

State of the Union (of Geometric Objects): A Review ^{*}

Pankaj K. Agarwal[†]

János Pach[‡]

Micha Sharir[§]

July 9, 2007

Abstract

Let \mathcal{C} be a set of geometric objects in \mathbb{R}^d . The combinatorial complexity of the union $\mathcal{U}(\mathcal{C})$ of \mathcal{C} is the total number of faces of all dimensions, of the arrangement of the boundaries of the objects, which lie on its boundary. We survey the known upper bounds on the complexity of the union of n geometric objects satisfying various natural conditions. These problems play a central role in the design and analysis of many geometric algorithms arising in robotics, molecular modeling, solid modeling, and shape matching, and the techniques used for their solutions are interesting in their own right.

1 Introduction

Let $\mathcal{C} = \{C_1, \dots, C_n\}$ be a set of n geometric objects, such as disks or convex polygons in the plane, or balls, cylinders, or convex polyhedra in three and higher dimensions. Let $\mathcal{U}(\mathcal{C}) = \bigcup_{i=1}^n C_i$ denote the union of the objects in \mathcal{C} . The combinatorial complexity (or complexity for brevity) of $\mathcal{U}(\mathcal{C})$ is the number of faces of the arrangement of the boundaries of the objects, which lie on the boundary of the union; see below for a formal definition. Several combinatorial and algorithmic problems in a wide range of applications, including linear programming, robotics, solid modeling, molecular modeling, and geographic information systems, can be formulated as problems that seek to calibrate the complexity of the union of a set of objects, or to compute their union. We begin by reviewing some of these applications.

^{*}Work on this paper was partially supported by a grant from the U.S.-Israel Binational Science Foundation. Work by J.P. and M.S. was also supported by NSF Grant CCF-05-14079. Work by P.A. was also supported by NSF under grants CCR-00-86013, EIA-01-31905, CCR-02-04118, and DEB-04-25465, by the ARO grant W911NF-04-1-0278. Work by J.P. was also supported by grants from PSC-CUNY Research Foundation and OTKA (Hungarian National Research Foundation). Work by M.S. was also supported by grant 155/05 from the Israel Science Fund, Israeli Academy of Sciences, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

[†]Department of Computer Science, Duke University, Durham NC 27708, USA. E-mail: pankaj@cs.duke.edu

[‡]City College, CUNY, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA. E-mail: pach@cims.nyu.edu

[§]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. E-mail: michas@post.tau.ac.il




Figure 1. An instance of two-dimensional linear programming: (a) The shaded region denotes the feasible region $\bigcap_{i=1}^n \bar{C}_i$; (b) The shaded region denotes $\bigcup_{i=1}^n C_i$.

Linear programming. Given a family $\bar{\mathcal{C}} = \{\bar{C}_1, \dots, \bar{C}_n\}$ of n halfspaces in \mathbb{R}^d , we want to maximize a linear function over $\bigcap_{i=1}^n \bar{C}_i$. Since the maximum (if it exists) is achieved at the boundary of the common intersection, the problem can be reformulated as minimizing a linear function over the boundary of $\bigcup_{i=1}^n C_i$, where C_i is the (closed) halfspace complementary to \bar{C}_i ; see Figure 1. The worst-case running time of the simplex algorithm, as well as many other naïve solutions to linear programming, is proportional to the total number of vertices of $\mathcal{U}(\bar{\mathcal{C}})$. According to McMullen’s Upper Bound Theorem [96, 97], this number cannot exceed

$$\binom{n - \lceil d/2 \rceil}{\lfloor d/2 \rfloor} + \binom{n - 1 - \lceil (d-1)/2 \rceil}{\lfloor (d-1)/2 \rfloor},$$

with equality for cyclic polytopes and for all other simplicial neighborly polytopes. Regarding the dimension d as a constant, an assumption that we will follow throughout this paper, we can write this bound as $\Theta(n^{\lfloor d/2 \rfloor})$.

Robotics. Assume that we have a robot system B with d degrees of freedom, i.e., we can represent each placement of B as a point in \mathbb{R}^d . We call the space of all placements the *configuration space* of B . Suppose the (say, three-dimensional) workspace of B is cluttered with a family $\mathcal{O} = \{O_1, \dots, O_m\}$ of polyhedral obstacles whose shapes and locations are known. B is allowed to move freely in a motion that traces a continuous path in the configuration space, but B has to avoid collision with the obstacles. For each O_i , let $C_i \subseteq \mathbb{R}^d$ be the set of placements of B at which it collides with the obstacle O_i . C_i is referred to as the *C-obstacle* (or *expanded obstacle*) induced by O_i . Set $\mathcal{C} = \{C_1, \dots, C_m\}$. The *free configuration space* $\mathbb{F} = \mathbb{R}^d \setminus \mathcal{U}(\mathcal{C})$ is the set of all *free* placements of B , i.e., placements at which B does not intersect any obstacle.

For instance, let B be a convex polygonal object with r vertices, which is only allowed to translate in \mathbb{R}^2 . Let $\mathcal{O} = \{O_1, \dots, O_m\}$ be a set of m convex polygonal obstacles in \mathbb{R}^2 . Fix a reference point o (the origin) within B . A placement of B can be represented by specifying the x - and y -coordinates of o . B intersects an obstacle O_i if and only if o belongs to the “expanded obstacle” $C_i = O_i \oplus (-B)$, where \oplus denotes the *Minkowski sum*, i.e.,

$$C_i = \{x - b \mid x \in O_i, b \in B\}.$$

Hence, $\mathbb{F} = \mathbb{R}^2 \setminus \mathcal{U}(\mathcal{C})$; see Figure 2.




Figure 2. The space of free placements of the robot B is the complement of the union of the expanded obstacles C_i .

Going back to the general case, let $Z \in \mathbb{R}^d$ be a given initial free placement of B . Then the set of all free placements of B that can be reached from Z via a collision-free continuous motion corresponds to the connected component of \mathbb{F} containing Z . The problem of determining whether there exists a collision-free path from an initial configuration I to a final configuration F is equivalent to determining whether I and F lie in the same connected component of \mathbb{F} .

This close relationship between union of regions and motion planning has been a major motivation for studying the former problem, and has led to considerable work on various aspects of the union problem [11, 66, 89, 107, 109]. The complexity of $\mathcal{U}(\mathcal{C})$ serves as a trivial lower bound for the running time of many motion-planning algorithms that compute the entire free space. However, in view of the preceding discussion, there is also considerable interest in bounding the combinatorial complexity of, and constructing, a single connected component of the complement of $\mathcal{U}(\mathcal{C})$ [65, 109].

Molecular modeling. A molecule can be modeled as the union of a family of balls, where the radius of each ball depends on the atom that it models and the position of each ball depends on the molecular structure. In the *Van der Waals model*, a molecule is a family of possibly overlapping balls, where the radius of each ball is determined by the van der Waals radius of the corresponding atom in the molecule; see Figure 3 (a). Lee and Richards [87] proposed another model, called *solvent accessible model*, which is used to study the interaction between the protein and solvent molecules. A molecule is modeled as a family of balls in this model as well, but the balls representing the solvent molecules are shrunk to points and the balls representing atoms in the protein are inflated by the same amount [105]. See Figure 3 (b). Even though these models ignore various additional properties of molecules, they have been useful in a variety of applications. Many problems in molecular modeling can be formulated as problems related to geometric, combinatorial, or topological properties of the union of balls. See [46, 67, 95] for more details.




Figure 3. Representing (chain A of) the protein 1A22 as the union of a set of balls: (a) atoms are drawn using van der Waals radii, and (b) solvent accessible model.

Constructive solid geometry. Constructive solid geometry (CSG), a widely used technique in CAD and computer graphics, is a method for representing a complex object as a Boolean function of simple objects (called primitives); see Figure 4. Often CSG provides a rather simple representation of a visually complex object, using a Boolean formula cleverly. A challenging problem in this area is to compute the boundary representation of the complex object from the given Boolean function, which basically reduces to the problem of computing the union or intersection of two (or more) objects. Much work has been done in CSG on developing simple, robust, efficient algorithms for computing the boundary representation. See [61, 86] for more details.




Figure 4. Representing a complex object as a Boolean function of primitives. The figure is taken from [1].

Proximity problems. Let P and Q be two finite point sets in \mathbb{R}^d . The *directed Hausdorff distance* from P to Q , denoted by $h(P, Q)$, is (here $\|\cdot\|$ denotes the Euclidean norm, but other metrics can also be considered)

$$h(P, Q) = \max_{p \in P} \min_{q \in Q} \|p - q\|.$$

The *Hausdorff distance* between P and Q is $H(P, Q) = \max\{h(P, Q), h(Q, P)\}$. It is a widely used metric to measure similarity between two point sets. Let $B(x, r)$ denote the ball of radius r centered at x . Then $h(P, Q) \leq r$ if and only if P is contained in the union $\bigcup_{q \in Q} B(q, r)$. Hence, the decision problem of computing the Hausdorff distance, i.e., whether $H(P, Q) \leq r$, can be formulated as point location in the union of a set of congruent balls (or, more generally, of translates of the unit ball of the given norm) [10, 72, 73].

Small-size ε -nets. Given a point set P , and an admissible collection \mathcal{R} of *ranges* (subsets of P), and a parameter $\varepsilon > 0$, an ε -*net* of (P, \mathcal{R}) is a subset $N \subseteq P$ with the property that any range that contains at least $\varepsilon|P|$ points of P contains at least one point of N . By now, ε -nets are a standard and useful tool in the design and analysis of geometric algorithms; see [92, 99] for more general definitions and reviews. If the *VC-dimension* of the range space has a finite value δ (in geometry, this is the case when the ranges have simple shape, such as halfspaces, balls, tetrahedra, etc.), there exist ε -nets of size $(c\delta/\varepsilon) \log(\delta/\varepsilon)$, for some absolute constant c [71, 83]. A challenging question is to identify the situations in which the logarithmic factor can be removed or replaced by a smaller factor. See, e.g., Matoušek et al. [94] for a result of this kind, for the case when the ranges are halfplanes in the plane or halfspaces in three dimensions. Clarkson and Varadarajan [34] have recently shown that if the complexity of the union of any r ranges in \mathcal{R} is sufficiently close to $O(r)$, then the above general bound on the size of the smallest ε -net for (P, \mathcal{R}) can be improved.

Conflict-free colorings. A coloring of a family \mathcal{C} of regions in the plane is called *conflict-free* if for each point $p \in \mathcal{U}(\mathcal{C})$, there is at least one region containing p whose color is unique among all regions in \mathcal{C} that contain p . This definition was motivated by a frequency allocation problem for cellular telephone networks [54]. Minimizing the number of frequencies used by the system requires minimizing the number of colors in a conflict-free coloring of the transmission ranges of the base-stations. Alon and Smorodinsky [12] showed that whenever the family \mathcal{C} has the property that the complexity of the union of any r ranges in \mathcal{C} is $O(r)$, there is a conflict-free coloring using only $O(\log^3 D)$ colors, where D denotes the maximum number of regions in \mathcal{C} that intersect a given region. For other results on conflict-free coloring that exploit the complexity of the union of the regions to be colored, see Har-Peled and Smorodinsky [70].

These examples illustrate the wide scope of problems that can be formulated in terms of, or are closely related to, the union of a collection of geometric objects. Before proceeding further, we formalize our notation and introduce additional terminology.

Preliminaries and notation. We assume that each object C_i in the given collection \mathcal{C} is a semi-algebraic set.¹ In many cases we will also assume that each C_i has constant description complexity, which is the case, e.g., for balls, cylinders, or tetrahedra. However, we will also consider objects of non-constant description complexity, such as convex polyhedra. Also, in many planar instances, we will relax the condition, by considering fairly arbitrary curves, with the main restriction that each pair of them intersect in a constant number of points.

Each *face* of $\mathcal{U}(\mathcal{C})$ (or, more precisely, of $\partial\mathcal{U}(\mathcal{C})$) is a maximal connected (relatively open) subset of $\partial\mathcal{U}(\mathcal{C})$ that lies in the intersection of the boundaries of a fixed subset of objects, and avoids all other objects of \mathcal{C} . As usual, we refer to faces of dimension 0 and 1 as *vertices* and *edges* (or *elementary arcs*), respectively. The *combinatorial complexity* of $\mathcal{U}(\mathcal{C})$, denoted by $\kappa(\mathcal{C})$, is the total number of faces, of all dimensions, that appear on $\partial\mathcal{U}(\mathcal{C})$. Note that, in certain cases, this notion of a face is too “liberal”: if the boundary of an object $C \in \mathcal{C}$ is not a single algebraic surface, we typically regard each maximal connected portion of it that lies on a single surface (variety) as a separate “face” (this is the case, e.g., for convex polygons or polyhedra). In this case one may want to define a face of $\mathcal{U}(\mathcal{C})$ to be a maximal connected region that lies in the intersection of a fixed subset of faces of individual objects in \mathcal{C} (and avoids all other such faces and objects). In such cases, we will continue to use the notation $\kappa(\mathcal{C})$ to denote the combinatorial complexity of $\mathcal{U}(\mathcal{C})$ under this refined definition of a face.

The study of the union of geometric objects falls into the broad topic of *arrangements* of geometric objects, which has been studied since the seminal paper by J. Steiner in 1826 [112], and which has received much attention in the last quarter century. Slightly modifying the traditional definition, the *arrangement* of a finite collection \mathcal{C} of (full-dimensional) geometric objects in \mathbb{R}^d , denoted as $\mathcal{A}(\mathcal{C})$, is the decomposition of d -space into relatively open connected cells of dimensions $0, \dots, d$ induced by \mathcal{C} , where each cell is a maximal connected set of points lying in the intersection of the interiors of a fixed subset of \mathcal{C} and of the boundaries of another fixed subset, and avoids all other sets of \mathcal{C} ; lower-dimensional faces are also referred to as *faces*. (As above, if the boundaries of the objects of \mathcal{C} do not have constant description complexity, the arrangement itself is refined accordingly.) Note that $\mathcal{U}(\mathcal{C})$ is a substructure of $\mathcal{A}(\mathcal{C})$, in the sense that each face of $\mathcal{U}(\mathcal{C})$ is also a face of $\mathcal{A}(\mathcal{C})$. $\mathcal{U}(\mathcal{C})$ typically contains in its interior many faces of $\mathcal{A}(\mathcal{C})$, but they are ignored in the analysis of its complexity. As such, $\kappa(\mathcal{C})$ is bounded by the combinatorial complexity of $\mathcal{A}(\mathcal{C})$, which, in the worst case, is $\Theta(n^d)$ if the objects in \mathcal{C} are semi-algebraic sets of constant description complexity. In the worst case, the asymptotic bound on $\kappa(\mathcal{C})$ can indeed be $\Theta(n^d)$. This is the case, for example, when \mathcal{C} is a family of n large and flat “plates” in \mathbb{R}^d , each being the region enclosed between a pair of parallel and sufficiently close hyperplanes. See Figure 5 for a simple planar construction involving triangles. However, if \mathcal{C} satisfies certain natural conditions, $\kappa(\mathcal{C})$ may be smaller. For example, the case of halfspaces, mentioned above, yields the particularly favorable bound $\Theta(n^{\lfloor d/2 \rfloor})$ on $\kappa(\mathcal{C})$. The challenge is thus to identify classes of objects for which the bound on $\kappa(\mathcal{C})$ is substantially smaller than $\Theta(n^d)$. As we shall see, in most of the cases that we will review here, $\kappa(\mathcal{C})$ is close to $O(n^{d-1})$. Easily constructed matching lower bounds indicate that this

¹A subset of \mathbb{R}^d is called a *real semi-algebraic set* if it is described in terms of a bounded number of polynomials of bounded maximum degree in a bounded number of variables.

is the best “order of magnitude” one can hope for in most of these favorable instances. We will occasionally use the shorthand notation $O^*(f(n))$ to denote bounds of the form $C_\varepsilon f(n) \cdot n^\varepsilon$, which hold for any $\varepsilon > 0$, where the constant of proportionality C_ε depends on ε , and typically tends to ∞ as ε decreases to 0.




Figure 5. n pairwise crossing triangles with $\Theta(n^2)$ intersection points on the boundary of their union.

The rest of the survey is organized as follows. We review the known results on the complexity of the union of planar objects in Section 2, and of three-dimensional objects in Section 3. We also sketch proofs of some of the main results. We then briefly review in Section 4 the (very few) known results in higher dimensions. Section 5 discusses the relationship between the union of objects and generalized Voronoi diagrams, and gives a brief review of the recent progress in the analysis of the complexity of these diagrams. We conclude in Section 6 with a short discussion of the topic and its relatives.

2 Union of Planar Objects

In this section we review the known results on the union of geometric objects in the plane. The study of the union of planar objects goes back to at least the early 1980s, when researchers were interested in the union of rectangles or disks, motivated by VLSI design, biochemistry, and other applications [26, 77, 84, 104]. However, the early work focused on computing the union or its measure, rather than bounding its complexity.

2.1 Union of pseudo-halfplanes

Let $\mathcal{F} = \{f_1, \dots, f_n\}$ be a set of n totally defined continuous univariate functions. For each f_i , let C_i be the set of points lying on one of the sides of (above or below) the graph of f_i . We refer to C_i as a *pseudo-halfplane*. If C_i lies below (resp., above) f_i , it is called a *lower* (resp., *upper*) pseudo-halfplane. Set $\mathcal{C} = \{C_1, \dots, C_n\}$. If each f_i is a linear function, then $\partial\mathcal{U}(\mathcal{C})$ is the boundary of a convex polygon, so $\kappa(\mathcal{C})$ is linear. For more general functions, the bounds on $\kappa(\mathcal{C})$ are more involved, and are related to lower and upper envelopes, defined as follows.

The *lower envelope* of a collection \mathcal{F} of functions, as above, denoted by $L_{\mathcal{F}}$, is the pointwise minimum of the functions in \mathcal{F} , i.e.,

$$L_{\mathcal{F}}(x) = \min_{1 \leq i \leq n} f_i(x).$$

The *upper envelope* is defined as the pointwise maximum of \mathcal{F} , i.e.,

$$U_{\mathcal{F}}(x) = \max_{1 \leq i \leq n} f_i(x).$$

If each C_i is a lower pseudo-halfplane, then $\mathcal{U}(\mathcal{C})$ is the region lying below the upper envelope of \mathcal{F} . Similarly, if each C_i is an upper pseudo-halfplane, then $\mathcal{U}(\mathcal{C})$ is the region lying above the lower envelope of \mathcal{F} . A fundamental observation (see [109]) is that if the graphs of any pair of functions in \mathcal{F} intersect in at most s points, for any fixed constant s , then the graph of the lower or upper envelope of \mathcal{F} consists of at most $\lambda_s(n)$ elementary arcs, where $\lambda_s(n)$ is the maximum length of an (n, s) *Davenport-Schinzel sequence*; see [109] for more details. Letting $\alpha(n)$ denote the extremely slow-growing inverse Ackermann function, the best known bounds on $\lambda_s(n)$ are

$$\begin{aligned} \lambda_1(n) &= n, \\ \lambda_2(n) &= 2n - 1, \\ \lambda_3(n) &= \Theta(n\alpha(n)), \\ \lambda_4(n) &= \Theta(n \cdot 2^{\alpha(n)}), \\ \lambda_{2s+2}(n) &= n \cdot 2^{\Theta(\alpha^s(n))} && \text{for } s > 1, \\ \lambda_{2s+3}(n) &= n\alpha(n)^{O(\alpha^s(n))} && \text{for } s \geq 1. \end{aligned}$$




Figure 6. Functions in \mathcal{F}^- (resp., \mathcal{F}^+) are drawn with dashed (resp., solid) lines. The sandwich region between $U_{\mathcal{F}^-}$ and $L_{\mathcal{F}^+}$, the complement of $\mathcal{U}(\mathcal{C})$, is shaded.

The case when some of the regions of \mathcal{C} are lower pseudo-halfplanes and some are upper pseudo-halfplanes is not that much harder. Let \mathcal{F}^- (resp., \mathcal{F}^+) denote the subset of those functions in \mathcal{F} that bound lower (resp., upper) pseudo-halfplanes in \mathcal{C} . Then $\mathcal{U}(\mathcal{C})$ is the complement of the *sandwich region*, consisting of those points that lie above the upper envelope $U_{\mathcal{F}^-}$ and below the lower envelope $L_{\mathcal{F}^+}$. See Figure 6. It is known (and easy to show) that the complexity of the sandwich region is proportional to the sum of the complexities of $U_{\mathcal{F}^-}$ and of $L_{\mathcal{F}^+}$. We thus have the following result.

Theorem 2.1. *Let \mathcal{C} be a set of n pseudo-halfplanes so that the boundaries of any pair of them intersect in at most s points. Then $\kappa(\mathcal{C}) = O(\lambda_s(n))$.*

2.2 Regions with few pairwise boundary intersections

Let $\mathcal{C} = \{C_1, C_2, \dots, C_n\}$ be a family of n simply connected regions in the plane, each bounded by a simple closed Jordan curve. Assume, for simplicity, that these curves are in *general position*, i.e., any two of them cross only a finite number of times (two curves γ_1 and γ_2 are said to *cross* each other at a point, if γ_1 passes from one side of γ_2 to the other side at this point), no two curves touch each other, and no three curves pass through a common point.²

In this subsection we consider the case in which the boundaries of any pair of regions in \mathcal{C} cross in a small number of points, and derive linear, or near-linear bounds for the complexity of their union.

Union of pseudo-disks. If the boundaries of any two distinct regions in \mathcal{C} cross at most twice, then \mathcal{C} is called a family of *pseudo-disks*. See Figure 7. In this especially favorable case, we have the following result.




Figure 7. (a) A family of pseudo-disks. (b) Another family of n pseudo-disks with $6n - 12$ elementary arcs on the boundary of its union.

Theorem 2.2 (Kedem et al. [76]). *Let $\mathcal{C} = \{C_1, C_2, \dots, C_n\}$ be a family of $n \geq 3$ pseudo-disks in the plane. Then the boundary of $\mathcal{U}(\mathcal{C})$ consists of at most $6n - 12$ elementary arcs, and this bound is tight in the worst case.*

We present the proof of Theorem 2.2 for the case of circular disks. (A more direct proof for the union of circular disks based on the so-called *lifting transform*, which extends to higher dimensions, is given in Section 4.) Assign to each C_i its center, p_i , and connect p_i to p_j by a straight-line segment if and only if ∂C_i and ∂C_j cross each other, and at least one of their crossing points belongs to $\partial \mathcal{U}(\mathcal{C})$; see Figure 8. It is easy to verify that no two segments in the resulting geometric graph G cross each other, i.e., G is *planar*. Therefore, G has at most $3n - 6$ edges, each of which corresponds

²One can extend the general position assumption to other instances and to higher dimensions; see [109]. There is a general perturbation-based argument [109] that shows that the asymptotic upper bound on $\kappa(\mathcal{C})$ is not affected by the general position assumption.

to at most two vertices of $\partial\mathcal{U}(\mathcal{C})$. Consequently, the number of crossings on $\partial\mathcal{U}(\mathcal{C})$, and hence the number of elementary arcs, is at most $6n - 12$. The proof for the case of general pseudo-disks also uses planarity, and follows as a special case of the proof of a more general result (Theorem 2.5), given later in this section.

In other words, the *complexity* of $\mathcal{U}(C)$ is at most linear in n . A lower-bound construction (which can also be realized using normal disks), in which the number of elementary arcs is exactly $6n - 12$, is shown in Figure 7(b). The proof for general pseudo-disks can be found later in this subsection, where a generalization of Theorem 2.2 is stated and proved (Theorem 2.5).




Figure 8. The proof of Theorem 2.2 for disks.

We conclude the discussion on pseudo-disks by giving two examples of pseudo-disks that arise in practice. First, recall the example of translational motion planning in the plane. Because of this application, the lemma is formulated in terms of the reflection $-B$ of B , but it holds of course for B too.

Lemma 2.3 (Kedem et al. [76]). *Let O_1, O_2 , be two disjoint convex bodies in the plane, and let B be another convex body in the plane. Then the boundaries of the Minkowski sums $C_1 = O_1 \oplus (-B)$ and $C_2 = O_2 \oplus (-B)$ cross at most twice.*

Proof: We argue that C_1 and C_2 have exactly two common outer tangents, from which the lemma follows easily. For a convex body C and for each $\theta \in [0, 2\pi)$, define $f(C, \theta)$ to be the signed distance from the origin o to the unique tangent $\tau(C, \theta)$ to C at orientation θ , which has C lying to its left; $f(C, \theta)$ is positive (resp., negative) if o lies to the left (resp., right) of $\tau(C, \theta)$. It easily follows from the definition of Minkowski sums that

$$\begin{aligned} f(C_1, \theta) &= f(O_1, \theta) - f(B, \theta) \\ f(C_2, \theta) &= f(O_2, \theta) - f(B, \theta). \end{aligned}$$




Figure 9. The proof that C_1 and C_2 have only two common outer tangents.

See Figure 9. Hence, $\tau(C_1, \theta) = \tau(C_2, \theta)$, i.e., C_1 and C_2 have a common outer tangent at orientation θ , if and only if $\tau(O_1, \theta) = \tau(O_2, \theta)$, i.e., O_1 and O_2 have a common outer tangent at orientation θ . Since O_1 and O_2 are disjoint, they have exactly two common outer tangents, and the claim follows. \square

Lemma 2.3 implies that if $\mathcal{O} = \{O_1, \dots, O_n\}$ is a set of $n \geq 3$ pairwise-disjoint convex obstacles and B is a convex “robot” translating in the plane, then $\partial\mathbb{F}$, the boundary of the free space, has at most $6n - 12$ elementary arcs. If B and the obstacles are convex polygons, so that B has k vertices, and the total number of obstacle vertices is s , then \mathbb{F} has $O(kn + s)$ vertices, of which at most $6n - 12$ are reflex intersection vertices.

Another commonly occurring example of pseudo-disks is the case of *homothets*. Let B be a convex body in the plane, and for $1 \leq i \leq n$, let C_i be a homothetic copy of B , i.e., $C_i = \lambda_i B + x_i$ for arbitrary parameters $\lambda_i > 0$ and $x_i \in \mathbb{R}^2$. Set $\mathcal{C} = \{C_1, \dots, C_n\}$. It is known (and easy to show) that \mathcal{C} is a family of pseudo-disks. Hence, $\partial\mathcal{U}(\mathcal{C})$ has at most $6n - 12$ elementary arcs.

Allowing three intersections. What happens if we somewhat weaken the condition in Theorem 2.2, by assuming that the boundaries of any two members of \mathcal{C} cross at most *three* times, rather than twice? At first glance this problem seems to be foolish because two closed curves in general position can cross only an *even* number of times. However, by a slight modification we obtain a meaningful question with a somewhat surprising answer.

Theorem 2.4 (Edelsbrunner et al. [44]). *Let $\{\gamma_1, \gamma_2, \dots, \gamma_n\}$ be a family of n simple curves in general position in the upper halfplane. Assume that the endpoints of each curve are on the x -axis, and that any two curves cross at most three times. Let C_i denote the bounded region enclosed by γ_i and the x -axis (see Figure 10 (a)). Then $\kappa(\mathcal{C}) = O(n\alpha(n))$, and this bound is asymptotically tight.*

Note that if each γ_i is an x -monotone curve, then Theorem 2.4 follows from Theorem 2.1. However, as seen in Figure 10 (a), nonmonotone curves may cause holes in the union (i.e., bounded components of the complement of the union), which makes the proof of the above theorem less obvious and quite technical. The proof of Edelsbrunner et al. [44] proceeds by constructing a




Figure 10. (a) Union of 3-intersecting regions. (b) The curve Γ ; it switches from one input curve to another at hollow circles, and the filled circles denote the vertices of the $\mathcal{U}(\mathcal{C})$ that are not switching points of Γ ; $\Sigma = \langle 1, 1, 2, 2, 3, 3, 4, 4, 4, 5, 4, 6, 6, 7, 7 \rangle$.

curve Γ that starts at $-\infty$ on the x -axis and proceeds to the right, always following one of the γ_i , consistently with its orientation, possibly switching arcs at intersection points, but never visiting a point more than once, and eventually ending at $+\infty$ on the x -axis. The curve Γ traces each arc of $\mathcal{U}(\mathcal{C})$ exactly once, consistently with the orientation of the corresponding input curve, and all holes of $\mathcal{U}(\mathcal{C})$ lie outside Γ , i.e., Γ can be continuously deformed within $\mathcal{U}(\mathcal{C})$, so as to coincide with the x -axis; see Figure 10 (b). The proof then continues by labeling each elementary arc of Γ that appears on $\partial\mathcal{U}(\mathcal{C})$ with the curve to which it belongs, producing a sequence Σ of labels. One can then show that if one removes every symbol of Σ which is equal to its predecessor, then the remaining sequence is an $(n, 3)$ Davenport-Schinzel sequence, and thus its length is $O(n\alpha(n))$. One can also show that the number of deleted labels is $O(n\alpha(n))$, which completes the proof of Theorem 2.4. The details can be found in [44].

Beyond three intersections. If we allow the boundaries of two objects in \mathcal{C} to cross at most *four* times, then the situation completely deteriorates. As illustrated in Figure 5, every pair of the n triangles intersect in precisely four points, and all $4\binom{n}{2}$ intersection points belong to the boundary of their union. However, Whitesides and Zhao [117] discovered that by excluding certain types of crossings between the members of \mathcal{C} , it is still possible to prove a linear upper bound on the complexity of $\mathcal{U}(\mathcal{C})$ even if pairs of members of \mathcal{C} may intersect in more than two points. More precisely, a family \mathcal{C} of simply connected regions bounded by simple closed curves in general position in the plane is called *k-admissible* (with k even) if for any pair $C_i, C_j \in \mathcal{C}$,

- (i) $C_i \setminus C_j$ and $C_j \setminus C_i$ are connected, and
- (ii) ∂C_i and ∂C_j cross in at most k points.

See Figure 11. Theorem 2.2 is a special case of the following theorem (with $k = 2$).

Theorem 2.5 (Whitesides and Zhao [117]). *Let $\mathcal{C} = \{C_1, C_2, \dots, C_n\}$ be a k -admissible family of $n \geq 3$ simply connected regions in general position in the plane. Then $\partial\mathcal{U}(\mathcal{C})$ consists of at most $k(3n - 6)$ elementary arcs, and this bound cannot be improved.*




Figure 11. A pair of regions belonging to a (a) 4-admissible family, (b) nonadmissible family ($C_1 \setminus C_2$ is disconnected).

Proof: We sketch the proof given in [101] (see also [103]). As usual, it suffices to bound the number of vertices of $\mathcal{U}(\mathcal{C})$. For every C_i that contributes at least one arc to $\partial\mathcal{U}(\mathcal{C})$, we fix a point p_i in the interior of such an arc. For any pair $C_i, C_j \in \mathcal{C}$ that generate a vertex q on $\partial\mathcal{U}(\mathcal{C})$, we draw an edge (but only one!) e_{ij} between p_i and p_j , as follows. Starting from p_i , follow ∂C_i to q (in any direction), and from there follow ∂C_j to p_j (in any direction). Let H be the resulting graph; see Figure 12.

We claim that any two edges of H not incident to the same vertex cross an even number of times. We sketch the proof of this claim for the case of pseudo-disks ($k = 2$). Let e_{ij} and e_{kl} be two edges of H , where the first (resp., second) edge passes through an intersection point q_{ij} (resp., q_{kl}) of the boundaries of C_i, C_j (resp., C_k, C_l), which lies on the boundary of the union. Each of the points q_{ij}, q_{kl} splits its respective edge into two “half-edges.” We claim that any pair of half-edges cross an even number of times, that is, either twice or not at all. If this were not the case, then the two half-edges would cross exactly once, and then the pseudo-disk property is easily seen to imply that one endpoint of each half-edge must lie in the interior of the other object, which is impossible, since each half-edge starts and ends at a point on the boundary of the union. This argument also applies to any even $k > 2$, exploiting condition (i) above.




Figure 12. The union of pseudo-disks via a planarity argument. Each point p_i is labeled as i , and points q_{ij} are labeled as ij . Here $e(1, 2)$ and $e(3, 4)$ cross each other six times.

A remarkable result by Chojnacki (alias Hanani) [33] (see also [116, 91], and [103] for a new proof), states that if a graph G can be drawn in the plane so that any two of its edges not incident to the same vertex cross an even number of times, then G is planar. Hence, we can conclude that H is

planar, so it has at most $3n - 6$ edges. That is, there are at most $3n - 6$ pairs $\{C_i, C_j\}$ contributing vertices to $\mathcal{U}(\mathcal{C})$, and each of them can contribute at most k such points. \square

Counting regular vertices. If ∂C_i and ∂C_j intersect in precisely two points, then we call these intersection points *regular*; otherwise their intersection points are called *irregular*. See Figure 13(a). A vertex of $\mathcal{U}(\mathcal{C})$ is regular if it is a regular intersection point, and irregular otherwise. If \mathcal{C} is a set of pseudo-disks, then $\mathcal{U}(\mathcal{C})$ does not have any irregular vertex. A natural way to generalize Theorem 2.2 is to obtain sharp bounds on the number of regular vertices in $\mathcal{U}(\mathcal{C})$ even if the boundaries of some pairs of objects in \mathcal{C} intersect at more than two points.




Figure 13. (a) Regular (darkly shaded circles) and irregular (lightly shaded circles) vertices of planar unions. (b) A union of convex polygons with quadratically many regular vertices.

Let \mathcal{C} be a family of $n \geq 3$ convex regions in general position in the plane, and let $R(\mathcal{C})$ and $I(\mathcal{C})$ denote, respectively, the number of regular and irregular vertices of $\mathcal{U}(\mathcal{C})$. Pach and Sharir [101] showed that

$$R(\mathcal{C}) \leq 2I(\mathcal{C}) + 6n - 12. \quad (1)$$

This result is sharper than Theorem 2.2, in the sense that, for establishing the upper bound $6n - 12$ on the number of elementary arcs (or the number of intersection points) on $\partial\mathcal{U}(\mathcal{C})$, one does not have to insist that all boundary intersection points of pairs of objects of \mathcal{C} be regular. It suffices to require that all vertices of $\mathcal{U}(\mathcal{C})$ be regular. The extension of the above result to nonconvex regions remains elusive:

Open Problem 1. *Is it true that for every set \mathcal{C} of n simply connected regions in general position in the plane, one has $R(\mathcal{C}) \leq 2I(\mathcal{C}) + 6n - 12$?*

It is not hard to show that the coefficient of $I(\mathcal{C})$ in (1) cannot be replaced by any constant smaller than 2. Moreover, in general $R(\mathcal{C})$ can be $\Theta(|\mathcal{C}|^2) = \Theta(n^2)$ in the worst case, as is illustrated in Figure 13(b), unless we limit the number of times the boundaries of a pair of curves in \mathcal{C}




Figure 14. The lower-bound construction for the number of regular vertices on the union of rectangles and disks.

are allowed to cross each other (this number is not bounded by a constant in Figure 13(b)). However, we cannot expect a *linear* upper bound even under such an assumption (unless we deal with pseudo-disks): For any n , we can construct a family \mathcal{C} of n disks and rectangles in general position in the plane satisfying $R(\mathcal{C}) = \Omega(n^{4/3})$, as follows. Take a system of $n/2$ lines and $n/2$ points with $\Theta(n^{4/3})$ incidences between them [99]. Fix two sufficiently small parameters $0 < \varepsilon < \varepsilon' < 2\varepsilon$. Shift each line by distances ε and 2ε , and create a sufficiently long rectangle bounded by the shifted copies. Expand each point into a disk of radius ε' . See Figure 14. With an appropriate choice of $\varepsilon, \varepsilon'$, the resulting family of rectangles and disks has $\Theta(n^{4/3})$ regular vertices on the boundary of their union. For the special case of rectangles and disks, this bound is asymptotically tight [17]. If \mathcal{C} is a set of n simply connected regions so that the boundaries of any pair of them intersect in at most s points, for some constant $s > 0$, then there exists $\delta = \delta(s) > 0$ such that $\mathcal{U}(\mathcal{C})$ has $O(n^{2-\delta})$ regular vertices [17]. Recently, the bound has been improved to $O^*(n^{4/3})$, where the constant of proportionality depends on s (and on the hidden $\varepsilon > 0$), if the objects in \mathcal{C} are *convex* [58]. See also [56] for some related results.

Open Problem 2. Let \mathcal{C} be a set of simply connected regions in general position in the plane, so that the boundaries of any pair of them intersect in at most some constant number, s , of points. Obtain a sharp bound on $R(\mathcal{C})$, which depends only on n (and s), and not on $I(\mathcal{C})$.

2.3 Union of fat objects

The construction depicted in Figure 5, showing that the union of n triangles may have *quadratic* complexity, uses extremely narrow triangles. On the other hand, as we saw in Section 2.2, the complexity of the union of n circular disks or (axis-parallel) squares is linear, thereby raising the question whether the union of “fat” objects has small complexity. In the last fifteen years this question has been answered in the affirmative under various notions of fatness [13, 48, 50, 51, 93, 102]. In fact, these results have motivated the study of faster geometric algorithms, for a variety of applications, for fat objects in two and three dimensions [4, 7, 35, 75, 85, 110, 111]. In this section we review the known results on the complexity of the union of fat planar objects, starting with the simplest but important case of fat triangles.

Union of fat triangles. For any fixed $\alpha > 0$, a triangle is called α -fat if each of its angles is at least α . Matoušek et al. [93] proved that the complexity of the union of n α -fat triangles is $O(n \log \log n)$, for any fixed $\alpha > 0$. Their proof is based on the fact that every family \mathcal{C} of n α -fat triangles in the plane determines at most a linear number of *holes*, namely, bounded components of the complement of $\mathcal{U}(\mathcal{C})$. The strongest known bound on the number of holes (in terms of its dependence on α) is the following.

Theorem 2.6 (Pach and Tardos [102]). *Any family \mathcal{C} of n α -fat triangles in the plane determines $O((n/\alpha) \log(1/\alpha))$ holes. This bound is tight up to the logarithmic factor.*

Proof: We sketch the earlier proof, given in [93], of an upper bound on the number of holes, which is linear in n , but in which the dependence of the constant of proportionality on α is $1/\alpha^3$.




Figure 15. Replacing a fat triangle by three canonical triangles.

We first replace each triangle $\Delta \in \mathcal{C}$ by three $(\alpha/2)$ -fat triangles contained in Δ , by bending the edges of Δ inwards, as depicted in Figure 15, so that the directions of the edges of the new triangles belong to the family of the $O(1/\alpha)$ so-called “canonical” directions $j\alpha/2$, $j = 0, 1, \dots$. During the bending, the holes of the union *expand*, so their number can decrease only when two holes merge into a common hole. However, this can happen only when the bending sweeps through a triangle vertex, which can happen only once per vertex, and thus implies that the number of holes can go down by at most $3n$.

Thus, we obtain $O(1)$ *canonical families* of $(\alpha/2)$ -fat triangles with fixed edge directions, so that each family consists of *homothetic triangles*, i.e., similar triangles in parallel positions. It suffices to bound the number of holes in the union of these families. Since any vertex of the union is also a vertex of the union of just *two families*, it suffices to establish a linear upper bound on the number of holes determined by the union of two canonical families.

As stated in Section 2.2, the union of homothetic triangles has linear complexity, so the union of all members of a *single* canonical family has linear size. Consider then the union of two families R and G of “red” and “green” triangles, where the triangles in each family are $((\alpha/2)$ -fat and) homothetic to each other. For simplicity, assume $|R| = |G| = n$. We may ignore holes with a “monochromatic” vertex (i.e., a vertex of $\mathcal{U}(R)$ or of $\mathcal{U}(G)$), as there are only $O(n)$ such holes. It is not hard to see that any remaining hole is either a *convex quadrangular hole* (see Figure 16) or a *convex hexagonal hole*. An easy application of Euler’s polyhedral formula implies that the number of holes is dominated by the number of quadrangular holes, which we now proceed to bound.

We draw a graph H , whose vertices are represented by the elementary arcs of the boundary of




Figure 16. Shaded regions denote bichromatic quadrangular holes; a, b, c, d, e are the elementary arcs of $\mathcal{U}(\mathcal{R})$ that induce these holes; $(a, b), (c, e), (d, e)$ are the edges of H .

$\mathcal{U}(R)$ and whose edges connect pairs of elementary arcs that bound a common quadrangular hole; we draw such an edge by connecting (only once!) the pair of red arcs across such a common hole. Obviously, H is a planar graph. See Figure 16.

Hence, there are only $O(n)$ pairs of elementary arcs of $\mathcal{U}(R)$ that form a common quadrangular hole. There are $O(n)$ quadrangular holes that are *first* or *last* along some edge of a green triangle; let us call them *extremal*. On the other hand, it can be shown by a simple trigonometric calculation, illustrated in Figure 17, that each pair $\{e, e'\}$ of elementary arcs of $\mathcal{U}(R)$ can generate only $O(1)$ nonextremal quadrangular holes. \square




Figure 17. A fixed pair of elementary arcs of $\mathcal{U}(R)$ generates only $O(1/\alpha^2)$ (nonextremal, bichromatic, and quadrangular) holes.

This result can be used to establish a more general upper bound for the number of holes determined by a family of triangles with given angles.

Theorem 2.7 (Pach and Tardos [102]). Let $\mathcal{C} = \{C_1, C_2, \dots, C_n\}$ be a family of $n > 1$ trian-

gles in the plane, and let α_i denote the smallest angle of C_i , for $1 \leq i \leq n$. Suppose $0 < \alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_n$, and let $k \leq n$ be the largest integer satisfying $\sum_{i=1}^k \alpha_i < \pi$. Then \mathcal{C} determines $O(nk \log k)$ holes. Furthermore, there exists a family $\mathcal{C}' = \{C'_1, C'_2, \dots, C'_n\}$, where C'_i is congruent to C_i and \mathcal{C}' determines $\Omega(nk)$ holes.

Proof: Note that each C_i , for $k > i$, is (π/k) -fat, so the union of C_{k+1}, \dots, C_n , denoted by \mathcal{U}' , has $O(nk \log k)$ holes. Adding C_1, \dots, C_k to \mathcal{U}' creates at most $O(nk)$ new holes. \square

If we consider *infinite wedges* (i.e., convex cones) rather than triangles, then the same bound holds not only for the number of *holes*, but also for the *complexity* of the union. The following result strengthens some earlier bounds in [13, 50].

Theorem 2.8 (Pach and Tardos [102]). *Let \mathcal{C} be a family of n wedges in the plane with angles $0 < \alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_n < \pi$. Let $k \leq n$ be the largest integer satisfying $\sum_{i=1}^k \alpha_i < \pi$. If $k \geq 2$, then $\kappa(\mathcal{C})$ is $O(nk \log k)$. Furthermore, there exists a family of n wedges with angles $\alpha_1, \alpha_2, \dots, \alpha_n$, which determines $\Omega((\pi - \alpha_n)nk)$ holes.*

Using Theorem 2.6, we state a slightly strengthened version of the result in [93].

Theorem 2.9. *For any fixed $\alpha > 0$, the boundary of the union of n α -fat triangles in the plane consists of at most $O((n/\alpha) \log \log n \log(1/\alpha))$ elementary arcs.*

Matoušek et al. [93] also proved that if, in addition to being fat, all triangles have roughly the same size (i.e., the ratio between any pair of diameters is bounded by a constant), then their union has linear complexity. On the other hand, by slightly modifying the $\Omega(n\alpha(n))$ lower-bound construction for the lower envelopes of n segments [118], one can construct n equilateral ($\pi/3$ -fat) triangles, whose union has a slightly superlinear (i.e., $\Omega(n\alpha(n))$) complexity.

We conclude the discussion on fat triangles by mentioning an obvious open problem.

Open Problem 3. *What is the maximum complexity of the union of n α -fat triangles?*

Union of fat convex objects. Extending the notion of fatness to more general objects, we call a *convex body* C in the plane α -fat, for $\alpha \geq 1$, if there exist two disks D, D' , such that $D \subseteq C \subseteq D'$, and the ratio between the radii of D' and D is at most α . See Figure 18. Note that this extends the definition of fatness for triangles: an α -fat triangle is easily seen to be α' -fat as a convex body, for a suitable $\alpha' \geq 1$, and vice versa. Efrat and Sharir [51] showed that the complexity of the union of n simply shaped convex α -fat objects in the plane is $O^*(n)$, where the constant of proportionality depends on (the hidden ε , and on) the maximum number of intersections between any pair of boundaries. The proof uses both the bound on the complexity of the union of fat triangles, and the bound on the number of regular vertices of the union (see (1)).

We also remark that the complexity of the union of n *arbitrary* convex *polygons* with a total of s vertices is $O(n^2 + s\alpha(n))$ [19].




Figure 18. Fat-like planar objects with near-linear union complexity.

Union of fat non-convex objects. There are other, more general, notions of “fatness” that extend to non-convex objects, and for which the combinatorial complexity of the union of n “fat” planar objects remains $O^*(n)$. For instance, call a possibly non-convex object C α -round if for each point $p \in \partial C$, there exists a disk D of radius $\alpha \text{diam}(C)$ which passes through p and which is contained in C ; see Figure 18. Informally, α -round objects cannot have convex corners, nor can they have very thin bottlenecks (but they can have reflex corners). Efrat and Katz [49] have shown that the complexity of the union of n α -round objects is $O(\lambda_s(n) \log n)$, where s is a constant that depends on the description complexity of the input objects. This result has been further extended by Efrat [48] to so-called (α, β) -covered objects: An object C is (α, β) -covered if for each point $p \in \partial C$, there exists an α -fat triangle T that has p as a vertex, is contained in C , and each of its edges is at least $\beta \text{diam}(C)$ long; see Figure 18. Thus, these objects are not necessarily smooth, but their corners cannot be too sharp. Efrat [48] has shown that if \mathcal{C} is a collection of n (α, β) -covered objects, each pair of whose boundaries intersect in at most $s = O(1)$ points, then $\kappa(\mathcal{C}) = O(\lambda_{s+2}(n) \log^2 n \log \log n)$. See also [85, 110, 111] for other related results.

3 Union of Objects in Three Dimensions

3.1 Overview

Starting in the mid 1990s, research on the complexity of the union of geometric objects has shifted to the study of instances in three and higher dimensions. As mentioned in the introduction, the maximum complexity of the union of n simply shaped objects in \mathbb{R}^3 is $\Theta(n^3)$, and this bound can already be attained by flat boxes. There are very few particularly favorable cases for which the union complexity is linear in n , including the cases of halfspaces and of axis-parallel unit cubes [31, 28]. In general, though, the goal is to find classes of objects for which the maximum complexity of the union is nearly quadratic. Indeed, in most of the cases studied so far (as will be reviewed below), the complexity of the union can be quadratic (and sometimes slightly super-quadratic) in the worst case. This is the case, e.g., for balls, cubes, congruent cylinders, and halfspaces bounded by xy -monotone surfaces of constant description complexity.

As the evidence discovered so far suggests, there are several important classes of objects in \mathbb{R}^3

whose union has at most nearly-quadratic complexity, in complete analogy to the planar situation. One such class is the class of *fat* objects, where, in complete analogy with the planar case, a compact convex object C is called α -*fat* if the ratio between the radii of the smallest enclosing ball and of the largest inscribed ball of C is at most α . Other notions of fatness, such as α -*roundness*, have also been extended to \mathbb{R}^3 . A prevailing conjecture is that the maximum complexity of the union of such fat objects is indeed at most nearly 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 geometry [37, Problem 4]. Nevertheless, considerable progress towards establishing this bound has recently been made, as we will shortly review.

As in the plane, another candidate class of objects with small union complexity are Minkowski sums of pairwise disjoint convex objects with a fixed convex object. In the plane, this class was handled by showing that its members are pseudo-disks, and then by applying the general linear bound of [76] (Theorem 2.2). However, the analysis of the union of such Minkowski sums is considerably harder in 3-space (because they are not “pseudo-balls”—see below), and there are only a few (albeit important) instances for which a near-quadratic bound has been established [11, 20]; see Section 3.4.

A third class of objects with small union complexity are *pseudo-halfspaces*, i.e., regions lying above or below an xy -monotone surface (the graph of a continuous totally defined function). This extends the class of pseudo-halfplanes, and was one of the first classes to be studied.

We note that extending the notion of pseudo-disks to three dimensions does not seem to lead to any new insights. A family of regions in \mathbb{R}^3 is said to consist of *pseudo-balls*, if the boundaries of any two members intersect in a single closed curve, and the boundaries of any three members intersect in at most two points. It is trivial to show that the complexity of the union of a collection \mathcal{C} of n pseudo-balls is $O(n^2)$, by considering the portion of the union boundary on the boundary of each member of \mathcal{C} separately, and by applying Theorem 2.2. Hence, in particular, the complexity of the union of n balls in \mathbb{R}^3 is $O(n^2)$; it is easy to construct examples where the union has $\Theta(n^2)$ vertices, even with unit balls, and even when the unit balls all have a common point; see [29]. Somewhat surprisingly, Minkowski sums of disjoint convex bodies with a fixed convex body are not *pseudo-balls*; see a more detailed discussion below.

3.2 Union of pseudo-halfspaces

Let $\mathcal{F} = \{f_1, \dots, f_n\}$ be a family of n continuous totally defined bivariate functions (in x, y). As in Section 2.1, we refer to the region lying below (resp., above) the graph of f_i as the lower (resp., upper) *pseudo-halfspace* bounded by that graph. For each $1 \leq i \leq n$, let C_i be one of these two pseudo-halfspaces, and let \mathcal{C} denote the collection $\{C_1, \dots, C_n\}$.

Halperin and Sharir [68] proved that if each function in \mathcal{F} is of constant description complexity, then the complexity of the lower or upper envelope of \mathcal{F} is $O^*(n^2)$. This immediately implies that if all the C_i 's are lower (or all are upper) pseudo-halfspaces, then their union has $O^*(n^2)$ complexity. Agarwal et al. [9] have established an $O^*(n^2)$ bound on the complexity of the sandwich region

between the lower and upper envelopes of two respective families of a total of n bivariate functions, each of constant description complexity. The proof is based on the following interesting result: Let \mathcal{F} and \mathcal{G} be two collections of a total of n bivariate functions, as above, and let $M_{\mathcal{F}}$ (resp., $M_{\mathcal{G}}$) denote the *minimization diagram* of \mathcal{F} (resp., \mathcal{G}), namely, the xy -projection of the lower envelope of \mathcal{F} (resp., \mathcal{G}). Then the *overlay* of the two minimization diagrams has $O^*(n^2)$ complexity. Note that we make no assumption on any relation between \mathcal{F} and \mathcal{G} . Also, the result continues to hold when one or both diagrams are replaced by the respective *maximization diagram*, i.e., the xy -projection of the respective upper envelope. This implies the following result.

Theorem 3.1 (Agarwal et al. [9]). *Let \mathcal{C} be a set of n pseudo-halfspaces in \mathbb{R}^3 , each of which is a semi-algebraic set of constant description complexity. Then the complexity of $\mathcal{U}(\mathcal{C})$ is $O^*(n^2)$.*

3.3 Union of convex polyhedra

As already remarked, an easy extension of the planar construction shown in Figure 5 shows that the maximum complexity of the union of n (axis-aligned or arbitrarily aligned) boxes (or wedges, or tetrahedra) in \mathbb{R}^3 is $\Theta(n^3)$; see Figure 19. Moreover, we can easily adapt this construction to show that the maximum union complexity of three *nonconvex* polyhedra with a total of s facets is $\Theta(s^3)$. A natural question is whether a similar lower bound also exists for the complexity of the union of convex polyhedra, i.e., a bound that is cubic in the number of facets. The following result by Aronov et al. [21] answers this question in the negative, and calibrates, more or less, the true maximum complexity of such a union.




Figure 19. Union of n boxes in \mathbb{R}^3 with $\Theta(n^3)$ complexity.

Theorem 3.2 (Aronov et al. [21]). *The complexity of the union of n convex polyhedra in \mathbb{R}^3 with a total of s facets is $O(n^3 + sn \log n)$. This complexity can be $\Omega(n^3 + sn\alpha(n))$ in the worst case.*

It is interesting to note that the above bound is cubic only in the number of polyhedra, but it is only *linear* in s . (Compare with the bound $O(n^2 + s\alpha(n))$ for the case of convex polygons in the plane [19].) The cubic term disappears in the special case where the polyhedra in \mathcal{C} are

Minkowski sums of pairwise-disjoint convex polyhedra with another fixed convex polyhedron—see the following subsection for details.

The proof of Theorem 3.2 given in [21] is rather technical; we highlight two of its key ingredients that have been proved useful in some other contexts too. We note that techniques for analyzing the union of objects in \mathbb{R}^3 (and in higher dimensions) are rather scarce; we will mention some of these techniques as we encounter instances in which they can be exploited.

Special quadrilaterals and special cubes—Junctions in the union. Let \mathcal{C} be a family of n convex polyhedra with a total of s facets, and let C_1, C_2, C_3 be three members of \mathcal{C} with the following property. There exists a facet F_1 of C_1 , such that $Q = F_1 \cap C_2 \cap C_3$ is a quadrilateral, having two opposite edges on ∂C_2 and two opposite edges on ∂C_3 , and no other member of \mathcal{C} intersects Q . In this case, we call Q a *special quadrilateral*; see Figure 20.




Figure 20. A special quadrilateral in the union of polyhedra.

Aronov et al. [20, 21] have introduced this notion, and showed that, for arbitrary collections \mathcal{C} as above, the complexity of the union of \mathcal{C} is $O^*(n^2 + Q(n, s))$, where $Q(n, s)$ is an upper bound on the number of special quadrilaterals in any subcollection of \mathcal{C} . They have then shown that, for collections \mathcal{C} of Minkowski sums of pairwise disjoint convex polyhedra with another fixed polyhedron, $Q(n, s) = O(ns)$.

Pach et al. [100] have extended this notion to that of *special cubes*, where a special cube is an intersection of three members of \mathcal{C} , which has the combinatorial structure of a cube, where each of the three intersecting polyhedra contributes a pair of opposite facets to the intersection, and no other member of \mathcal{C} meets the “cube”. Pach et al. have shown that the union complexity of \mathcal{C} is roughly n^2 + the number of special cubes in any subcollection of \mathcal{C} .

Thus, the problem of bounding the complexity of the union reduces to that of bounding the number of special quadrilaterals or cubes. This has been done for special quadrilaterals, in the context of Minkowski sums of pairwise disjoint convex polyhedra, in [20], using a fairly intricate topological argument, and for special cubes, in the context of arbitrarily aligned nearly congruent cubes in [100], using a plane sweeping argument.

Charging schemes. This technique can be used in a variety of scenarios. Here we sketch in a special case how it can be applied to convex polyhedra.

Let \mathcal{C} be a family of n convex polyhedra in \mathbb{R}^3 , each with a constant number of facets, and consider the problem of bounding the complexity of $\mathcal{U}(\mathcal{C})$. Clearly, the number of vertices of $\mathcal{U}(\mathcal{C})$ that are vertices of some member of \mathcal{C} or are double-intersection points lying on an edge of some member and on a facet of another is $O(n^2)$. Therefore, we have to bound the number of triple-intersection points on the boundary of $\mathcal{U}(\mathcal{C})$, i.e., points that belong to the boundaries of three distinct members of \mathcal{C} . (Assuming that the sets are in general position, no point can belong to the boundaries of more than three distinct members.)

Consider the arrangement \mathcal{A} induced by the boundaries of the polyhedra in \mathcal{C} . Define the *level* of a vertex of this arrangement to be the number of members of \mathcal{C} that contain v in their interior. The number of triple-intersection vertices at level i is denoted by $V_i = V_i(\mathcal{C})$. We have to bound $V_0(\mathcal{C})$, that is, the number of triple-intersection vertices at level 0.

Each vertex v of the union is incident to three edges of the arrangement \mathcal{A} , each leading *away* from the union boundary; that is, each such edge is contained in the intersection segment of two of the facets containing v , and leads into the interior of the third polyhedron. We follow each of these edges, and charge v to the three vertices that are the other endpoints of these edges. See Figure 21.




Figure 21. The charging scheme. The three dashed edges emanating from v lead into the interior of the union.

The favorable situation is when all three charged vertices are triple-intersection vertices at level 1. In this case, each of them can be charged at most three times (see Figure 21), so the number of charging vertices v of this kind is at most $V_1(\mathcal{C})$, the number of triple-intersection vertices at level 1. The case where one of the charged vertices is not a triple intersection is easy, because there are only $O(n^2)$ such vertices (in the entire arrangement), and each is charged only a constant number of times, so there can be at most $O(n^2)$ charging vertices v of this kind.

Let us denote by $V_0^*(\mathcal{C})$ the number of vertices v at level 0 for which at least one of the charged vertices is a triple intersection vertex that also lies at level 0. We thus obtain the inequality

$$V_0(\mathcal{C}) \leq V_1(\mathcal{C}) + V_0^*(\mathcal{C}) + O(n^2).$$

The main difficulty is in obtaining a nearly quadratic bound on $V_0^*(\mathcal{C})$. In general, this is impossible: for instance, when the members of \mathcal{C} are large and thin plates that form a grid, one can easily check that $V_0(\mathcal{C}) = V_0^*(\mathcal{C}) = \Theta(n^3)$. Suppose, however, that we are in a favorable situation, and have somehow managed to show that $V_0^*(\mathcal{C}) = O^*(n^2)$. Then we get

$$V_0(\mathcal{C}) \leq V_1(\mathcal{C}) + O^*(n^2).$$

Let \mathcal{R} be a random subset of \mathcal{C} , obtained by removing one element uniformly at random. An easy calculation shows that

$$\mathbf{E}(V_0(\mathcal{R})) = \frac{n-3}{n}V_0(\mathcal{C}) + \frac{1}{n}V_1(\mathcal{C}).$$

Combining this with the preceding inequality, and writing $V_0(m)$ for the maximum value of $V_0(\mathcal{C})$ for $|\mathcal{C}| = m$, we obtain

$$\frac{1}{n}V_0(\mathcal{C}) \leq \frac{1}{n}V_1(\mathcal{C}) + O^*(n) = \mathbf{E}(V_0(\mathcal{R})) - \frac{n-3}{n}V_0(\mathcal{C}) + O^*(n) \leq V_0(n-1) - \frac{n-3}{n}V_0(\mathcal{C}) + O^*(n),$$

or

$$\frac{n-2}{n}V_0(n) \leq V_0(n-1) + O^*(n).$$

Dividing this by $(n-1)(n-2)$, we obtain a telescoping recurrence that solves to $V_0(n) = O^*(n^2)$. If the overhead term $V_0^*(\mathcal{C})$ is strictly $O(n^2)$, the recurrence solves to $O(n^2 \log n)$.

The above scheme is a special instance of a technique developed by Tagansky [113, 114], built upon earlier cruder charging schemes. As already noted, the real challenge is to bound $V_0^*(\mathcal{C})$. One way of doing so is to apply the charging scheme repeatedly, where in the next stage we want to bound the number of level-1 edges of the arrangement with both endpoints at level 0, by charging them to more complex local structures that have three level-0 vertices connected by two level-1 edges, and so on. This multi-stage scheme ends when the overhead term is the number of special quadrilaterals defined above (or can be pushed further until the overhead term counts the number of special cubes). See [100, 113] for details.

An interesting feature, hidden in this quick review, is that the only bottleneck in the analysis is to bound the number of special quadrilaterals. In contrast, it is relatively easy to give a quadratic upper bound for the number of “special polygons” with more than four vertices, where such a polygon Q is the intersection of a facet of one member of \mathcal{C} with two other members of \mathcal{C} , so that no fourth member of \mathcal{C} meets Q .

3.4 Robots with three degrees of freedom: Complexity of the free space

A special class of problems that involve unions in three dimensions arises in motion planning for robots with three degrees of freedom. Recall that in this case the *configuration space*, which represents all possible placements of the given robot B , is 3-dimensional, and each obstacle O in the physical environment (the workspace of B) generates an *expanded obstacle* (or C -obstacle) O^* ,

which is the locus of all placements of B at which it intersects O . The free portion \mathbb{F} of the configuration space is then the complement of the union of the C -obstacles.

In this subsection we review several results that arise in this context. As already discussed in the general setting, the naive bound on the complexity of \mathbb{F} is cubic in the number of possible contacts between features of B and features of the obstacles. In many instances, this bound can be attained, but there are several special cases where better, nearly quadratic, bounds can be established.

Let B be a robot with three degrees of freedom, so that each placement of B can be parametrized by three real parameters. For simplicity, let us assume that the configuration space, the set of all placements of B , is \mathbb{R}^3 . Two special cases of such a robot that we consider are: a planar object that is allowed to translate and rotate amid obstacles in \mathbb{R}^2 , and a three-dimensional object allowed only to translate amid obstacles in \mathbb{R}^3 . Bounding the complexity of \mathbb{F} in the former case was one of the first applications that led to the study of the union of objects in \mathbb{R}^3 [89, 90].

Translation and rotation in 2D. Let B be a convex polygon in \mathbb{R}^2 that is allowed to translate and rotate in the plane amid a set $\mathcal{O} = \{O_1, \dots, O_n\}$ of pairwise openly disjoint obstacles, each of which is a convex polygon, with a total of s vertices. To parametrize the configuration space, we fix a point $o \in B$ and a ray ρ emanating from o and rigidly attached to B . A placement of B is then parametrized by a point $(a, b, \tan(\theta/2)) \in \mathbb{R}^3$, where (a, b) are the coordinates of o and θ is the counterclockwise angle from the x -axis to ρ ; see Figure 22(a). A placement of B is *free* if B does not intersect any obstacle at this placement, and *semi-free* if B makes contact with one or more obstacles at this placement but does not intersect the interior of any obstacle. A generic contact between the boundaries of B and an obstacle can be represented by a pair (σ, ω) where σ is a vertex of B and ω is an edge of the obstacle, or σ is an edge of B and ω is a vertex of the obstacle.




Figure 22. (a) Representation of a placement of B . (b) A triple contact.

For each obstacle O_i , let C_i denote the corresponding expanded obstacle, which is the set of placements at which B intersects O_i ; C_i is a semi-algebraic set whose complexity depends on that of B and O_i . As noted, putting $\mathcal{C} = \{C_1, \dots, C_n\}$, we have $\mathbb{F} = \mathbb{R}^3 \setminus \mathcal{U}(\mathcal{C})$, and $\partial\mathbb{F}$ is the locus of all semi-free placements. A vertex of \mathbb{F} formed by the intersection of the boundaries of three expanded obstacles corresponds to a placement of B at which it makes three distinct contacts with the obstacles (see Figure 22(b)); these placements are referred to as *critical* semifree placements

or *critical vertices* (of \mathbb{F}). It can easily be argued that if B is a polygon with k vertices then the complexity of \mathbb{F} is proportional to $k^2 s^2$ plus the number of critical vertices.




Figure 23. A nonconvex polygon with $\Omega(n^3)$ critical semi-free placements.

If B is a nonconvex polygon, then \mathbb{F} can have $\Omega(n^3)$ critical vertices, as shown in Figure 3.4. However, the bound improves considerably when B is convex. For instance, if B is a line segment, then, as shown in several early works (around the mid 1980s), \mathbb{F} has $O(s^2)$ vertices [27, 90]. Recently, Agarwal et al. [2] improved the bound to $O(ns)$; this improved bound holds even if the obstacles in \mathcal{O} are not pairwise disjoint. In fact, if the obstacles are pairwise disjoint, then the number of critical vertices of \mathbb{F} is only $O(n^2 + s)$; however, the number of vertices formed by a pair of expanded obstacles (edge-face intersection points) can be $\Theta(ns)$.

The main result for this scenario is:

Theorem 3.3 (Leven and Sharir [89]). *If B is a convex k -gon, then the complexity of \mathbb{F} is $O(ks\lambda_6(ks))$.*

Since the number of combinatorially different contacts between B and the obstacles is $\Theta(ks)$, this bound is nearly quadratic in the number of contacts. Here is a brief sketch of the analysis in [89]. Let ϕ be a (vertex-edge or edge-vertex) contact between the boundaries of B and of an obstacle, and let $\Gamma_\phi \subseteq \mathbb{R}^3$ denote the set of all placements of B at which the contact ϕ is made; Γ_ϕ is a two-dimensional algebraic surface patch. For each contact ϕ , we define a family \mathcal{C}_ϕ of $O(ks)$ pseudo-halfplanes in Γ_ϕ , where each pseudo-halfplane h represents placements at which ϕ is made and another contact ϕ' is “violated”— ϕ' is made at placements on ∂h , and B and the corresponding obstacle intersect at placements within h . The boundaries of any pair of these pseudo-halfplanes intersect in at most six points. The main observation in the analysis is that if B and the obstacles are in general position, and if the complexity of \mathbb{F} is more than $k^2 s^2$, then at least a fraction of the vertices of \mathcal{C} are vertices of $\mathcal{U}(\mathcal{C}_\phi)$, over all contacts ϕ . By Theorem 2.1, $\kappa(\mathcal{C}_\phi)$ is $O(\lambda_6(ks))$, for each ϕ , and thus the complexity of \mathbb{F} is $O(ks\lambda_6(ks))$. The details of the proof can be found in [89].

Combining this overall approach with a few new observations, and performing a more careful analysis, Agarwal and Gujgunte [5] have recently improved the bound on the complexity of \mathbb{F} to $O(ks\lambda_6(kn))$.

Van der Stappen et al. [111] studied the case in which the obstacles are fat, and proved a linear bound on the complexity of \mathbb{F} , under certain reasonable assumptions. We refer the reader to their paper for more details.

Translational motion planning in \mathbb{R}^3 . Let B be a convex object in \mathbb{R}^3 that is allowed to translate amid a set $\mathcal{O} = \{O_1, \dots, O_n\}$ of n obstacles, each of which is a convex polytope. We fix a point $o \in B$ and represent a placement of B by specifying the coordinates (x, y, z) of o . As mentioned in the introduction, the expanded obstacle C_i generated by O_i is now the Minkowski sum $O_i \oplus (-B)$ of O_i and the reflected image $-B$ of B , and, as usual, $\mathbb{F} = \mathbb{R}^3 \setminus \mathcal{U}(\mathcal{C})$, where $\mathcal{C} = \{C_1, \dots, C_n\}$. This has led to the extensive study of the complexity of (and algorithms for constructing) the union of a family \mathcal{C} of Minkowski sums of this kind in \mathbb{R}^3 .

In the planar case, the crucial property of such a collection of Minkowski sums was that each pair of boundaries cross at most twice, so the collection is a family of pseudo-disks. The corresponding property in \mathbb{R}^3 (assuming general position) is that each pair of boundaries intersect in a single connected closed curve [76]. However, a triple of boundaries can intersect in an arbitrarily large number of points, which makes the analysis of the union complexity considerably harder than in the plane. Near-quadratic bounds have been established for only a few special cases, summarized in the following theorems. (In each part, s effectively denotes the overall complexity of the individual Minkowski sums in \mathcal{C} , but its precise definition is slightly different in each case.)

Theorem 3.4 (Halperin and Yap [69]). *If B is a cube, the complexity of $\mathcal{U}(\mathcal{C})$ (and thus of \mathbb{F}) is $O(s^2\alpha(s))$, where s denotes the overall number of faces of the original polytopes in \mathcal{O} .*

Theorem 3.5 (Aronov and Sharir [20]). *If B is a convex polytope, the complexity of $\mathcal{U}(\mathcal{C})$ is $O(ns \log n)$, where s denotes the overall number of faces of the polytopes in \mathcal{C} . There exist constructions where the union complexity is $\Omega(ns\alpha(n))$.*

Theorem 3.6 (Agarwal and Sharir [11]). *If B is a ball, the complexity of $\mathcal{U}(\mathcal{C})$ is $O^*(s^2)$, where s is the total number of faces of the polytopes in \mathcal{O} . In particular taking \mathcal{O} to be a set of n lines in 3-space, the complexity of the union of n congruent infinite cylinders in \mathbb{R}^3 is $O^*(n^2)$.*

The proofs of these theorems are rather different, and each of them is very technical. The proof of Theorem 3.4 is based on ideas similar to those used by Leven and Sharir [89]. The proof of Theorem 3.5 is a special case of the analysis of the union of arbitrary convex polyhedra, given in [21], where the main new ingredient is an intricate topological argument that shows that the number of special quadrilaterals in the union is $O(ns)$. The proof of Theorem 3.6 is the most involved; it uses a rather complicated charging scheme, and is based on several geometric observations that reduce the problem to that of bounding the complexity of sandwich regions between upper and lower envelopes of bivariate functions.

These results lead to a few natural questions that remain elusive:

Open Problem 4. *What is the maximum complexity of the union of n congruent cones or tori?*

Open Problem 5. *What is the maximum complexity of the union of n cylinders of different radii?*

Although the upper bound for all these cases is conjectured to be $O^*(n^2)$, no subcubic upper bounds are known to date.

3.5 Union of fat objects

Similar to the planar case, a compact convex object C is called α -fat, for some constant $\alpha \geq 1$, if the ratio between the radii of the smallest enclosing ball and of the largest inscribed ball of C is at most α . In this subsection we review some of the recent (and slightly less recent) developments in the analysis of the complexity of the union of fat objects in 3-space.

Union of axis-aligned cubes. We begin by considering the simple case of axis-aligned cubes.

Theorem 3.7 (Boissonnat et al. [28]; see also [31]). *The complexity of the union of n axis-aligned cubes in \mathbb{R}^3 is $O(n^2)$. The bound reduces to $O(n)$ if the cubes are of the same (or nearly the same) size. Both bounds are tight in the worst case.*

Proof: This result is sufficiently simple to allow us to provide a complete proof. We only need to count the number of vertices of the union that are incident to three facets of three distinct respective cubes; the number of all other vertices (of the entire arrangement of the cube boundaries) is only $O(n^2)$. Let v be such a vertex, incident to facets F_1, F_2, F_3 of three distinct respective cubes C_1, C_2, C_3 , so that C_1 is the largest cube among them. Follow the intersection segment $F_2 \cap F_3$ from v into C_1 . This segment has to end within C_1 , at a point that lies on an edge of C_2 or C_3 , and on the remaining facet F_3 or F_2 . The number of such terminal points is clearly only $O(n^2)$, and each of them can be encountered in such a tracing from only a constant number of vertices v of the union. Hence, the number of these vertices, and thus the complexity of the union, is $O(n^2)$. The proof for congruent cubes is also simple, but we omit it. \square

Union of arbitrary nearly congruent cubes. If the cubes are not axis-parallel, the problem becomes much harder. Pach et al. [100] have studied the case where the cubes have equal (or “almost equal”) size, and have shown³ that the complexity of their union is $O^*(n^2)$. The key observation in their analysis is that one can lay out a regular grid, where the size of its cells is somewhat smaller than that of the given cubes, so that (a) each cube meets only a constant number of cells, and (b) no two opposite facets of a cube meet the same cell. This allows us to consider the union separately in each cell, and observe that in each cell the union becomes a union of unbounded halfspaces, (right-angle) dihedral wedges, and (orthant-like) trihedral wedges. The analysis thus reduces to that of bounding the complexity of the union of such wedges. The main technical ingredient in the analysis of [100] is:

³We do not highlight this result, because it is now subsumed by the result of Ezra and Sharir [60], which we will shortly present.

Theorem 3.8 (Pach et al. [100]). *The complexity of the union of n α -fat dihedral wedges is $O^*(n^2)$, where the constant of proportionality depends on (the hidden ε and on) α .*

Pach et al. were not as successful in analyzing the complexity of the union of α -fat trihedral wedges (wedges whose solid angle is at least α), for any constant $\alpha > 0$, and managed to establish a nearly quadratic bound only when the wedges are “substantially fat”, a case that includes wedges formed at a vertex of a cube, but not wedges formed at a vertex of a regular tetrahedron.

A major observation in the analysis of [100] is that, for any triple of α -fat dihedral wedges, there are many directions u , such that a plane orthogonal to u cuts each of the three wedges in a cross-section which is itself α' -fat, for some $\alpha' > 0$ that depends on α . This allows the analysis to proceed by sweeping the given wedges by a plane, considering only those wedges that meet the plane in fat cross-sections, and by analyzing critical events when the boundaries of three of the swept wedges become concurrent. (Finding such a good sweeping direction for triples of trihedral wedges is harder; in general this is impossible unless the wedges are really “substantially fat”.) The analysis then combines the study of special cubes (as reviewed above) with some other tricks, to conclude that the complexity of the union of such wedges is nearly quadratic.

To recap, the technique of [100], powerful as it were, could not handle cubes of arbitrary sizes (the grid reduction does not work then), nor could it handle other kinds of fat polyhedra (for which the wedges formed at their vertices are not sufficiently fat); even the special case of regular tetrahedra remained open. Both of these shortcomings have recently been overcome by Ezra and Sharir [60], who have obtained a nearly quadratic bound for the complexity of the union of n arbitrary fat tetrahedra. We will review this result below, and we note that it immediately implies a nearly quadratic bound for the union complexity of n arbitrary cubes in \mathbb{R}^3 (of arbitrary sizes).

Union of fat tetrahedra and of cubes. We say that a tetrahedron is α -fat if each of its solid angles is at least α . This definition is compatible with the other standard definitions of fatness. Specifically, the ratio between the radii of the smallest enclosing ball and the largest inscribed ball of an α -fat tetrahedron is at most $\alpha' = O(1/\sqrt{\alpha})$. Conversely, if this ratio is at most α' for some tetrahedron, then it must be α -fat with $\alpha = \Omega(1/(\alpha')^2)$.




Figure 24. An α -fat trihedral wedge and an α -fat dihedral wedge.

To simplify the presentation, let us assume for the moment that we are given a collection \mathcal{C} of n α -fat tetrahedra of *nearly equal size*, meaning that the diameters of the tetrahedra in \mathcal{C} are within

some constant ratio of each other. Then there is an easy grid-based argument, similar to the one used for nearly equal cubes, to reduce the analysis of their union to that of the union of α -fat trihedral wedges, namely, trihedral wedges whose solid angles are at least α (see Figure 24). Specifically, assume, for simplicity, that all the diameters lie in the interval $[1, c]$, for a fixed constant c . We lay out a grid of sufficiently small (but constant) cell size, so that (a) for any tetrahedron τ of \mathcal{C} and any grid cell Δ , at most three facets of τ meet Δ , and (b) each tetrahedron in \mathcal{C} crosses only $O(1)$ grid cells. Hence, within each grid cell Δ , we need to bound the complexity of the union of some n_Δ α -fat trihedral wedges (which can also degenerate further to dihedral wedges or halfspaces).

Suppose that we have a bound of $O^*(m^2)$ on the complexity of the union of m α -fat trihedral wedges, with a constant of proportionality that depends on α . This bound, combined with the above reduction, implies that the complexity of the union of n nearly-equal α -fat tetrahedra is $\sum_\Delta O^*(n_\Delta^2) = O^*(n^2)$. (The case of nearly equal cubes is now an easy corollary of this result.)

The analysis in [60] applies also to the case where the tetrahedra have arbitrary sizes (diameters). It is somewhat involved, and we sketch here only some of its highlights. To simplify the presentation, we only consider the case of fat trihedral wedges. Let then \mathcal{C} be a family of n α -fat trihedral wedges. The main technical tool in the analysis of [60] is the following lemma.

Lemma 3.9 (Ezra and Sharir [60]). (a) Let R be a set of r planes in \mathbb{R}^3 , and let W be an arbitrary trihedral wedge. The number of cells of $\mathcal{A}(R)$ that meet all three facets of W is only $O(r)$.

(b) Let P be a convex polyhedron with r facets in \mathbb{R}^3 , and let Σ be the family of tetrahedra into which P is decomposed by the Dobkin-Kirkpatrick hierarchical decomposition scheme [38]. The number of tetrahedra in Σ that meet all three facets of W is only $O(\log r)$.

See Figure 25(b) for an illustration. Note that the lemma applies to any trihedral wedge, not necessarily fat. Note also that the planar version of the lemma is trivial: In an arrangement of r lines in the plane, at most *one* cell can meet all three edges of a given triangle (Figure 25(a)). As another trivial variant in the plane, the number of cells that meet both sides of a *wedge* is $O(r)$.




Figure 25. (a) In the plane, only one cell of the arrangement can meet all three edges of a given triangle. (b) In \mathbb{R}^3 , many cells (but only $O(r)$) can meet all three facets of a given trihedral wedge (a schematic view from the apex of the wedge).

Lemma 3.9 suggests the following recursive decomposition scheme. Take a random sample

R of r planes that support the facets of the wedges of \mathcal{C} . Construct the arrangement $\mathcal{A}(R)$ and decompose each of its cells into tetrahedra, using the Dobkin-Kirkpatrick scheme. We obtain a decomposition Ξ of 3-space into $O(r^3)$ tetrahedra, with the property that for each wedge W of \mathcal{C} , the number of tetrahedra that meet all three facets of W is only $O(r \log r)$. Hence, on average, each tetrahedron is crossed by at most $O\left(\frac{n}{r^2} \log r\right)$ wedges of \mathcal{C} with this property. Moreover, the standard theory of random sampling [71] allows us to assume that R has the property that each of the simplices of Ξ is crossed by the boundaries of at most $O\left(\frac{n}{r} \log r\right)$ wedges of \mathcal{C} . To recap, we obtain $O(r^3)$ subproblems, each involving at most $O\left(\frac{n}{r} \log r\right)$ wedges, of which, on average, only $O\left(\frac{n}{r^2} \log r\right)$ are trihedral wedges, and the rest are dihedral wedges (or halfspaces).

To obtain the asserted near-quadratic bound, the analysis in [60] applies the decomposition repeatedly, taking r to be a sufficiently large constant, and involves a rather careful counting of the vertices that are not passed down the recursion. Instead of reconstructing this somewhat involved analysis, let us consider the following simpler quick-and-dirty approach. If we choose $r = \sqrt{n}$, we obtain $O(n^{3/2})$ subproblems, each involving some number m of trihedral wedges (which is only logarithmic on average), and $O^*(n^{1/2})$ dihedral wedges. The number of vertices of the union that are formed by three dihedral wedges is $O^*((n^{1/2})^2) = O^*(n)$ [100], and the number of vertices that lie on the boundary of at least one trihedral wedge is $O^*(mn)$ (using a rough quadratic bound for each trihedral wedge separately). Summing over the tetrahedra, and using the fact that the m 's sum to $O(nr \log r) = O^*(n^{3/2})$, yields the overall bound of $O^*(n^{5/2})$ for the complexity of the union. With the more careful and recursive analysis in [60], this bound drops to $O^*(n^2)$.

The above analysis can also be applied to the case of fat tetrahedra rather than wedges (Lemma 3.9 obviously carries over to this case), but then considerably more effort is needed to count vertices that are not passed down the main recursion. The analysis of [60] culminates at the following result.

Theorem 3.10 (Ezra and Sharir [60]). *The complexity of the union of n arbitrary α -fat tetrahedra in \mathbb{R}^3 is $O^*(n^2)$, where the constant of proportionality depends on α (and on the hidden $\varepsilon > 0$).*

Union of α -round objects. Let \mathcal{C} be a family of n α -round objects in \mathbb{R}^3 . That is, for each $C \in \mathcal{C}$, any point $p \in \partial C$ is incident to a ball of radius α times the diameter of C , which is fully contained in C . We first consider a special case of this problem, in which we further assume that the diameter of each member of \mathcal{C} is between 1 and D , for some constant D . This will allow us to introduce one of the techniques for analyzing unions in three dimensions. We may therefore assume that all the balls used in the definition of roundness are of the same radius α .

These assumptions are easily seen to imply that if v is a vertex of the union, incident to the boundaries of three sets C_1, C_2, C_3 , then, with at least some constant probability, a random direction u has the property that the line through p at direction u intersects each of the sets C_1, C_2, C_3 in an interval of length at least $\alpha' = \beta\alpha$, for some sufficiently small but absolute constant $\beta > 0$. We call a vertex satisfying the above property for a direction u a u -feasible vertex. To prove that $\kappa(\mathcal{C})$ is $O^*(n^2)$, it suffices to establish a near-quadratic bound on the number of u -feasible vertices for any fixed direction u . Suppose, without loss of generality, that u is the z -direction.

Partition \mathbb{R}^3 into horizontal slabs of width α' . For each slab S , let $\mathcal{C}_S^+, \mathcal{C}_S^-$ denote the family




Figure 26. Reducing the union of nearly equal α -round objects in \mathbb{R}^3 to sandwich regions.

of objects of \mathcal{C} that intersect the top and bottom boundaries of S respectively, clipped to within S . Moreover, retain, for each object $C \in \mathcal{C}_S^-$, the portion of its top boundary consisting of those points $x \in S$ for which the vertical segment from x to the bottom boundary of S is fully contained in C . Apply a symmetric trimming process to the bottom boundaries of the objects of \mathcal{C}_S^+ . Put $n_s = |\mathcal{C}_S^+| + |\mathcal{C}_S^-|$. It can be checked that each u -feasible vertex that lies in the slab S is a vertex of the sandwich region between (the trimmed portions of) the upper envelope of the top boundaries of objects in \mathcal{C}_S^- and the lower envelope of the bottom boundaries of objects in \mathcal{C}_S^+ . See Figure 26. As mentioned in Section 3.2, the number of vertices on the sandwich region is $O^*(n_s^2)$. However, $\sum_S n_s = O(n)$, because each object in \mathcal{C} can cross only $O(1)$ slabs. Therefore $\kappa(\mathcal{C}) = \sum_S O^*(n_s^2) = O^*(n^2)$.

The above argument fails when the diameters of the objects in \mathcal{C} differ significantly. This has been overcome by Aronov et al. [18], who have extended the nearly quadratic bound to this case, using a somewhat more involved technique, which is also based on reducing the problem to sandwich regions between envelopes.

Theorem 3.11 (Aronov et al. [18]; Agarwal and Sharir [11]). *Let \mathcal{C} be a family of n 3-dimensional α -round objects of constant description complexity. Then the complexity of $\mathcal{U}(\mathcal{C})$ is $O^*(n^2)$.*

In spite of all the progress reviewed in this section, the following general question is still open.

Open Problem 6. *What is the maximum complexity of the union of n α -fat (convex) objects of constant description complexity in \mathbb{R}^3 ?*

4 Beyond Three Dimensions

In higher dimensions, the problem of bounding the complexity of the union of geometric objects becomes even more complicated, and only very few results are known, which we duly review here.

Union of pseudo-halfspaces. As already mentioned in the introduction, the complexity of the union of n halfspaces (each bounded by a hyperplane) in \mathbb{R}^d is $O(n^{\lfloor d/2 \rfloor})$. For pseudo-halfspaces

(regions lying above or below the graph of some continuous function of constant description complexity), the bounds are not that small. As shown by Sharir [108], the complexity of the lower (or upper) envelope of n $(d - 1)$ -variate functions of constant description complexity is $O^*(n^{d-1})$. Hence, the union of n pseudo-halfspaces, all of which are lower (or all upper) is $O^*(n^{d-1})$.

However, this is not known to hold in the mixed case, where some pseudo-halfspaces are lower and some are upper, in $d \geq 5$ dimensions. As in two and three dimensions, we seek bounds on the complexity of the sandwich region between a lower and an upper envelope, which turns out to be a hard problem when the dimension d increases. An $O^*(n^3)$ bound on the complexity of the sandwich region in \mathbb{R}^4 was proved by Koltun and Sharir [81]. (As in the three-dimensional case, this is based on a nearly cubic bound, established in [81], on the complexity of the overlay of two minimization and/or maximization diagrams, this time of trivariate functions.) This yields a nearly cubic bound on the complexity of the union of n pseudo-halfspaces of constant description complexity in \mathbb{R}^4 . The problem of whether sandwich regions have asymptotically smaller complexity than that of the entire arrangement is still open for $d \geq 5$.

Open Problem 7. *What is the maximum complexity of the union of n pseudo-halfspaces of constant description complexity in \mathbb{R}^d , for $d \geq 5$?*

Linearization. The so-called linearization technique can be used to bound the complexity of the union of certain classes of regions, by transforming these regions to halfspaces. Specifically, let $f(x, a)$ be a $(d + p)$ -variate polynomial, with $x \in \mathbb{R}^d$ and $a \in \mathbb{R}^p$. Let a^1, \dots, a^n be n points in \mathbb{R}^p , and set $\mathcal{F} = \{f_i(x) \equiv f(x, a^i) \mid 1 \leq i \leq n\}$; thus \mathcal{F} is a collection of d -variate polynomials. For each i , let C_i be one of the two regions $f_i \geq 0$ or $f_i \leq 0$, and set $\mathcal{C} = \{C_1, \dots, C_n\}$. Suppose that $f(x, a)$ can be expressed in the form

$$f(x, a) = \psi_0(a) + \psi_1(a)\varphi_1(x) + \dots + \psi_k(a)\varphi_k(x), \quad (2)$$

where ψ_0, \dots, ψ_k are p -variate polynomials and $\varphi_1, \dots, \varphi_k$ are d -variate polynomials. We define the map $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^k$ by

$$\varphi(x) = (\varphi_1(x), \dots, \varphi_k(x)).$$

Then the image $\Gamma = \{\varphi(x) \mid x \in \mathbb{R}^d\}$ of \mathbb{R}^d is a d -dimensional surface in \mathbb{R}^k (assuming $k \geq d$), and for any $a \in \mathbb{R}^p$, $f(x, a)$ maps to the k -variate linear function

$$h_a(y_1, \dots, y_k) = \psi_0(a) + \psi_1(a)y_1 + \dots + \psi_k(a)y_k,$$

in the sense that for any $x \in \mathbb{R}^d$, $f(x, a) = h_a(\varphi(x))$. The region C_i maps to one of the two halfspaces bounded by the hyperplane h_{a_i} (more precisely, to the intersection of Γ with such a halfspace). We refer to k as the *dimension of the linearization* φ , and say that \mathcal{F} admits a linearization of dimension k . Agarwal and Matoušek [8] describe an algorithm that computes a linearization of the smallest dimension under certain mild assumptions. If \mathcal{F} admits a linearization of dimension k , then the complexity of $\mathcal{U}(\mathcal{C})$ is bounded by the complexity of the union of n halfspaces in \mathbb{R}^{k+1} ,

and it is therefore $O(n^{\lceil k/2 \rceil})$. The most popular example of linearization is perhaps the so-called *lifting transform* (see Section 2.2 for the planar case), which is constructed from the polynomial

$$f(x, a) = (x_1 - a_1)^2 + \cdots + (x_d - a_d)^2 - a_{d+1}^2,$$

for $x \in \mathbb{R}^d$ and $a \in \mathbb{R}^{d+1}$. The resulting lifting transformation itself is then

$$\varphi(x) = (x_1, x_2, \dots, x_d, x_1^2 + \cdots + x_d^2).$$

This maps \mathbb{R}^d to the standard paraboloid $x_{d+1} = x_1^2 + \cdots + x_d^2$ in \mathbb{R}^{d+1} , and a ball in \mathbb{R}^d is mapped to a halfspace in \mathbb{R}^{k+1} , which implies that the complexity of the union of n balls in \mathbb{R}^d is $O(n^{\lceil d/2 \rceil})$.

Axis-aligned cubes. Boissonnat et al. [28] provide an upper bound of $O(n^{\lceil d/2 \rceil})$ 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 (or nearly equal) size. The complexity of the union of n simply-shaped convex bodies in \mathbb{R}^d with a common interior point o is $O^*(n^{d-1})$, which follows from the observation that the boundary of their union can be interpreted as the upper envelope of n $(d-1)$ -variate functions (in spherical coordinates about o). A slightly refined bound for polyhedra in \mathbb{R}^3 with a common interior point was given in [73].

Koltun and Sharir [81] extended Theorem 3.11 to \mathbb{R}^4 , by proving that the complexity of the union of n convex α -round objects in \mathbb{R}^4 with equal diameters is $O^*(n^3)$. These results have been further generalized by Aronov et al. [18] for (not necessarily convex) α -round objects. The only obstacle to obtaining analogous results in \mathbb{R}^d , for $d \geq 5$, is our inability to establish sharp upper bounds on the complexity of sandwich regions (as discussed above) in $d \geq 5$ dimensions.

Open Problem 8. *What is the maximum complexity of the union of n α -round objects of constant description complexity in \mathbb{R}^d , with (or without) nearly equal diameters, for $d \geq 5$?*

5 Generalized Voronoi Diagrams

Voronoi diagrams are closely related to unions of geometric objects, in the following manner. Let \mathcal{C} be a set of n pairwise disjoint convex objects in \mathbb{R}^d , each of constant description complexity, and let ρ be a metric (or a *convex distance function* [115]). For a point $x \in \mathbb{R}^d$, let $\Phi(x)$ denote the set of objects of \mathcal{C} that are nearest to x , i.e.,

$$\Phi(x) = \{C \in \mathcal{C} \mid \rho(x, C) \leq \rho(x, C') \text{ for each } C' \in \mathcal{C}\}.$$

The *Voronoi diagram* $\text{Vor}_\rho(\mathcal{C})$ of \mathcal{C} under the metric ρ (sometimes also simply denoted as $\text{Vor}(\mathcal{C})$) is the partition of \mathbb{R}^d into maximal connected regions of various dimensions, so that, for each region V , the set $\Phi(x)$ is the same for all $x \in V$. For each full-dimensional region (cell), $\Phi(\cdot)$ generally consists of a single site C , and the cell is called the Voronoi cell of C . For $i = 1, \dots, n$, let γ_i be the graph of the function $x_{d+1} = \rho(x, C_i)$, for $x \in \mathbb{R}^d$, and set $\Gamma = \{\gamma_i\}_{i=1}^n$. Edelsbrunner and Seidel

[47] made the rather obvious observation that $\text{Vor}_\rho(\mathcal{C})$ is the *minimization diagram* of Γ , that is, the projection onto \mathbb{R}^d of the lower envelope of the surfaces in Γ .

To see the connection between generalized Voronoi diagrams and unions of objects, let \mathcal{C} and ρ be as above (say, for the 3-dimensional case). For an object $C \in \mathcal{C}$ and a parameter $r \geq 0$, define $\mathbb{B}(C, r) = \{x \in \mathbb{R}^3 \mid \rho(x, C) \leq r\}$. For a fixed r , the union $K_r = \bigcup_{C \in \mathcal{C}} \mathbb{B}(C, r)$ is the region consisting of all points $x \in \mathbb{R}^3$ whose smallest ρ -distance from a site in \mathcal{C} is at most r . This in turn can be interpreted as a “cross-section” of $\text{Vor}(\mathcal{C})$ —it is in fact a cross-section at height $x_4 = r$ of the lower envelope of the corresponding collection Γ . Moreover, for each site $C_i \in \mathcal{C}$, the intersection of ∂K_r with the Voronoi cell of C_i is equal to the intersection of $\partial \mathbb{B}(C_i, r)$ with that cell.

In general, if the metric ρ is a norm or a distance function induced by some convex body B , that is, $\rho(x, y) = \min\{\lambda \mid y \in x + \lambda B\}$, the resulting “balls” $\mathbb{B}(C, r)$ are the *Minkowski sums* $C \oplus (-rB)$, for $C \in \mathcal{C}$ (the minus sign is superfluous if ρ is a metric). Thus the union of Minkowski sums of this kind is a substructure of the corresponding Voronoi diagram. Of course, this connection also holds in any higher dimension.

One immediate conclusion is that the complexity of $\text{Vor}(\mathcal{C})$ is at least as large as that of K_r . In practice, establishing a tight bound on the latter complexity is a considerably easier task, and in many instances the corresponding question concerning the complexity of the entire Voronoi diagram is still open. For instance, consider the case in which \mathcal{C} is a set of lines in 3-space, and ρ is the Euclidean metric. Then the expanded sites $\mathbb{B}(C, r)$, for $C \in \mathcal{C}$, are n congruent infinite cylinders in \mathbb{R}^3 , of radius r . As mentioned above, it is shown in [11] that the complexity of the union K_r of these cylinders is $O^*(n^2)$, but it is a major open problem to establish a similar nearly quadratic bound on the complexity of $\text{Vor}_\rho(\mathcal{C})$ (see an Open Problem below). There are (rare) cases in which the complexity of the entire Voronoi diagram is an order of magnitude larger than that of K_r . For example, the complexity of the multiplicatively weighted Voronoi diagram of a point set in the plane can have quadratic complexity [25], while the size of K_r in this case is only linear.

In the classical case, when ρ is the Euclidean metric and the objects in \mathcal{C} are singletons (points), the graphs of the distance functions $\rho(x, C_i)$ can be replaced by a collection of n hyperplanes in \mathbb{R}^{d+1} , using a straightforward linearization technique, without affecting the minimization diagram. Hence, the maximum possible complexity of $\text{Vor}(\mathcal{C})$ is $O(n^{\lceil d/2 \rceil})$, and this is tight in the worst case (see, e.g., [78, 106]). In more general settings, though, this reduction is not possible, and the complexity of the Voronoi diagram can be much higher. Applying the observation of [47], and the bounds in Section 4 on the complexity of lower (or upper) envelopes, we obtain that, under reasonable assumptions on ρ and on the objects in \mathcal{C} , the complexity of the Voronoi diagram is $O^*(n^d)$. While this bound is nontrivial (the trivial one is $O(n^{d+1})$), in general it is not expected to be tight. For example, in the case of planar Voronoi diagrams, this bound is near-quadratic, but the complexity of “almost every” planar Voronoi diagram is only $O(n)$. Nevertheless, as mentioned above, for certain “pathological” distance functions, the corresponding planar Voronoi diagram can indeed have quadratic complexity [25].

Voronoi diagrams of points in \mathbb{R}^3 . As noted above, the complexity of the Euclidean Voronoi diagram of n points in \mathbb{R}^3 is $\Theta(n^2)$. It has been a long-standing open problem to determine whether a similar quadratic or nearly quadratic bound holds in \mathbb{R}^3 for more general objects and metrics (here the known bounds on the complexity of lower envelopes only give an upper bound of $O^*(n^3)$). The problem stated above calls for improving this bound by roughly another factor of n . Since we are aiming for a bound that is “two orders of magnitude” better than the complexity of $\mathcal{A}(\Gamma)$, this appears to be a considerably more difficult problem than the problem of bounding the complexity of lower envelopes. The only hope of making progress here is to exploit the special structure of the distance functions $\rho(x, C)$.

Boissonnat et al. [28] have shown that the maximum complexity of the L_1 -Voronoi diagram of a set of n points in \mathbb{R}^3 is $\Theta(n^2)$. Tagansky [113] proved that the complexity of the three-dimensional Voronoi diagram of point sites under a general polyhedral convex distance function (induced by a polytope with $O(1)$ facets) is $O(n^2 \log n)$. The bound was subsequently improved by Icking and Ma [74] to $O(n^2)$.

Voronoi diagrams of lines in \mathbb{R}^3 . Let ρ be a convex distance function in \mathbb{R}^3 whose unit ball is a convex polytope with a constant number of facets. (Recall that not every distance function ρ is necessarily a metric— ρ fails to be symmetric if the defining polytope (its unit ball) is not centrally symmetric.) Chew et al. [32] showed that the complexity of the Voronoi diagram of n lines in \mathbb{R}^3 with respect to ρ is $O(n^2 \alpha(n) \log n)$. Clearly, the L_1 and L_∞ metrics satisfy the above assumptions. In these special cases, the best known lower bound for the complexity of the diagram is $\Omega(n^2 \alpha(n))$. Koltun and Sharir [82] extended the theorem of Chew et al. [32] to arbitrary collections of pairwise disjoint line segments and triangles, where the respective upper bounds on the complexity of the diagram are $O(n^2 \alpha(n) \log n)$ and $O^*(n^2)$.

As already mentioned, in spite of some recent progress, little is known about the complexity of the Euclidean Voronoi diagram of lines in 3-space.

Open Problem 9. *What is the maximum complexity of the Euclidean Voronoi diagram of n lines or triangles in \mathbb{R}^3 ?*

If the input lines have a constant number of orientations, then the complexity of their Euclidean Voronoi diagram is $O^*(n^2)$, as shown by Koltun and Sharir [80]. Dwyer [40] has shown that the expected complexity of the (Euclidean) Voronoi diagram of a set of n randomly selected lines in \mathbb{R}^3 is only $O(n^{3/2})$. For the general case, a recent work by Everett *et al.* [55] sheds some light on the geometric and topological structure of bisectors and trisectors defined by a pair (resp., triple) of lines in space.

Voronoi diagram of moving points in the plane. An interesting special case of generalized Voronoi diagrams are *dynamic Voronoi diagrams* for moving points in the plane. Let \mathcal{C} be a set of n points in the plane, each moving along some line at some fixed velocity. The goal is to bound the number of combinatorial changes of the Euclidean diagram $\text{Vor}(\mathcal{C})$ over time. This dynamic

Voronoi diagram can easily be transformed into a static three-dimensional Voronoi diagram, by adding the time t as a third coordinate. The points become lines in \mathbb{R}^3 , and the “metric”⁴ is a distance function induced by a horizontal disk (that is, the distance from a point $p(x_0, y_0, t_0)$ to a line ℓ is the Euclidean distance from p to the point of intersection of ℓ with the horizontal plane $t = t_0$). Cubic or nearly cubic bounds are known for this problem, even under more general settings [62, 64, 108], but subcubic bounds are known only in some very special cases [30, 79].

Open Problem 10. *What is the maximum complexity of the dynamic (Euclidean) Voronoi diagram of n moving points in the plane? What if all points move at the same speed?*

A recent study by Agarwal et al. [3] presents some necessary conditions for the diagram to have large complexity, and thereby offers some intuition on why “typical” dynamic Voronoi diagrams are expected to have small complexity. The expected complexity of the dynamic Voronoi diagram of n points moving *randomly* in the plane is $O(n^{3/2})$ [41].

Voronoi diagrams in higher dimensions. Next, consider the problem of bounding the complexity of generalized Voronoi diagrams in higher dimensions. As mentioned above, when the objects in \mathcal{C} are n points in \mathbb{R}^d and the metric is Euclidean, the complexity of $\text{Vor}(\mathcal{C})$ is $O(n^{\lceil d/2 \rceil})$. As d increases, this becomes significantly smaller than the naive $O(n^{d+1})$ bound or the improved bound, $O^*(n^d)$, obtained by viewing the Voronoi diagram as a lower envelope in \mathbb{R}^{d+1} . The same bound of $O(n^{\lceil d/2 \rceil})$ has been obtained in [28] for the complexity of the L_∞ -diagram of n points in \mathbb{R}^d ; this bound too was shown to be tight in the worst case. It was thus tempting to conjecture that the maximum complexity of generalized Voronoi diagrams in higher dimensions is close to $n^{\lceil d/2 \rceil}$. However, this conjecture was disproved by Aronov [16], who established a lower bound of $\Omega(n^{d-1})$. The sites used in his construction are lower-dimensional flats, and the distance is either Euclidean or a polyhedral convex distance function. (It is interesting that the lower bound in Aronov’s construction depends on the affine dimension $0 \leq k \leq d - 2$ of the sites: It is $\Omega(\max\{n^{k+1}, n^{\lceil (d-k)/2 \rceil}\})$.) Thus, for $d = 3$, this lower bound does not contradict the conjecture made above, that the complexity of generalized Voronoi diagrams should be at most near-quadratic in this case. Also, in higher dimensions, the conjecture mentioned above is still not refuted when the sites are singleton points. However, very little is known about this problem. For instance, the following problem remains open.

Open Problem 11. *What is the maximum complexity of the Voronoi diagram of a set of points in \mathbb{R}^d under polyhedral metrics or convex distance functions whose unit balls have $O(1)$ facets?*

Finally, for the general case, Aronov’s construction still leaves a gap of roughly a factor of n between the best known upper and lower bounds, and thus suggests the conjecture that the complexity of such diagrams is $O^*(n^{d-1})$. This is still a major open problem.

⁴This is not really a metric, because the distance between two points is defined only when they have the same t -coordinate.

Open Problem 12. *Is it true that, for a set \mathcal{C} of n pairwise disjoint convex sites of constant description complexity in \mathbb{R}^d , and for a metric (or convex distance function) ρ whose unit ball has constant description complexity, the complexity of $\text{Vor}(\mathcal{C})$ is always $O^*(n^{d-1})$?*

Medial axis. A special case of Voronoi diagrams is the *medial axis*. Here we are given a region C with a complex boundary, e.g., a (not necessarily convex) polyhedron with many faces. We regard each feature of ∂C (vertex, edge, face) as a separate site, and consider the Voronoi diagram of these sites within the interior of C . The lower-dimensional faces of the diagram yield a “skeletal” representation of C , which has several advantages in practice [23]. A particularly difficult, and still open, special case is the following.

Open Problem 13. *Let \mathcal{C} be a collection of n balls in \mathbb{R}^3 . What is the maximum complexity of the medial axis of $\mathcal{U}(\mathcal{C})$? What is the maximum complexity when all the balls have the same radius?*

In fact (see Amenta and Kolluri [15]), it suffices to bound the complexity of the Voronoi diagram of the vertices of $\mathcal{U}(\mathcal{C})$ within the union. Since the union may have $\Theta(n^2)$ vertices in the worst case, and the complexity of the Voronoi diagram of that many points in \mathbb{R}^3 can in general be quadratic in their number, a naive upper bound on the complexity of the medial axis is $O(n^4)$. However, the best known lower bound is only quadratic, and closing the gap between the bounds is a challenging open problem.

Voronoi diagrams of regularly sampled points. Dwyer [39] proved that the expected size of the (Euclidean) Voronoi diagram of a set of uniformly distributed random points inside a ball in \mathbb{R}^d is linear. Later, Erickson [52, 53] studied the complexity of the Voronoi diagram of a point set P in \mathbb{R}^3 in terms of the *spread* of P , which is the ratio of the largest and the smallest pairwise distances between the points of P . Erickson proved that the complexity of the Voronoi diagram of a set of points in \mathbb{R}^3 with spread Δ is $O(\Delta^3)$. He also proved that this bound is tight in the worst case, by showing an $\Omega(n^{3/2})$ lower bound for a set of n point nicely distributed on a cylinder, so that their spread is $O(\sqrt{n})$.

Motivated by the problem of surface reconstruction from a set of sample points, a considerable amount of work has been dedicated to bounding the complexity of the Voronoi diagram of a set of regularly sampled points on a surface Γ in \mathbb{R}^3 . Golin and Na [63] showed that the expected complexity of the Voronoi diagram of n uniformly distributed random points on a polyhedral surface Γ in \mathbb{R}^3 is $O(n \log^4 n)$. A set $S \subset \Gamma$ is called an (ε, λ) -*sample* if any ball of radius ε centered at a point of Γ contains at least one and at most λ points of S . Attali and Boissonnat [22] proved that if $S \subset \mathbb{R}^3$ is an (ε, λ) -sample on a polyhedral surface, then the size of its Voronoi diagram is linear. Attali *et al.* [24] proved that if $S \subset \mathbb{R}^3$ is an (ε, λ) -sample of size n on a generic (smooth) surface, then its Voronoi diagram has $O(n \log n)$ complexity. Roughly speaking, a surface is *generic* if the points on the surface at which one of the principal curvature is locally maximal, form a finite set of curves with bounded length; spheres and cylinders are not generic surfaces. Note that the assumption of genericity is probably crucial in the proof of [24], because of Erickson’s lower-bound

construction for points on cylinders [52]. Recently, Amenta *et al.* [14] proved that the complexity of the Voronoi diagram of n nicely distributed points on a convex p -dimensional polyhedron in \mathbb{R}^d is $O(n^{(d-1)/p})$; see the original paper for details on the sampling condition and other issues. We conclude this discussion by mentioning the following open problem:

Open Problem 14. *What is the maximum complexity of the Voronoi diagram of a set of n points regularly sampled on (or sufficiently near) a smooth manifold in \mathbb{R}^d ?*

6 Discussion

In this survey we have reviewed the extensive work concerning the complexity of the union of a family of geometric objects in two, three, and higher dimensions. We also reviewed the state of the art concerning the complexity of generalized Voronoi diagrams in three and higher dimensions.

However, we have not discussed algorithms for computing the union of geometric objects. Several deterministic divide-and-conquer, randomized divide-and-conquer, and randomized incremental algorithms have been proposed to compute the union for a variety of special cases [11, 57, 59]. Motivated by many applications, considerable work has addressed related issues, such as computing the volume of the union, the gradient of the volume of union of balls in \mathbb{R}^3 (regarding the volume of the union as a function from \mathbb{R}^{3n} to \mathbb{R}), or certain geometric or topological properties of the union of balls. It is beyond the scope of this survey to review these results, and we refer the reader to [6, 43, 45, 46, 88, 98] and the references therein for a sample of such results.

References

- [1] http://en.wikipedia.org/wiki/Constructive_solid_geometry
- [2] P. K. Agarwal, D. Chen, and M. Sharir, On short segment stabbers, unpublished manuscript, 2007.
- [3] P. K. Agarwal, J. Gao, L. Guibas, V. Koltun, and M. Sharir, Stable Delaunay graphs, unpublished manuscript, 2007.
- [4] P. K. Agarwal, E. F. Grove, T. M. Murali, J. S. Vitter, Binary space partitions for fat rectangles, *SIAM J. Comput.* **29** (2000), 1422–1448.
- [5] P. K. Agarwal and S. Gujgunte, On the complexity of the free space of a convex polygon moving amid convex planar obstacles, unpublished manuscript, 2007.
- [6] P.K. Agarwal, H. Kaplan, and M. Sharir, Computing the volume of the union of cubes in three dimensions, *Proc. 23rd Annu. Sympos. Comput. Geom.*, 2007, 294–301.
- [7] P. K. Agarwal, M. Katz, and M. Sharir, Computing depth order and related problems, *Comput. Geom. Theory Appl.* **5** (1995), 187–206.
- [8] P. K. Agarwal and J. Matoušek, On range searching with semi-algebraic sets, *Discrete Comput. Geom.*, **11** (1994), 393–418.

- [9] 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.
- [10] P. K. Agarwal and M. Sharir, Efficient algorithms for geometric optimization, *ACM Comput. Surveys* **30** (1998), 412–458.
- [11] 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.
- [12] N. Alon and S. Smorodinsky, Conflict-free colorings of shallow disks, *Proc. 22nd Annu. Sympos. Comput. Geom.*, 2006, 41–43.
- [13] 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.
- [14] N. Amenta, D. Attali, and O. Devillers, Complexity of the Delaunay triangulation for points on lower-dimensional polyhedra, *Proc. 18th Annu. ACM-SIAM Sympos. Discrete Algo.*, 2007, 1106–1113.
- [15] N. Amenta and R. K. Kolluri, The medial axis of a union of balls, *Comput. Geom. Theory Appl.*, **20** (2001), 25–37.
- [16] B. Aronov, A lower bound on Voronoi diagram complexity, *Inform. Process. Lett.* **83** (2002), 183–185.
- [17] B. Aronov, A. Efrat, D. Halperin, and M. Sharir, On the number of regular vertices of the union of Jordan regions, *Discrete Comput. Geom.* **25** (2001), 203–220.
- [18] B. Aronov, A. Efrat, V. Koltun, and M. Sharir, On the union of κ -round objects in three and four dimensions, *Discrete Comput. Geom.* **36** (2006), 511–526.
- [19] B. Aronov and M. Sharir, The common exterior of convex polygons in the plane, *Comput. Geom. Theory Appl.* **8** (1997), 139–149.
- [20] B. Aronov and M. Sharir, On translational motion planning of a convex polyhedron in 3-space, *SIAM J. Comput.* **26** (1997), 1785–1803.
- [21] B. Aronov, M. Sharir, and B. Tagansky, The union of convex polyhedra in three dimensions, *SIAM J. Comput.* **26** (1997), 1670–1688.
- [22] D. Attali and J.-D. Boissonnat, A linear bound on the complexity of the Delaunay triangulation of points on polyhedral surfaces, *Discrete Comput. Geom.* **31** (2004), 369–384.
- [23] D. Attali, J.-D. Boissonnat, and H. Edelsbrunner, Stability and computation of the medial axis — a state-of-the-art report, in *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration* (T. Möller, B. Hamann, and B. Russell, eds.), Springer-Verlag, to appear.
- [24] D. Attali, J.-D. Boissonnat, and A. Lieutier, Complexity of the Delaunay triangulation of points on surfaces the smooth case, *Proc. 19th Annu. Sympos. Comput. Geom.*, 2003, 201–210.
- [25] F. Aurenhammer and R. Klein, Voronoi diagrams, in *Handbook of Computational Geometry* (J.-R. Sack and J. Urrutia, eds.), Elsevier Science, Amsterdam, 2000, pp. 201–290.

- [26] J. L. Bentley, Algorithms for Klee’s rectangle problems. Unpublished notes, Computer Science Department, Carnegie Mellon University, 1977.
- [27] B. K. Bhattacharya and J. Zorbas, Solving the two-dimensional findpath problem using a line-triangle representation of the robot, *J. Algorithms* **9** (1988), 449–469.
- [28] 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.
- [29] H. Brönnimann and O. Devillers, The union of unit balls has quadratic complexity, even if they all contain the origin, Tech. Report 3758, INRIA, Sophia-Antipolis, 1999.
- [30] L. P. Chew, Near-quadratic bounds for the L_1 Voronoi diagram of moving points, *Comput. Geom. Theory Appl.* **7** (1997), 73–80.
- [31] L. P. Chew, D. Dor, A. Efrat, and K. Kedem, Geometric pattern matching in d -dimensional space, *Discrete Comput. Geom.* **21** (1999), 257–274.
- [32] 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.
- [33] Ch. Chojnacki (A. Hanani), Über wesentlich unplättbare Kurven im dreidimensionalen Raume, *Fund. Math.* **23** (1934), 135–142.
- [34] K. L. Clarkson and K. Varadarajan, Improved approximation algorithms for geometric set covers, *Discrete Comput. Geom.* **37** (2007), 43–58.
- [35] M. de Berg and C. Gray, Vertical ray shooting and computing depth orders for fat objects *Proc. 17th Annu. ACM-SIAM Sympos. Discrete Algo.*, 2006, 494–503.
- [36] M. de Berg, M. Katz, F. van der Stappen, and J. Vleugels, Realistic input models for geometric algorithms, *Algorithmica* **34** (2002), 81–97.
- [37] E. D. Demaine, J. S. B. Mitchell, and J. O’Rourke, The Open Problems Project, <http://www.cs.smith.edu/~orourke/TOPP/>.
- [38] D. P. Dobkin and D. G. Kirkpatrick, Determining the separation of preprocessed polyhedra — a unified approach. *Proc. 17th Intl. Colloq. Auto. Lang. Program.*, 1990, 400–413.
- [39] R. A. Dwyer: Higher-dimensional Voronoi diagrams in linear expected time, *Discrete Comput. Geom.* **6** (1991), 343–367.
- [40] R. A. Dwyer, Voronoi diagrams of random lines and flats, *Discrete Comput. Geom.* **17** (1997), 123–136.
- [41] R. A. Dwyer, Voronoi diagrams and convex hulls of random moving points, *Discrete Comput. Geom.* **23** (2000), 343–365.
- [42] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer Verlag, Heidelberg, 1987.
- [43] H. Edelsbrunner, The union of balls and its dual shape, *Discrete Comput. Geom.* **13** (1995), 415–440.
- [44] H. Edelsbrunner, L. Guibas, J. Hershberger, J. Pach, R. Pollack, R. Seidel, M. Sharir, and J. Snoeyink, On arrangements of Jordan arcs with three intersections per pair, *Discrete Comput. Geom.* **4** (1989), 523–539.

- [45] H. Edelsbrunner and P. Koehl, The weighted volume derivative of a space filling diagram, *Proc. Natl. Acad. Sci.* **100** (2003), 2203–2208.
- [46] H. Edelsbrunner and P. Koehl, The geometry of biomolecular solvation, in *Combinatorial and Computational Geometry* (J. E. Goodman, J. Pach, and E. Welzl, eds.), Cambridge Univ. Press, New York, 2005, 243–276.
- [47] H. Edelsbrunner and R. Seidel, Voronoi diagrams and arrangements, *Discrete Comput. Geom.* **1** (1986), 25–44.
- [48] A. Efrat, The complexity of the union of (α, β) -covered objects, *SIAM J. Comput.* **34** (2005), 775–787.
- [49] A. Efrat and M. Katz, On the union of κ -curved objects, *Comput. Geom. Theory Appl.* **14** (1999), 241–254.
- [50] 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.
- [51] A. Efrat and M. Sharir, The complexity of the union of fat objects in the plane, *Discrete Comput. Geom.* **23** (2000), 171–189.
- [52] J. Erickson, Nice point sets can have nasty Delaunay triangulations, *Discrete Comput. Geom.* **30** (2003), 109–132.
- [53] J. Erickson, Dense point sets have sparse Delaunay triangulations, *Discrete Comput. Geom.* **33** (2005), 83–115.
- [54] G. Even, Z. Lotker, D. Ron, and S. Smorodinsky, Conflict-free colorings of simple geometric regions with applications to frequency assignment in cellular networks, *SIAM J. Comput.* **33** (2003), 94–136.
- [55] H. Everett, S. Lazard, D. Lazard, and M. Safey El Din, The Voronoi diagram of three lines, *Proc. 23rd Annu. Sympos. Comput. Geom.*, 2007, 255–264.
- [56] E. Ezra, *Geometric Arrangements: Substructures and Algorithms*, Ph.D. Dissertation, School of Computer Science, Tel Aviv University, August 2007.
- [57] E. Ezra, D. Halperin, and M. Sharir, Speeding up the incremental construction of the union of geometric objects in practice, *Comput. Geom. Theory Appl.* **27** (2004), 63–85.
- [58] E. Ezra, J. Pach, and M. Sharir, On regular vertices on the union of planar convex objects, *Proc. 23rd Annu. Sympos. Comput. Geom.*, 2007, 220–226.
- [59] E. Ezra and M. Sharir, Output sensitive construction of the union of triangles, *SIAM J. Comput.* **34** (2005), 1331–1351.
- [60] E. Ezra and M. Sharir, Almost tight bound for the union of fat tetrahedra in three dimensions, *Proc. 48th Annu. IEEE Sympos. Foundat. Vcomput. Sci.*, 2007, to appear.
- [61] J. D. Foley, A. van Dam, S. Feiner, and J. F. Hughes, *Computer Graphics — Principles and Practice*, 2nd ed., Addison-Wesley, 1992.
- [62] J.-J. Fu and R.C.T. Lee, Voronoi diagrams of moving points in the plane, *Int. J. Comput. Geom. Appl.* **1** (1991), 23–32.

- [63] M. J. Golin and H.-S. Na, The probabilistic complexity of the Voronoi diagram of points on a polyhedron, *Proc. 18th Annu. Sympos. Comput. Geom.*, 2002, 209–216.
- [64] L. J. Guibas, J. S. B. Mitchell, and T. Roos, Voronoi diagrams of moving points in the plane, *Proc. 17th Internat. Workshop Graph-Theoret. Concepts Comput. Sci., Lecture Notes Comput. Sci.*, Vol. 570, Springer-Verlag, 1991, 113–125.
- [65] L. Guibas, M. Sharir, and S. Sifrony, On the general motion planning problem with two degrees of freedom, *Discrete Comput. Geom.* **4** (1989), 491–521.
- [66] D. Halperin, *Algorithmic Motion Planning via Arrangements of Curves and of Surfaces*, Ph.D. Dissertation, Computer Science Department, Tel-Aviv University, Tel Aviv, 1992.
- [67] D. Halperin and M. H. Overmars, Spheres, molecules, and hidden surface removal, *Comput. Geom. Theory Appl.* **11** (1998) 83–102.
- [68] D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions, with applications to visibility in terrains, *Discrete Comput. Geom.* **12** (1994), 313–326.
- [69] D. Halperin and C.-K. Yap, Combinatorial complexity of translating a box in polyhedral 3-space, *Comput. Geom. Theory Appl.* **9** (1998), 181–196.
- [70] S. Har-Peled and S. Smorodinsky, Conflict-free coloring of points and simple regions in the plane. *Discrete Comput. Geom.* **34** (2005), 47–70.
- [71] D. Haussler and E. Welzl, Epsilon-nets and simplex range queries, *Discrete Comput. Geom.* **2** (1987), 127–151.
- [72] D. P. Huttenlocher and K. Kedem, Computing the minimum Hausdorff distance for point sets under translation, *Proc. 6th Annu. Sympos. Comput. Geom.*, 1990, pp. 340–349.
- [73] D. Huttenlocher, K. Kedem, and M. Sharir, The upper envelope of Voronoi surfaces and its applications, *Discrete Comput. Geom.* **9** (1993), 267–291.
- [74] C. Icking and L. Ma, A tight bound for the complexity of Voronoi diagrams under polyhedral convex distance functions in 3D, *Proc. 33rd Annu. ACM Sympos. Theory of Comput.*, 2001, 316–321.
- [75] M. J. Katz, 3-D vertical ray shooting and 2-D point enclosure, range searching, and arc shooting amidst convex fat objects, *Comput. Geom. Theory Appl.* **8** (1997), 299–316.
- [76] 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.
- [77] V. Klee, Can the measure of $\bigcup [a_i, b_i]$ be computed in less than $O(n \log n)$ steps? *Amer. Math. Monthly* **84** (1977), 284–285.
- [78] V. Klee, On the complexity of d -dimensional Voronoi diagrams, *Arch. Math. (Basel)*, **34** (1980), 75–80.
- [79] V. Koltun, Ready, set, go! The Voronoi diagram of moving points that start from a line, *Inf. Process. Lett.* **89** (2004), 233–235.
- [80] V. Koltun and M. Sharir, Three-dimensional Euclidean Voronoi diagrams of lines with a fixed number of orientations, *SIAM J. Comput.* **32** (2003), 616–642.

- [81] V. Koltun and M. Sharir, The partition technique for the overlay of envelopes, *SIAM J. Comput.* **32** (2003), 841–863.
- [82] V. Koltun and M. Sharir, Polyhedral Voronoi diagrams of polyhedra in three dimensions, *Discrete Comput. Geom.* **31** (2004), 83–124.
- [83] J. Komlós, J. Pach, and G. Woeginger, Almost tight bounds for ε -nets, *Discrete Comput. Geom.* **7** (1992), 163–173.
- [84] K. Kratky, The area of intersection of n equal circular disks, *J. Phys. A: Math. Gen.* **11** (1978), 1017–1024.
- [85] M. van Kreveld, On fat partitioning, fat covering and the union size of polygons, *Comput. Geom. Theory Appl.* **9** (1998), 197–210.
- [86] D. H. Laidla, W. B. Trumbore, and J. F. Hughes, Constructive solid geometry for polyhedral objects, *SIGGRAPH Proceedings*, 1986, 161–170.
- [87] B. Lee and F. M. Richards, The interpretation of protein structure: Estimation of static accessibility, *J. Molecular Biology*, **55** (1971), 379–400.
- [88] J. van Leeuwen and D. Wood, The measure problem for rectangular ranges in d -space, *J. Algorithms* **2** (1981), 282–300.
- [89] D. Leven and M. Sharir, On the number of critical free contacts of a convex polygonal object moving in two-dimensional polygonal space, *Discrete Comput. Geom.*, **2** (1987), 255–270.
- [90] D. Leven and M. Sharir, An efficient and simple motion planning algorithm for a ladder amidst polygonal barriers, *J. Algorithms* **8** (1987), 192–215.
- [91] L. Lovász, J. Pach, and M. Szegedy, On Conway’s thrackle conjecture, *Discrete Comput. Geom.* **18** (1997), 369–376.
- [92] J. Matoušek, Epsilon-nets and computational geometry, in: *New Trends in Discrete and Computational Geometry*, (J. Pach, ed.), Springer Verlag, Berlin, 1993, 69–89.
- [93] J. Matoušek, J. Pach, M. Sharir, S. Sifrony, and E. Welzl, Fat triangles determine linearly many holes, *SIAM J. Comput.* **23** (1994), 154–169.
- [94] J. Matoušek, R. Seidel, and E. Welzl, How to net a lot with little: Small epsilon nets for disks and halfspaces, in *Proc. 6th Annu. Sympos. Comput. Geom.*, 1990, 16–22.
- [95] P. G. Mezey, Molecular surfaces, in: *Reviews in Computational Chemistry* (K. B. Lipkowitz and D. B. Boyd, eds.), Vol. **1**, VCH Publishers, 1990.
- [96] P. McMullen, On the upper-bound conjecture for convex polytopes, *J. Combinat. Theory, Ser. B* **10** (1971), 187–200.
- [97] P. McMullen and G. C. Shephard, *Convex Polytopes and the Upper Bound Conjecture*, Cambridge University Press, Cambridge, England, 1971.
- [98] M. Overmars and C.K. Yap, New upper bounds in Klee’s measure problem, *SIAM J. Comput.* **20** (1991), 1034–1045.

- [99] J. Pach and P. K. Agarwal, *Combinatorial Geometry*, John Wiley & Sons, New York, NY, 1995.
- [100] J. Pach, I. Safruti and M. Sharir, The union of congruent cubes in three dimensions, *Discrete Comput. Geom.* **30** (2003), 133–160.
- [101] J. Pach and M. Sharir, On the boundary of the union of planar convex sets, *Discrete Comput. Geom.* **21** (1999), 321–328.
- [102] J. Pach and G. Tardos, On the boundary complexity of the union of fat triangles, *SIAM J. Comput.* **31** (2002), 1745–1760.
- [103] M. J. Pelsmajer, M. Schaefer, D. Štefankovič, Removing even crossings, *J. Combinat. Theory, Ser. B* **97** (2007), 489–500.
- [104] F. P. Preparata and M. I. Shamos, *Computational Geometry: An Introduction*, Springer-Verlag, New York, 1985.
- [105] F. M. Richards, Areas, volumes, packing, and protein structure, *Annu. Rev. Biophys. Bioeng.* **6** (1977), 151–176.
- [106] R. Seidel, Exact upper bounds for the number of faces in d -dimensional Voronoi diagrams, in *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, DIMACS Ser. Discrete Math. Theoret. Comput. Sci. **4**, (P. Gritzman and B. Sturmfels, eds.), AMS, Providence, RI, 1991, 517–530.
- [107] M. Sharir, Algorithmic motion planning, in: *Handbook of Discrete and Computational Geometry, 2nd ed.* (J. E. Goodman and J. O’Rourke, eds.), CRC Press LLC, Boca Raton, FL, 2004, pp. 1037–1064.
- [108] M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* **12** (1994), 327–345.
- [109] M. Sharir and P. K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, NY, 1995.
- [110] F. van der Stappen, *Motion Planning amidst Fat Obstacles*, Ph.D. Dissertation, Faculteit Wiskunde & Informatica, Universiteit Utrecht, 1994.
- [111] F. van der Stappen, D. Halperin, and M. Overmars, The complexity of the free space for a robot moving amidst fat obstacles, *Comput. Geom. Theory Appl.* **3** (1993), 353–373.
- [112] J. Steiner, Einige Gesetze über die Theilung der Ebene und des Raumes, *J. Reine Angew. Math.*, **1** (1826), 349–364.
- [113] B. Tagansky, *The Complexity of Substructures in Arrangements of Surfaces*, Ph.D. Dissertation, Tel Aviv University, Tel Aviv, 1996.
- [114] B. Tagansky, A new technique for analyzing substructures in arrangements of piecewise linear surfaces, *Discrete Comput. Geom.* **16** (1996), 455–479.
- [115] A. C. Thompson, *Minkowski Geometry*, Cambridge University Press, Cambridge, 1996.
- [116] W. T. Tutte, Toward a theory of crossing numbers, *J. Combinat. Theory* **8** (1970), 45–53.

- [117] S. Whitesides and R. Zhao, K -admissible collections of Jordan curves and offsets of circular arc figures, *Technical Report SOCS 90.08*, McGill University, Montreal, 1990.
- [118] A. Wiernik and M. Sharir, Planar realizations of nonlinear Davenport-Schinzel sequences by segments, *Discrete Comput. Geom.* **3** (1988), 15–47.