (1997), 1785–1803.
9. B. Aronov, M. Sharir, and B. Tagansky, The union of convex polyhedra in three dimensions, *SIAM J. Comput.* **26** (1997), 1670–1688.
10. S. Basu, On the combinatorial and topological complexity of a single cell, *Proc. 39th IEEE Symp. Foundations of Computer Science*, 1998, pp. 606–616.
11. M. de Berg, H. Everett, and L. J. Guibas, The union of moving polygonal pseudodiscs—Combinatorial bounds and applications, *Comput. Geom. Theory Appl.* **11** (1998), 69–81.
12. 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.

13. 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.
14. K. Clarkson and P. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* **4** (1989), 387–421.
15. H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.
16. A. Efrat, G. Rote, and M. Sharir, On the union of fat wedges and separating a collection of segments by a line, *Comput. Geom. Theory Appl.* **3** (1993), 277–288.
17. A. Efrat and M. Sharir, The complexity of the union of fat objects in the plane, *Discrete Comput. Geom.* **23** (2000), 171–189.
18. J.-J. Fu and R. C. T. Lee, Voronoi diagrams of moving points in the plane, *Internat. J. Comput. Geom. Appl.* **1** (1991), 23–32.
19. A. Hernandez-Barrera, S. Nandy, and T. Asano, Translating a convex polyhedron over monotone polyhedra, unpublished manuscript, 2000.
20. 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.
21. M. van Kreveld, On fat partitioning, fat covering, and the union size of polygons, *Comput. Geom. Theory Appl.* **9** (1998), 197–210.
22. J. Matoušek, J. Pach, M. Sharir, S. Sifrony, and E. Welzl, Fat triangles determine linearly many holes, *SIAM J. Comput.* **23** (1994), 154–169.
23. M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* **12** (1994), 327–345.
24. M. Sharir and P. Agarwal, *Davenport–Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.
25. S. Sifrony and M. Sharir, A new efficient motion planning algorithm for a rod in two-dimensional polygonal space, *Algorithmica* **2** (1987), 367–402.
26. B. Tagansky, The Complexity of Substructures in Arrangements of Surfaces, Ph.D. Dissertation, Computer Science Department, Tel Aviv University, 1996.

Received July 6, 1999, and in revised form April 25, 2000. Online publication August 21, 2000.