

On the Union of κ -Round Objects in Three and Four Dimensions*

Boris Aronov[†] Alon Efrat[‡] Vladlen Koltun[§] Micha Sharir[¶]

July 6, 2005

Abstract

A compact set c in \mathbb{R}^d is κ -round if for every point $p \in \partial c$ there exists a closed ball that contains p , is contained in c , and has radius $\kappa \operatorname{diam} c$. We show that, for any fixed $\kappa > 0$, the combinatorial complexity of the union of n κ -round, not necessarily convex objects in \mathbb{R}^3 (resp., in \mathbb{R}^4) of constant description complexity is $O(n^{2+\varepsilon})$ (resp., $O(n^{3+\varepsilon})$) for any $\varepsilon > 0$, where the constant of proportionality depends on ε , κ , and the algebraic complexity of the objects. The bound is almost tight in the worst case.

1 Introduction

Given a set \mathcal{C} of n geometric objects in \mathbb{R}^d , let $\mathcal{U} = \mathcal{U}(\mathcal{C}) := \bigcup_{c \in \mathcal{C}} c$ denote their union, and let $\mathcal{A} = \mathcal{A}(\mathcal{C})$ denote the arrangement [39] of the (boundaries of the) objects in \mathcal{C} . The (combinatorial) *complexity* of \mathcal{U} is defined to be the number of faces of \mathcal{A} of all dimensions on the boundary $\partial \mathcal{U}$ of the union. The study of the complexity of the union of objects in

*A preliminary version of this paper has appeared in *Proc. 20th Annu. ACM Sympos. Comput. Geom.* (2004), 383–390. Work on this paper by B.A. and M.S. has been supported by a grant from the U.S.-Israeli Binational Science Foundation. Part of the work by B.A., V.K., and M.S. on this paper has been carried out while visiting the Mathematical Science Research Institute in Berkeley, California. Work by B.A. has also been supported by NSF Grant ITR CCR-00-81964. Work by A.E. has also been supported by NSF Grant CNS-03-12443 and NSF CAREER Grant 03-48000. Work by V.K. has also been supported by NSF Grant CCR-01-21555. Work by M.S. has also been supported by NSF Grants CCR-97-32101 and CCR-00-98246, by a grant from the Israel Science Fund (for a Center of Excellence in Geometric Computing), and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

[†]Department of Computer and Information Science, Polytechnic University, Brooklyn, NY 11201-3840, USA; <http://cis.poly.edu/~aronov>.

[‡]Department of Computer Science, University of Arizona, Tucson, AZ 85721-0077, USA; <http://www.cs.arizona.edu/people/alon/>.

[§]Computer Science Department, 353 Serra Mall, Gates 464, Stanford University, Stanford, CA 94305; <http://cs.stanford.edu/~vladlen>.

[¶]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; <http://www.cs.tau.ac.il/~michas>.

two dimensions has a long and rich history in computational and combinatorial geometry, starting with the results of Kedem *et al.* [29] and Edelsbrunner *et al.* [19], who have shown that, if the boundaries of any two distinct objects $c_1, c_2 \in \mathcal{C}$ intersect at most twice (resp., three times), the maximum possible complexity of \mathcal{U} is $\Theta(n)$ (resp., $\Theta(n\alpha(n))$, where $\alpha(\cdot)$ is the inverse Ackermann function). For the latter result to be meaningful, it is assumed that every $c \in \mathcal{C}$ is a region bounded between the x -axis and a Jordan arc whose endpoints lie on the x -axis, and where we only count intersections between these arcs.

When the object boundaries are allowed to intersect four or more times, the complexity of \mathcal{U} can easily reach $\Omega(n^2)$, for instance when the objects form a grid of narrow strips. To make such a construction possible, though, the objects have to be “long and skinny.” In an attempt to analyze the behavior of geometric objects in “real life” (e.g., the prevailing geometric objects encountered in computer graphics, vision, manufacturing and robotics applications), various classes of *realistic input models* were introduced (see, e.g., [14]), one of the most prominent among which is that of *fat objects*. The fat-object model addresses the observation that in some applications long and skinny objects are rarely encountered and objects with bounded aspect ratio are predominant. Under the assumption that the input objects are fat (see below for the several possible precise definitions), improved results were obtained for various algorithmic problems [3, 4, 7, 13, 22, 27, 28, 30, 33, 34, 40]. As this list of applications indicates, there is a strong practical motivation to study the complexity of the union of fat objects, in addition to the intrinsic interest in the problem itself.

The union of α -*fat wedges*, i.e., wedges whose opening angle is at least some constant $\alpha > 0$, has complexity $O(n)$ [7, 24]; here and hereafter the implied constants depend on the fatness parameters. The union of α -*fat triangles*, i.e., triangles all of whose angles are at least some positive constant α , has complexity $O(n \log \log n)$ [32, 37]. Extending the study into the realm of curved objects, Efrat and Sharir [23] have shown that the complexity of the union of n convex *fat objects* (i.e., convex objects for which the ratio between the radii of the smallest enclosing and the largest enclosed disks is bounded by a global constant), each of constant description complexity, is $O(n^{1+\varepsilon})$ for any $\varepsilon > 0$. (An object has *constant description complexity* if it is a semi-algebraic set defined as a Boolean combination of a constant number of polynomial equalities and inequalities of constant maximum degree in a constant number of variables.) Efrat and Katz [21] have then shown that the union of κ -*curved* (not necessarily convex) objects in the plane also has near-linear complexity. A planar object s is said to be κ -curved if, for any point $p \in \partial s$, there is a disk $d \subseteq s$ that is incident to p and has radius κ times the diameter of s . (It is this definition that we extend to higher dimensions and study in the present paper.) Finally, extending both results, Efrat [20] has introduced a further generalization of fatness, with the property that the union complexity remains near-linear. See also [9, 10, 36, 37] for other results concerning the union of objects in the plane.

Compared with this extensive research on the union of objects in \mathbb{R}^2 , the situation in three dimensions looks rather grim. Of course, the union of n thin plate-like objects that form a three-dimensional grid can have complexity $\Omega(n^3)$. However, as in the plane, it is important, and motivated by a number of practical applications, to consider realistic input models, and in particular to study the complexity of the union of *fat* objects. A prevailing conjecture

is that this complexity (under an appropriate definition of fatness) is near-quadratic. Such a bound has however proved quite elusive to obtain for general fat objects, and this has been recognized as one of the major open problems in computational and combinatorial geometry [18, Problem 4].

The complexity of the union of axis-parallel cubes is $O(n^2)$, and it drops to $O(n)$ if the cubes have the same size [15]. If the cubes are not axis-parallel but have equal (or “almost equal”) size, the complexity of their union is $O(n^{2+\varepsilon})$ for any $\varepsilon > 0$ [35]. A general sub-cubic bound on the complexity of the union of cubes in three dimensions is currently not known. Motivated by motion planning applications and the analysis of Voronoi diagrams, Aronov and Sharir [11] have proved a near-quadratic bound for the complexity of the union of Minkowski sums of disjoint convex polyhedra of overall complexity n with a common convex polyhedron of constant complexity. This refines a more general bound, obtained by Aronov *et al.* [12] for the case of the union of arbitrary convex polyhedra in 3-space. Guided by similar motivations, Agarwal and Sharir [6] have shown that the union of Minkowski sums of disjoint polyhedra of overall complexity n with a ball has complexity $O(n^{2+\varepsilon})$ for any $\varepsilon > 0$.

The only rather general result on the complexity of the union of fat objects in 3-space stems from the analysis technique of Agarwal and Sharir [6] and appears in their paper: The complexity of $\mathcal{U}(\mathcal{C})$ is $O(n^{2+\varepsilon})$ for any $\varepsilon > 0$ if \mathcal{C} consists of n convex objects of near-equal size, with C^2 -continuous boundaries, bounded mean curvature, and constant description complexity.

In $d \geq 4$ dimensions, the results become even more scarce. The complexity of the union of n halfspaces (each bounded by a hyperplane) in \mathbb{R}^d is $O(n^{\lfloor d/2 \rfloor})$, as follows from the Upper Bound Theorem. The complexity of the union of n balls in d -space is $O(n^{\lfloor d/2 \rfloor})$, as follows by lifting them to hyperplanes in \mathbb{R}^{d+1} . Boissonnat *et al.* [15] present an upper bound of $O(n^{\lfloor d/2 \rfloor})$ for the union of n axis-parallel cubes in \mathbb{R}^d , which improves to $O(n^{\lfloor d/2 \rfloor})$ when the cubes have equal size. The union complexity of n convex bodies in \mathbb{R}^d of constant description complexity *with a common interior point* is $O(n^{d-1+\varepsilon})$ for any $\varepsilon > 0$, which follows from the results of Sharir [38] on the complexity of upper envelopes of $(d-1)$ -variate functions (see also a refined bound for polyhedra in \mathbb{R}^3 in [26]). Finally, Koltun and Sharir [31] extended the above-mentioned result of Agarwal and Sharir [6] to four dimensions and proved that the complexity of the union of n convex objects of near-equal size, with C^2 -continuous boundaries, bounded mean curvature, and constant description complexity in \mathbb{R}^4 is $O(n^{3+\varepsilon})$ for any $\varepsilon > 0$.

Our results: We say that a set $c \subseteq \mathbb{R}^d$, $d \geq 3$, is an *object* if it is a compact connected set with nonempty interior. In a complete analogy with the definition of κ -curved objects in two dimensions, a compact object c is κ -*round* (for a fixed $\kappa > 0$) if for every point $p \in \partial c$ there exists a closed ball $B(p, c, \kappa)$ of radius $\kappa \operatorname{diam} c$, which contains p and is contained in c . We call $B(p, c, \kappa)$ a *witness ball* for c at p . If c is convex, $B(p, c, \kappa)$ is unique. The definition, though, allows c to be non-convex and to have *reflex* edges and vertices, although c cannot have any *convex* edge or vertex. See Figure 1. Recall that an object c has *constant description complexity* if it is a semi-algebraic set defined by a constant number of polynomial

equalities and inequalities of constant maximum degree in a constant number of variables. We refer to the largest of these constants as the *algebraic complexity* of c .




Figure 1: Examples of (planar analogues of) κ -round and non- κ -round objects.

Our main result is the following:

Theorem 1.1. *Let \mathcal{C} be a set of n κ -round (not necessarily convex) objects of constant description complexity in \mathbb{R}^3 or in \mathbb{R}^4 . Then the combinatorial complexity of $\mathcal{U}(\mathcal{C})$ is $O(n^{2+\varepsilon})$ in \mathbb{R}^3 and $O(n^{3+\varepsilon})$ in \mathbb{R}^4 , for any $\varepsilon > 0$, where the constant of proportionality depends on ε , κ , and the algebraic complexity of the objects in \mathcal{C} .*

The bound is nearly tight; for a construction see the concluding remarks.

We note that our analysis applies in any dimension $d \geq 3$, except for its last step, where we reduce the problem to that of bounding the number of vertices of the *sandwich region* [39] between the upper envelope of a collection of $(d - 1)$ -variate functions and the lower envelope of another such collection. Sharp bounds on the number of such sandwich vertices are known only for $d = 3$ and $d = 4$, which is the only reason for our present inability to extend Theorem 1.1 to $d > 4$.

Our analysis extends the results of Agarwal and Sharir [6] and of Koltun and Sharir [31], as it allows the objects to be non-convex and to have drastically different sizes.

We also note that our result implies that standard randomized divide-and-conquer techniques [2, 11, 12] can be used to construct the union of n κ -round objects with constant description complexity in 3-space in time $O(n^{2+\varepsilon})$ for any $\varepsilon > 0$.

2 The Complexity of the Union

2.1 Fixing a Good Direction

Let \mathcal{C} be a collection of n κ -round, not necessarily convex objects in \mathbb{R}^d , $d \geq 3$, of constant description complexity, and let $\mathcal{U} = \mathcal{U}(\mathcal{C})$ denote their union. In what follows, we estimate the combinatorial complexity of \mathcal{U} by the number of *vertices* of $\partial\mathcal{U}$, namely, the number of intersection points of d boundaries of objects of \mathcal{C} that lie on $\partial\mathcal{U}$. Let $V = V(\mathcal{C})$ denote the set of these vertices. We assume *general position* of the objects in \mathcal{C} , meaning, in particular, that no $d + 1$ boundaries have a point in common, and no vertex of the union is a seam point (see the discussion below) of any of the boundary surfaces of its incident objects. This involves no real loss of generality, since the maximum complexity of the union is attained for sets in general position, as follows, e.g., from the discussion in [39]. It is indeed sufficient to estimate the complexity of the union by the number of its vertices: Any face of $\partial\mathcal{U}$ that has a vertex can be charged to one of its vertices, and the general position assumption implies that no vertex is charged more than $2^d - 1$ times. The number of faces with no vertices can easily be shown to be $O(n^{d-1})$, by charging each such face f to the intersection of at most $d - 1$ boundaries, so that f is incident to, or coincides with, a connected component of that intersection; the number of faces obtained in this manner from a fixed set of boundaries is bounded by a function of d and of the algebraic complexity of the surfaces involved.

We recall and expand some additional notations. We consider objects in \mathbb{R}^d for any fixed $d \geq 3$. (We re-emphasize that most of our analysis applies to any $d \geq 3$, and we will present it in this generality.) For an object $c \subset \mathbb{R}^d$, $\text{diam } c$ denotes the *diameter* of c . Given $0 < \kappa \leq 1/2$, c is κ -*round* if through every point $p \in \partial c$ there exists a closed *witness ball* $B = B(p, c, \kappa)$ for c at p , which has radius $\kappa \text{diam } c$, contains p and is contained in c . If p is a *smooth* point of ∂c , $B(p, c, \kappa)$ is unique (see Lemma 2.1). The term *seam* of c refers to the set of all non-smooth points (*seam points*) on ∂c . Each seam is a finite union of algebraic arcs and singleton sets in three dimensions; in general it is a finite union of relatively open $(d - 2)$ -dimensional algebraic cells. For any seam point p we define $B(p, c, \kappa)$ to be one of the balls that meets the above conditions. Recall that our general position assumptions require that no vertex of $\mathcal{U}(\mathcal{C})$ lies in the seam of any incident boundary.

Henceforth, we fix d and κ . The following fact is immediate from definitions.

Lemma 2.1. *If p is a smooth point of ∂c , then there is a unique hyperplane $\pi = \pi(p, c)$ tangent to c at p and the ball $B = B(p, c, \kappa)$ is tangent to π at p and thus uniquely determined by c , $p \in \partial c$, and κ .*

Given a not necessarily smooth point $p \in \partial c$ and a fixed α , where $0 < \alpha < 2$, we say that direction \mathbf{n} is *good* (α -*good*, to be precise) for c and p if the (undirected) line $\ell(\mathbf{n}, p)$ through p in direction \mathbf{n} intersects the witness ball $B(p, c, \kappa)$ in a segment of length at least $\alpha \kappa \text{diam } c$, i.e., at least α times the radius of $B(p, c, \kappa)$. A direction that is not good is *bad*.

Lemma 2.2. *For a point p on the boundary of a κ -round object c , the measure of the set of bad directions for c and p can be upper bounded by an expression $\mu(\alpha)$ that depends only on α (and d) and not on the choice of p , c , or κ , and approaches zero as $\alpha \rightarrow 0$.*

Proof. As the definition of a bad direction is scale invariant, we scale c so that witness balls have unit radius. Fix a point $p \in \partial c$ and let $B = B(p, c, \kappa)$ be the corresponding unit witness ball. By definition of a bad direction, the line ℓ through p in that direction has to be close enough to being tangent to B at p , so that its intersection pq with B has length less than α . See Figure 2. Let θ be the angle between ℓ and the hyperplane tangent to B




Figure 2: A bad direction for c and p ; the cross section by the plane containing ℓ and the center of B is shown.

at p . We must have $\sin \theta < \alpha/2$ for the direction to be bad, so the bad directions lie in the band of half-width $\sin^{-1}(\alpha/2)$ about the great sphere of directions tangent to B at p . The $(d-1)$ -dimensional volume of this band is

$$\mu(\alpha) := v_{d-2} \int_{\pi/2 - \sin^{-1}(\alpha/2)}^{\pi/2 + \sin^{-1}(\alpha/2)} \sin^{d-2} \varphi d\varphi,$$

where v_{d-2} is the volume of the $(d-2)$ -dimensional unit sphere. Clearly, $\mu(\alpha)$ satisfies the properties asserted in the lemma. \square

Consider a vertex v of \mathcal{U} , incident to the boundaries of d objects $c_1, \dots, c_d \in \mathcal{C}$. As a consequence of Lemma 2.2, for each i , a random direction \mathbf{n} will be bad for v and c_i with probability at most $\mu(\alpha)$, so it will be good for v and all of the d incident boundaries (we will then say that \mathbf{n} is *good for v*) with probability at least $1 - d\mu(\alpha)$, which we can assume to be at least $\frac{1}{2}$, by choosing α sufficiently small. In other words, for this choice of α , the expected number of vertices in V for which \mathbf{n} is a good direction is at least $|V|/2$. Henceforth, we assume that \mathbf{n} is a fixed direction for which this property holds, and we proceed to establish the asserted upper bounds for the number of those vertices $v \in V$ for which \mathbf{n} is good; let V_0 be the set of these vertices. Without loss of generality, we take \mathbf{n} to be the positive x_d -direction and refer to it as *vertical*.

2.2 Decomposing Objects into Pillars

For each $c \in \mathcal{C}$ we decompose ∂c into $O(1)$ patches, such that the relative interior of each patch is smooth, and such that each patch is a semi-algebraic set of constant description complexity. We achieve this decomposition by constructing the *cylindrical algebraic decomposition* of c , as defined by Collins [17]. This decomposes c into $O(1)$ cells, referred to as

Collins cells, of dimensions $0, 1, \dots, d$, where each full-dimensional cell can be written as the set of points satisfying a conjunction of inequalities of the following form¹

$$\begin{aligned} \beta_1^- &< x_1 < \beta_1^+, \\ \beta_2^-(x_1) &< x_2 < \beta_2^+(x_1), \\ \beta_3^-(x_1, x_2) &< x_3 < \beta_3^+(x_1, x_2), \\ &\dots \\ \beta_d^-(x_1, \dots, x_{d-1}) &\leq x_d \leq \beta_d^+(x_1, \dots, x_{d-1}), \end{aligned} \tag{1}$$

where all the functions β_j^-, β_j^+ are smooth algebraic functions of constant description complexity. By the general position assumption, all vertices in V_0 lie on the top or bottom boundary of some full-dimensional cell of each of their incident objects. (We note that the number of cells in Collins' cylindrical algebraic decomposition is usually quite large, albeit constant under our assumptions. For our analysis, any decomposition of c into cells of the above form will do. For example, one may use the vertical decomposition method, described, e.g., in [39].)

We now construct the desired decomposition of ∂c as the collection of the top and bottom boundaries of all the full-dimensional Collins cells of c . That is, for each such Collins cell σ given by (1), we include in our decomposition the two patches σ^-, σ^+ , defined by replacing the last inequality in (1) with $x_d = \beta_d^-(x_1, \dots, x_{d-1})$ or $x_d = \beta_d^+(x_1, \dots, x_{d-1})$, respectively. To simplify the notation, we continue to refer to these cell boundaries as the *Collins cells* of ∂c . We distinguish between *top* Collins cells, which form the top boundaries of the original full-dimensional cells, and *bottom* cells, which form the bottom boundaries.

For any object $c \in \mathcal{C}$ and any smooth point $p \in \partial c$, we say that p is *good* if the vertical direction is good for p and c . Fix an object $c \in \mathcal{C}$, and let τ be a $(d-1)$ -dimensional Collins cell of ∂c . Let $G(\tau)$ denote the set of all good points p that lie in the relative interior of τ . Recall that being good means that the vertical line through $p \in G(\tau)$ enters c below p (if τ is a top Collins cell) or above p (if τ is a bottom Collins cell), and penetrates the (unique) corresponding witness ball B for a distance of at least α times its radius. Due to the constant description complexity of c and of τ , $G(\tau)$ consists of a constant number of connected components, each of constant description complexity. We view each of these components as the graph of a partial $(d-1)$ -variate function g .

Lemma 2.3. *Let g be one of the partial functions defined above, and let D denote its domain of definition. Then the graph of g is “not too steep,” in the sense that any pair of points $x, x' \in D$ satisfy $|g(x) - g(x')| \leq 2\omega \text{diam } D/\alpha$, where $\omega \geq 1$ depends only on d and on the algebraic complexity of D .*

Proof. By construction, the graph $\Gamma = \Gamma(g)$ of g is a portion of a smooth algebraic surface delimited by a constant number of smooth algebraic arcs in three dimensions, or, in higher dimensions, by a constant number of lower-dimensional smooth algebraic patches of constant description complexity. At a point p of its relative interior, Γ has a well-defined tangent

¹The weak inequalities for x_d constitute a slight variation of the traditional definition, made to simplify the presentation.

hyperplane that is “not too vertical,” as it forms an angle $\beta \geq \sin^{-1}(\alpha/2)$ with the vertical direction. This fact continues to hold at the boundary of Γ , with the natural definition of a limit tangent hyperplane. In particular, the slope of any smooth curve on the closure $\bar{\Gamma}$ of Γ at any point p (i.e., the value of $\frac{dx_d}{ds}$ at p , where s is the arc length) is at most

$$\cot \beta \leq \cot(\sin^{-1}(\alpha/2)) = \frac{\sqrt{1 - (\alpha/2)^2}}{\alpha/2} \leq \frac{2}{\alpha}.$$

Connect the given pair of points $x, x' \in D$ by a shortest path $\gamma \subset D$. It is a concatenation of algebraic arcs and straight line segments; their number is at most a constant that depends on the algebraic complexity of D . Further subdivide these algebraic arcs into subarcs that are monotone in all coordinates. Again, the number of resulting subarcs is bounded by a function of the algebraic complexity of D . Each subarc fits into a cube of edge length $\text{diam } D$. Therefore its L_1 -length, and therefore its Euclidean length too, is at most $(d - 1) \text{diam } D$. Hence the Euclidean length $|\gamma|$ of γ is at most $\omega \text{diam } D$, where ω is a constant that depends only on d and on the algebraic complexity of D .

To estimate $|g(x) - g(x')|$, we integrate along γ . “Lifting” each subarc of γ to Γ produces a piecewise smooth curve of slope at most $2/\alpha$, hence $|g(x) - g(x')| \leq 2|\gamma|/\alpha \leq 2\omega \text{diam } D/\alpha$, as claimed. \square

Lemma 2.4. *The good portion $G(\tau)$ of any Collins cell τ of ∂c can be subdivided into patches, so that*

- (i) *The x_d -variation of each patch (i.e., the length of its orthogonal projection to the x_d -axis) is at most $\alpha\kappa \text{diam } c/10$.*
- (ii) *The diameter of the projection of each patch to the hyperplane $x_d = 0$ is at most $\alpha^2\kappa \text{diam } c/(20\omega)$, where ω is the parameter provided in Lemma 2.3.*
- (iii) *The number of patches is bounded by a function of α, κ, d , and the algebraic complexity of the objects in \mathcal{C} only.*

Proof. Consider one of the (constant number of) connected components of $G(\tau)$ and let D denote its domain. D is a semi-algebraic set of constant description complexity. The statement in the proposition is scale invariant, so we assume D has diameter one. Overlay D with a $(d - 1)$ -dimensional orthogonal grid of hyperplanes with step $\alpha^2\kappa/(20\omega\sqrt{d - 1})$, shifted if necessary to avoid degeneracies. Define a family of sets δ_i that are the closures of the connected components of the intersections of the interior of D with the grid cells. Each δ_i is a semi-algebraic set of constant description complexity, with algebraic complexity comparable with that of D . Define a family of patches π_i , such that each π_i is the collection of points on $G(\tau)$ over δ_i . Each grid cell has diameter $\alpha^2\kappa/(20\omega)$ and thus, by Lemma 2.3, the x_d -coordinates of points on a patch π_i vary by at most $\alpha\kappa/10$. The number of patches is no larger than $O((\alpha^2\kappa/(20\omega\sqrt{d - 1}))^{-d+1})$, a function of α, κ, d , and the algebraic complexity of the objects in \mathcal{C} . Therefore, the union of the above families of patches, over all domains D , meets the conditions listed in the statement of the lemma. \square

Corollary 2.5. *There exists a constant t , depending only on κ , α , d , and the algebraic complexity of the objects of \mathcal{C} , so that, for any $c \in \mathcal{C}$ and for any top Collins cell τ of ∂c , there exists a collection of connected objects c_1, c_2, \dots, c_t in \mathbb{R}^d , referred to as top pillars, such that the following properties hold:*

- (i) *Each top pillar $c_i \subset c$ consists of all the points that lie below a patch of a top portion of $G(\tau)$, as provided in Lemma 2.4, and above some horizontal hyperplane. The portion of ∂c_i on ∂c (resp., the bottom hyperplane, the vertical boundary) is referred to as its top cap (resp., bottom flat, vertical sides). See Figure 3 for an illustration.*




Figure 3: A top pillar in three dimensions. Features of its top cap are highly exaggerated. The base of the pillar need not be a full square in general.

- (ii) $\bigcup_i \partial c_i \cap \partial c = G(\tau)$.
- (iii) *The x_d -variation of the top cap $c_i \cap \partial c$, for any $1 \leq i \leq t$, is at most $\alpha\kappa \text{ diam } c/10$.*
- (iv) *The vertical distance from the top cap of each c_i to its bottom flat is $2\alpha\kappa \text{ diam } c/5$. The total height of a pillar is thus between $2/5$ and $1/2$ of $\alpha\kappa \text{ diam } c$.*
- (v) *The diameter of the projection of c_i to the hyperplane $x_d = 0$ is at most $\alpha^2\kappa \text{ diam } c/(20\omega)$.*
- (vi) *The slope of any line connecting two points, one on the top cap and one on the bottom flat of c_i , is at least $8\omega/\alpha$.*

A symmetric statement holds when τ is a bottom Collins cell of ∂c , where the corresponding bottom pillars c'_i have symmetric properties when we reverse the direction of the x_d -axis.

Proof. The construction is carried out as follows: Lemma 2.4 provides a subdivision of $G(\tau)$ into patches that will serve as the top caps of the pillars. Consider such a patch π_i . The pillar c_i consists of points that lie vertically below π_i and above a horizontal base hyperplane lying at vertical distance $2\alpha\kappa \text{ diam } c/5$ below the lowest point of π_i . Thus the pillars satisfy (i) and (ii) by construction. By Lemma 2.4, the vertical span of each π_i is $\alpha\kappa \text{ diam } c/10$, and the diameter of its projection to the hyperplane $x_d = 0$ is $\alpha^2\kappa \text{ diam } c/(20\omega)$. These

properties, and the construction, imply properties (iii)–(v). The top pillar c_i is contained in c since each point of π_i is good. Finally, (vi) follows, since the vertical distance between a point on the top cap and on the bottom flat of a pillar is at least $2\alpha\kappa \text{diam } c/5$ by (iv), while their horizontal distance is at most $\alpha^2\kappa \text{diam } c/(20\omega)$ by (v), yielding a slope of at least $8\omega/\alpha$. \square

Recall that we now consider only the subset V_0 of those vertices in V for which the vertical direction is good. Any such vertex necessarily lies on the good part of the top or bottom boundary of each of the d objects that it is incident to.

We next consider the collection P^+ of all top pillars c_i of all Collins cells of all original objects $c \in \mathcal{C}$, and the analogous collection P^- of bottom pillars, and put $P := P^+ \cup P^-$. Any vertex $q \in V_0$ is also a vertex of the union $\mathcal{U}' := \bigcup P$, so it suffices to bound the complexity of \mathcal{U}' , or, more precisely, to bound the number of vertices of \mathcal{U}' that lie on the top or bottom cap of each of the pillars they are incident to. In fact, we can restrict our attention further to those vertices of \mathcal{U}' that appear on d pillar caps and *lie on the boundary* of the original union \mathcal{U} . Clearly, these are the only vertices of concern at this point.

For each pillar $\pi \in P$, consider the projection $h(\pi)$ of π to the x_d -axis. We obtain a set H of intervals on the x_d -axis, and construct a so-called *hereditary segment tree* T on H , as defined in [16]. Each node w of T represents an interval I_w along the x_d -axis, and stores a list L_w of *long* pillars π , whose projection to the x_d -axis contains I_w but does not fully contain I_{w_0} , where w_0 is the parent of w , and a list S_w of *short* pillars, which are long in some proper descendant of w ; that is, their x_d -span only partially overlaps I_w .

Any vertex $q \in V_0$ of the type under consideration is incident to the caps of d pillars, all stored as long in d nodes of T (not necessarily distinct) that lie on a common path to the root, namely, the path from the leaf whose associated interval contains the x_d -coordinate of q . Let w be the highest of these nodes (i.e., the closest to the root). We count q at w . More precisely, the subproblem that we solve at w is to bound the number of vertices q of the union of $L_w \cup S_w$, such that

- (i) q lies in the caps of all incident pillar boundaries,
- (ii) the x_d -coordinate of q lies in I_w ,
- (iii) q is incident to at least one long pillar in L_w , and
- (iv) $q \in \partial\mathcal{U}$.

Deriving this bound forms the focus of the following subsection. The sum of these bounds, over all nodes w of T , yields a bound for $|V_0|$ and is thus the quantity we wish to bound.

2.3 Bounding the Union Complexity of Pillars

Let σ_w denote the horizontal slab $\mathbb{R}^{d-1} \times I_w$. Suppose that one of the pillars incident to q is a long top pillar π in L_w whose cap contains q . By construction, any point in σ_w lying vertically below the cap of π is contained in π . Hence, q is a point on the upper envelope of

the caps of the top pillars in L_w . Symmetrically, if π' is a long bottom pillar in L_w whose cap contains q , then q is a point on the lower envelope of the caps of the bottom pillars in L_w .

Suppose next that q is incident to a short top pillar π' in S_w whose top cap is incident to q . Since π' is short, its bottom flat may lie in the interior of σ_w . A similar situation may arise when π' is a short bottom pillar whose bottom cap is incident to q . Therefore, at first sight it appears that short pillars cannot be handled by just considering vertices that lie on their upper or lower envelopes. However, the following two lemmas show that the relevant vertices *can* be captured along envelopes of certain “canonical” subsets of short pillars.

Lemma 2.6. *Let π' be a short top pillar, whose top cap is incident to a vertex q of the union \mathcal{U} that also lies on the top cap of a long top pillar π . Then each point in σ_w vertically below π' is contained in the interior of the original union \mathcal{U} and thus cannot contain a vertex of \mathcal{U} .*

Symmetrically, if π' is a short bottom pillar, whose bottom cap is incident to a vertex of the union that also lies on the bottom cap of a long bottom pillar, then each point in σ_w vertically above π' is contained in the original union \mathcal{U} .

Proof. It is enough to address the case of top pillars. Let $s \in \sigma_w$ be a point that lies vertically below (the flat bottom of) π' . Then the angle φ between qs and the x_d -axis is rather small. Specifically, by Corollary 2.5(vi), we have $\tan \varphi < \alpha/(8\omega) < \alpha/8$ (since $\omega \geq 1$). Refer to Figure 4.




Figure 4: The figure depicts a short top pillar π' whose top cap is incident to a vertex q of the union \mathcal{U} that also lies on the top cap of a long top pillar π . Every point s lying within σ_w below π' must be contained in the interior of \mathcal{U} .

Let c be the object that contains π . Consider the witness ball $B = B(q, c, \kappa)$ for c at q . We claim that $s \in B$. Indeed, since π is long in w , we have, by Corollary 2.5(iv), $|I_w| \leq \alpha\kappa \text{diam } c/2$. Hence the difference between the x_d -coordinates of q and s is less than $\alpha/2$ times the radius of B . Recall that q lies on $G^+(c)$ and thus it is the top endpoint of a vertical chord of B of length at least α times the radius. It therefore follows that s lies above the center of B . Refer to Figure 5, which depicts the two-dimensional vertical cross section of B through q and the center o of B . Suppose that the length of the vertical chord of B through q is exactly β times the radius of B , $\beta \geq \alpha$, and use the notation in the figure




Figure 5: Illustrating the proof that the point s lies inside B . The figure depicts a vertical 2-dimensional cross section of B through o and q ; s does not necessarily lie in that cross section.

to conclude that

$$\tan \angle bqa = \frac{1 - \sqrt{1 - (\beta/2)^2}}{\beta/2} > \beta/4 \geq \alpha/4.$$

However, we have shown that $\tan \angle sqa < \alpha/8$. This, and the fact that s lies above o , is easily seen to imply that s lies in the interior of B , as asserted, which clearly completes the proof of the lemma. \square

Remark. If, as depicted in Figure 4, π' “hangs over” the vertical boundary of π , the point s may lie outside the union of the pillars (e.g., when π is an extreme pillar of c). However, s still lies within the interior of the union \mathcal{U} of the original objects, as we have just shown.

Lemma 2.7. *Let π' be a short bottom pillar, whose bottom cap is incident to a vertex q of \mathcal{U} that also lies on the top cap of a long top pillar π , and let s be a point in σ_w vertically above π' . Then s cannot lie on the top cap of any long top pillar.*

Symmetrically, let π' be a short top pillar, whose top cap is incident to a vertex q of \mathcal{U} that also lies on the bottom cap of a long bottom pillar, and let s be a point in σ_w vertically below π' . Then s cannot lie on the bottom cap of any long bottom pillar.

Proof. It is sufficient to consider only the former scenario. The proof is very similar to that of the preceding lemma. Suppose to the contrary that s does lie on the top cap of another long pillar π'' , contained in an original object $c'' \in \mathcal{C}$. See Figure 6. One can show, arguing in much the same way as above, that the angle φ between sq and the x_d -direction satisfies $\tan \varphi < \alpha/8$. Let $B = B(s, c'', \kappa)$ be the witness ball for c'' at s . Then, repeating the calculation illustrated in Figure 5, one concludes that q must lie in the interior of B , which is impossible. \square

In view of Lemma 2.6, we can take each short top pillar π' , whose top cap is incident to a vertex that also lies on the top cap of a long top pillar, and extend it all the way to the bottom of σ_w , without covering any vertex of the union under consideration. Symmetrically, any short bottom pillar whose bottom cap is incident to a vertex that also lies on the bottom cap of a long bottom pillar, can be extended all the way to the top of σ_w , without




Figure 6: No point s within σ_w can lie (a) above a short bottom pillar π' whose bottom cap is incident to a vertex q of the union \mathcal{U} that also lies on the top cap of a long top pillar π , and (b) on the top cap of another long top pillar. Only a portion of the long pillar π is depicted; the full pillar crosses the slab from top to bottom.

covering any vertex of the union under consideration. Exploiting Lemma 2.7 is done a little differently—see below.

Denote by L_w^+ the set of long top pillars in L_w , and by L_w^- the set of long bottom pillars in L_w . Denote by S_w^+ the set of extended short top pillars in S_w , and by S_w^- the set of extended short bottom pillars in S_w ; only those short pillars that satisfy the conditions in Lemma 2.6 are included in S_w^+ and S_w^- .

Let $q \in \sigma_w$ be a vertex of the union that lies on the cap of at least one long pillar in L_w . We consider the following cases:

Case A: q is incident only to top caps of (long or extended short) top pillars. In this case, the preceding analysis implies that q is a vertex of the upper envelope of the top caps of the pillars in $L_w^+ \cup S_w^+$. Symmetrically, if q is incident only to bottom caps of (long or extended short) bottom pillars, then q is a vertex of the lower envelope of the bottom caps of the pillars in $L_w^- \cup S_w^-$.

Case B: q is incident to at least one top cap of a long top pillar in L_w^+ , and to at least one bottom cap of a long bottom pillar in L_w^- . Here, by Lemma 2.6, any short pillar incident to q can be extended up or down, as appropriate, which is easily seen to imply that q is a vertex of the *sandwich region* [5, 39] between the upper envelope of the caps of pillars in $L_w^+ \cup S_w^+$ and the lower envelope of the caps of pillars in $L_w^- \cup S_w^-$.

Case C: q is incident to at least one top cap of a long top pillar in L_w^+ , to *no* bottom cap of any long bottom pillar in L_w^- , and to some top or bottom caps of short pillars. Let V^* denote the set of these vertices. If q is incident to the bottom caps of some short bottom pillars, then, in view of Lemma 2.7, we can extend any such pillar π' upwards to the ceiling of σ_w , without covering any other vertex of V^* . Let \hat{S}_w^- denote the set of short bottom pillars π' that contain such a point q on their bottom caps. Hence, in this case q is a vertex of the sandwich region between the upper envelope of the top caps of pillars in $L_w^+ \cup S_w^+$ and the lower envelope of the bottom caps of pillars in \hat{S}_w^- . The symmetric case, in which the roles of L_w^+ and L_w^- are interchanged, is analyzed in much the same way, and implies that q is a

vertex of the sandwich region between the lower envelope of the bottom caps of pillars in $L_w^- \cup S_w^-$ and the upper envelope of the top caps of pillars in an analogously defined set \hat{S}_w^+ .

Hence, in all three cases, q is a vertex of either an upper envelope, or a lower envelope, or the sandwich region between two envelopes, of caps of pillars in certain combinations of the sets $L_w^+, L_w^-, S_w^+, S_w^-, \hat{S}_w^+, \hat{S}_w^-$. As argued above, the caps of the pillars in these sets all have constant description complexity. Hence, using the results in [1, 25, 31, 38], it follows that, for $d = 3$ or 4, the number of vertices q under consideration inside σ_w is

$$O\left((|L_w^+| + |L_w^-| + |S_w^+| + |S_w^-| + |\hat{S}_w^+| + |\hat{S}_w^-|)^{d-1+\varepsilon}\right)$$

for any $\varepsilon > 0$, where the constant of proportionality depends on ε, κ, d , and on the algebraic complexity of the objects in \mathcal{C} . (Note that this is the only stage in the analysis where we restrict the dimension d .)

We sum this bound over all nodes w of the segment tree T . By the properties of hereditary segment trees [16], we have

$$\sum_w (|L_w^+| + |L_w^-| + |S_w^+| + |S_w^-| + |\hat{S}_w^+| + |\hat{S}_w^-|) = O(|P| \log |P|),$$

where P is, as above, the set of all pillars. By construction, $|P| = O(n)$, where the constant of proportionality depends on κ, d , and the algebraic complexity of the objects in \mathcal{C} . This implies that

$$\sum_w (|L_w^+| + |L_w^-| + |S_w^+| + |S_w^-| + |\hat{S}_w^+| + |\hat{S}_w^-|)^{d-1+\varepsilon} = O(n^{d-1+\varepsilon}),$$

for any $\varepsilon > 0$, and so the proof of Theorem 1.1 is complete. \square

Remarks. (1) As already mentioned, the analysis holds in any dimension, except for the lack of a sharp bound on the complexity of the sandwich region between envelopes in five and higher dimensions. The availability of such a bound would immediately imply the extension of Theorem 1.1 to the respective dimension.

(2) The bound in Theorem 1.1 is nearly tight, in the following sense. For any $\varepsilon > 0$, there exists a family of n κ -round objects in \mathbb{R}^d , $d \geq 3$, of constant description complexity, whose union has complexity $\Omega(n^{d-1})$, with $\kappa = 1/(2(1+\varepsilon))$ and implied constant *independent* of ε . Here is one such construction. Put $m := n/(d-1)$ and $h := \sqrt{d-2}$. Consider a $(d-1)$ -dimensional grid in the hyperplane $x_d = 0$, formed by $d-1$ pairwise orthogonal families, each consisting of m parallel $(d-2)$ -flats $1/(m-1)$ apart. Truncate the grid to within the $(d-1)$ -cube $[0, 1]^{d-1} \times \{0\}$; it consists of $d-1$ pairwise orthogonal families of parallel $(d-2)$ -cubes. Let B be a d -ball of radius $h/(2\varepsilon)$ centered at the origin and consider the family of objects obtained by taking the Minkowski sum of each $(d-2)$ -cube with B . Each such object c has diameter $h + h/\varepsilon$ which is the sum of diameters of B and of the $(d-2)$ -cube that produced it. Moreover, by construction, every point $p \in \partial c$ has a translated copy of B contained in c and touching p , so the roundness factor κ of c is $(h/2\varepsilon)/(h + h/\varepsilon) = 1/(2(1+\varepsilon))$, as claimed. Moreover, all objects are tangent to the hyperplane $x_d = h$ and their intersections with it form a $(d-1)$ -dimensional grid of complexity $\Theta(m^{d-1}) = \Theta(n^{d-1})$. Thus in particular the union

of the objects has complexity $\Omega(n^{d-1})$. Degeneracies can be removed by slightly perturbing the individual objects without reducing the complexity of the union. This construction is closely related to the one presented in [8].

Acknowledgments

The authors would like to thank Carola Wenk for fruitful discussions.

References

- [1] P.K. Agarwal, O. Schwarzkopf, and M. Sharir, The overlay of lower envelopes and its applications, *Discrete Comput. Geom.* 15 (1996), 1–13.
- [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, E.F. Grove, T.M. Murali, and J.S. Vitter, Binary space partitions for fat rectangles, *SIAM J. Comput.* 29 (2000), 1422–1448.
- [4] P.K. Agarwal, M. Katz, and M. Sharir, Computing depth order and related problems, *Comput. Geom. Theory Appl.* 5 (1995), 187–206.
- [5] P.K. Agarwal and M. Sharir, Arrangements of surfaces in higher dimensions, in *Handbook of Computational Geometry*, J.R. Sack and J. Urrutia (Eds.), North-Holland, 2000, pp. 49–119.
- [6] P.K. Agarwal and M. Sharir, Pipes, cigars, and kreplach: The union of Minkowski sums in three dimensions, *Discrete Comput. Geom.* 24 (2000), 645–685.
- [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, A lower bound on Voronoi diagram complexity, *Inform. Process. Lett.* 83 (2002), 183–185.
- [9] B. Aronov, A. Efrat, D. Halperin, and M. Sharir, A subquadratic bound on the number of regular vertices of the union of Jordan regions, *Discrete Comput. Geom.* 25 (2001), 203–220.
- [10] B. Aronov and M. Sharir, The common exterior of convex polygons in the plane, *Comput. Geom. Theory Appl.* 8 (1997), 139–149.
- [11] B. Aronov and M. Sharir, On translational motion planning of a convex polyhedron in 3-space, *SIAM J. Comput.* 26 (1997), 1785–1803.

- [12] B. Aronov, M. Sharir, and B. Tagansky, The union of convex polyhedra in three dimensions, *SIAM J. Comput.* 26 (1997), 1670–1688.
- [13] M. de Berg, Linear size binary space partitions for uncluttered scenes, *Algorithmica* 28 (2000), 353–366.
- [14] M. de Berg, M.J. Katz, A.F. van der Stappen, and J. Vleugels, Realistic input models for geometric algorithms, *Algorithmica* 34 (2002), 81–97.
- [15] 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), 473–484.
- [16] B. Chazelle, H. Edelsbrunner, L. Guibas, and M. Sharir, Algorithms for bichromatic line segment problems and polyhedral terrains, *Algorithmica* 11 (1994), 116–132.
- [17] G.E. Collins, Quantifier elimination for real closed fields by cylindrical algebraic decomposition, In *Proceedings of the Second GI Conference on Automata Theory and Formal Languages*, 134–183. Springer Lecture Notes in Computer Science 33, 1975.
- [18] E.D. Demaine, J.S.B. Mitchell, and J. O’Rourke, The Open Problems Project, <http://www.cs.smith.edu/~orourke/TOPP/>.
- [19] H. Edelsbrunner, L.J. Guibas, J. Pach, R. Pollack, R. Seidel, and M. Sharir, Arrangements of curves in the plane: Topology, combinatorics, and algorithms, *Theoret. Comput. Sci.* 92 (1992), 319–336.
- [20] A. Efrat, The complexity of the union of (α, β) -covered objects, *SIAM J. Comput.*, 34 (2005), 775–787.
- [21] A. Efrat and M. Katz, On the union of κ -curved objects, *Comput. Geom. Theory Appl.* 14 (1999), 241–254.
- [22] A. Efrat, M.J. Katz, F. Nielsen, and M. Sharir, Dynamic data structures for fat objects and their applications, *Comput. Geom. Theory Appl.* 15 (2000), 215–227.
- [23] A. Efrat and M. Sharir, The complexity of the union of fat objects in the plane, *Discrete Comput. Geom.* 23 (2000), 171–189.
- [24] 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.
- [25] D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions with applications to visibility of terrains, *Discrete Comput. Geom.* 12 (1994), 313–326.
- [26] D. Huttenlocher, K. Kedem, and M. Sharir, The upper envelope of Voronoi surfaces and its applications, *Discrete Comput. Geom.* 9 (1993), 267–291.
- [27] M. Katz, 3-D vertical ray shooting and 2-D point enclosure, range searching, and arc shooting amidst convex fat objects, *Comput. Geom. Theory Appl.* 8 (1997), 299–316.

- [28] M. Katz, M. Overmars, and M. Sharir, Efficient output sensitive hidden surface removal for objects with small union size, *Comput. Geom. Theory Appl.* 2 (1992), 223–234.
- [29] 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.
- [30] M. van Kreveld, On fat partitioning, fat covering, and the union size of polygons, *Comput. Geom. Theory Appl.* 9 (1998), 197–210.
- [31] V. Koltun and M. Sharir, The partition technique for overlays of envelopes, *SIAM J. Comput.* 32 (2003), 841–863.
- [32] J. Matoušek, J. Pach, M. Sharir, S. Sifrony, and E. Welzl, Fat triangles determine linearly many holes, *SIAM J. Comput.* 23 (1994), 154–169.
- [33] M.H. Overmars, Point location in fat subdivisions, *Inform. Process. Lett.* 44 (1992), 261–265.
- [34] M.H. Overmars and A.F. van der Stappen, Range searching and point location among fat objects, *J. Algorithms* 21 (1996), 629–656.
- [35] J. Pach, I. Safruti, and M. Sharir, The union of congruent cubes in three dimensions, *Discrete Comput. Geom.* 30 (2003), 133–160.
- [36] J. Pach and M. Sharir, On the boundary of the union of planar convex sets, *Discrete Comput. Geom.* 21 (1999), 321–328.
- [37] J. Pach and G. Tardos, On the boundary complexity of the union of fat triangles, *SIAM J. Comput.* 31 (2003), 1745–1760.
- [38] M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* 12 (1994), 327–345.
- [39] M. Sharir and P.K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.
- [40] A.F. van der Stappen, D. Halperin, and M.H. Overmars, The complexity of the free space for a robot moving amidst fat obstacles, *Comput. Geom. Theory Appl.* 3 (1993), 353–373.