On the complexity of the union of geometric objects

János Pach*

Courant Institute, New York University
and Hungarian Academy of Sciences

Abstract. Given a family $\mathcal{C}$ of regions bounded by simple closed curves in the plane, the *complexity* of their union is defined as the number of points along the boundary of $\cup\mathcal{C}$, which belong to more than one curve. Similarly, one can define the complexity of the union of 3-dimensional bodies, as the number of points on the boundary of the union, belonging to the surfaces of at least *three* distinct members of the family. We survey some 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 motion planning and computer graphics.

1 Introduction

Given a family $\mathcal{C}$ of polygons in the plane (or, a family of polyhedra in $\mathbb{R}^d$), the number of sides of their union (resp., the total number of its faces of all dimensions) is called the *combinatorial complexity*, or, simply, the *complexity* of $\cup\mathcal{C}$. This notion can be easily generalized to families of other simply shaped geometric objects with piecewise smooth boundaries. The complexity of a set is closely related to its *description size*, i.e., the number of parameters needed for its description.

Many basic problems in computational geometry related to motion planning [39, 40, 41], range searching [22, 18], computer graphics [1], and geographic information systems [6] lead to questions about the complexity of the boundary of the union of certain geometric objects. This was the motivation behind a lot of research during the past 15 years, establishing upper bounds for the complexity of the union of various objects.

We mention three simple examples.

1.1. Linear programming. Given a family $\mathcal{C}$ of n half-spaces in $\mathbb{R}^d$, we want to maximize a linear function on the boundary of their union, $\mathrm{Bd}(\cup\mathcal{C})$. The running time of the simplex algorithm, as well as many other naive solutions to this problem, is proportional to the total number of vertices of $\mathrm{Bd}(\cup\mathcal{C})$. According

* Supported by the National Science Foundation (USA) and the National Fund for Scientific Research (Hungary). e-mail: `pach@cims.nyu.edu`

to McMullen's Upper Bound Theorem [31], 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.

1.2. Motion planning amidst obstacles. Let R be a "robot," i.e., a convex polygonal object with a small number r of sides, which is free to translate amidst a collection of n convex polygonal obstacles, $C_1, C_2, \ldots, C_n$. Fix a reference point O (the origin) within R. In order to decide whether the robot can be moved from a fixed position to another without colliding with any of the obstacles, and to plan such a motion if it exists, we want to describe the space F of *free placements* of R, i.e., the locus of all positions of the reference point which correspond to placements of the robot, in which it does not intersect any C_i (see Fig. 1). It is easy to see that R intersects C_i if and only if the corresponding reference point belongs to the "expanded obstacle" $C_i' = C_i \oplus (-R)$, where $\oplus$ stands for the *Minkowski sum*, i.e.,

$$C_i' = \{c - r \mid c \in C_i, r \in R\}.$$

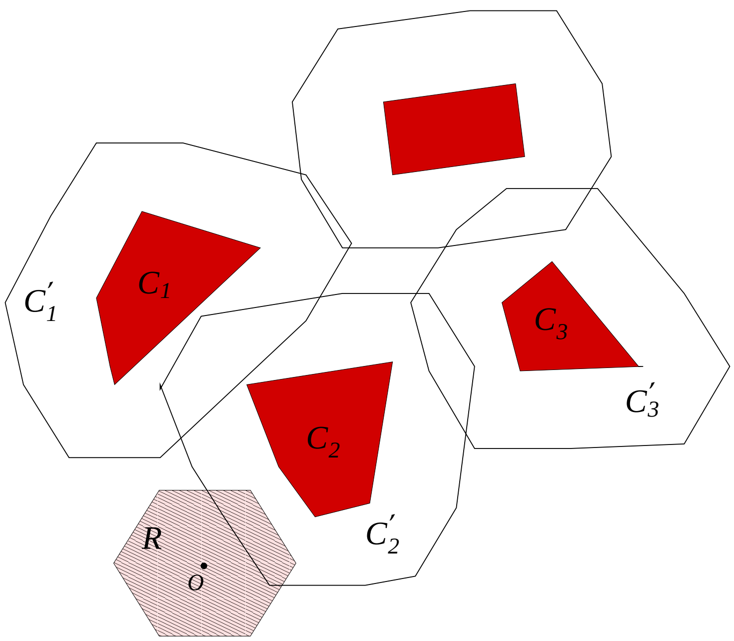

Figure 1: The space of free placements of the robot R.

In other words, the space F of free placements of the robot is equal to the complement of $\cup_{i=1}^{n} C_i'$. Therefore, the running time of any algorithm for the description of F is at least the number of vertices of the union of the expanded obstacles, and there are a number of efficient algorithms which almost achieve this time bound.

The above approach was initiated by Lozano-Pérez and Wesley [29], and it can also be applied to the situation when the robot is allowed to rotate [25], moves in three dimensions [4], has many arms [17], etc.

1.3. Overlay of maps. In most geographic information systems the data is stored in several thematic maps, each representing only one kind of information. E.g., there are separate maps for average temperature, for average precipitation, altitude, etc. A *face* of a map corresponds to a region where the given parameter is roughly constant. One often has to combine the information, i.e., to compute the overlay of two or more maps. This procedure results in a new map, more complex than the original ones. Consider, for example, two maps, and put their faces in a single list $F_1, F_2, \ldots, F_n$. Let F_i' denote a region obtained from F_i by very slightly shrinking it. It is easy to see now that the total number of vertices of the new combined map is proportional to the complexity of the boundary of $\cup_{i+1}^n F_i'$. Indeed, each vertex of the overlay will give rise to a "hole" determined by the F_i'.

2 An example – Translational motion planning

First we make some simple observations related to the motion planning problem stated above (1.2). Two curves, γ_1 and γ_2, are said to *cross* each other at a point, if at this point γ_1 passes from one side of γ_2 to the other.

Lemma 2.1. ([24]) *Let C_1, C_2, and R be convex bodies in the plane, and assume that C_1 and C_2 are disjoint. Then the boundaries of the Minkowski sums $C_1' = C_1 \oplus (-R)$ and $C_2' = C_2 \oplus (-R)$ cross at most twice.*

Let $\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$ be a family of simply connected regions in the plane bounded by simple closed curves. Assume, for simplicity, that these curves are in *general position*, i.e., any two of them cross only a finite number of times, no two touch each other, and no three pass through the same point. If any two distinct curves $\mathrm{Bd}(C_i)$ and $\mathrm{Bd}(C_j)$ cross at most twice, then $\mathcal{C}$ is usually called a family of *pseudo-disks*. A maximal connected piece of the boundary of $\cup\mathcal{C} = \cup_{i=1}^n C_i$, which belongs to the boundary of a member of $\mathcal{C}$, is said to be an *elementary arc*.

Theorem 2.2. ([24]) *Let $\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$ be a family of $n \geq 3$ pseudo-disks in the plane.*

Then the boundary of $\cup_{i=1}^n C_i$ consists of at most $6n - 12$ elementary arcs. This bound cannot be improved.

In the case when each C_i is a disk, the proof of Theorem 2.2 is straightforward (see Fig. 2). 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 $\mathrm{Bd}(C_i)$ and $\mathrm{Bd}(C_j)$ cross each other, and at least one of their crossing points belongs to $\mathrm{Bd}(\cup\mathcal{C})$. It is easy to verify that no two segments cross each other, i.e., the resulting graph G is *planar*. Therefore, G has at most $3n - 6$ edges, each of which corresponds to at most two crossings between the circles that belong to $\mathrm{Bd}(\cup\mathcal{C})$. Therefore, the number of crossings on $\mathrm{Bd}(\cup\mathcal{C})$, and hence the number of elementary arcs is at most $6n - 12$. Using the terminology introduced in the first paragraph of this paper, we obtain that the *complexity* of $\cup\mathcal{C}$ is at most linear in n.

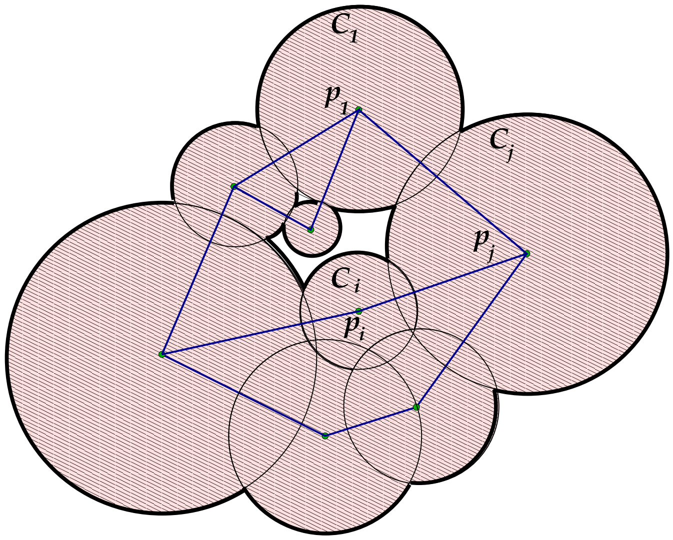

Figure 2: The proof of Theorem 2.2 for disks.

One can combine a well known algorithm of Bentley and Ottmann [7] for reporting all intersections in a collection of line segments with some standard divide-and-conquer techniques to obtain the following

Corollary 2.3. *There is an $O(nr\log^2 n)$-time algorithm for describing the space of free placements of a convex polygonal robot with r sides among a set of polygonal obstacles with a total of n sides, and for computing a collision-free path between any two given positions of the robot, if such a path exists.*

Leven and Sharir [27] reduced the $\log^2 n$ factor in the last result to $\log n$.

3 Allowing 3 intersections – Ackermann's function

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 if two closed curves are in general position, then they can cross only an *even* number of times. However, by a slight modification we obtain a meaningful question with a surprising answer.

Construct recursively an infinite sequence of integer-valued functions $A_1(n)$, $A_2(n), \ldots$ on the set of positive integers, as follows. Let $A_1(n) = 2n$ for every n. If A_k has already been defined for some $k \geq 1$, then let $A_{k+1}(n) = A_k^{(n)}(1)$. In other words, we n times iterate the function A_k, and take its value at 1. The function $A(n) = A_n(n)$ is called *Ackermann's function.* It grows so fast that $A(4)$ is a tower of 65536 2's! Consequently, the inverse of Ackermann's function is an extremely slowly growing funtion, whose all "practical" values are smaller than 4. For basic properties of these functions, see [43].

Theorem 3.1. ([11]) *Let $\{\gamma_1, \gamma_2, \ldots, \gamma_n\}$ be a family of simple curves in general position in the upper half-plane. 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.*

Then the boundary of $\cup_{i=1}^n C_i$ consists of at most $O(n\alpha(n))$ elementary arcs, where $\alpha(n)$ is the inverse of Ackermann's function. This bound is asymptotically tight.

Let $\{f_1, f_2, \ldots, f_n\}$ be a collection of real-valued functions, each defined on a subinterval of $\mathbb{R}$. For any $x \in \mathbb{R}$, consider the set $I(x)$ of all indices i, for which $f_i(x)$ is defined, and let

$$f(x) = \max_{i \in I(x)} f_i(x).$$

This partially defined function (and its graph) is called the *upper envelope* of the functions f_i (and of their graphs).

Perhaps the most important special case of the last theorem is the following result of Hart and Sharir, which answers a question of Atallah [5].

Theorem 3.2. ([20]) *The upper envelope of n non-vertical straight-line segments in the plane consists of at most $O(n\alpha(n))$ linear pieces.*

To verify that this result follows from Theorem 3.1, it is sufficient to notice that by attaching to each segment two vertical rays pointing downwards, one at each of its endpoints, we obtain a family of two-way infinite curves, any pair of which cross at most three times.

Wiernik and Sharir [46] showed that Theorem 3.2 is asymptotically tight.

A sequence of integers $a(i) \in \{1, 2, \ldots, n\}$ $(i = 1, 2, \ldots)$ is called a *Davenport-Schinzel sequence* (of order 3) if no two consecutive elements are the same, and there is no alternating subsequence of length 5, i.e., there are no indices $i_1 < i_2 < \ldots < i_5$ such that $a(i_1) = a(i_3) = a(i_5)$, $a(i_2) = a(i_4)$, but $a(i_1) \neq a(i_2)$.

In fact, Hart and Sharir proved that the length of any such sequence is at most $O(n\alpha(n))$. Theorem 3.2 can be easily deduced from this result. Indeed, if we order the linear pieces of the upper envelope of the segments $s_1, s_2, \ldots, s_n$ from left to right, and replace each piece by the index i of the segment s_i it belongs to, we obtain a Davenport-Schinzel sequence (see Fig. 3). To see this, we have to check only that this sequence has no alternating subsequence of length 5. However, this immediately follows from the fact that any two segments cross at most once.

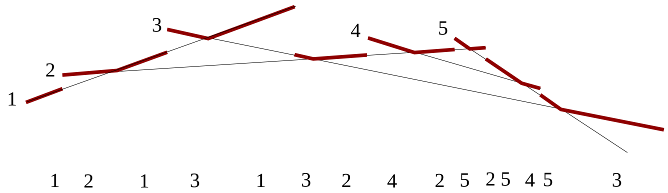

Figure 3: The upper envelope of segments and the corresponding Davenport-Schinzel sequence.

4 Well-behaved intersections – The role of parity

If we try to further weaken the condition in Theorem 2.2, assuming only that the boundaries of any two members of $\mathcal{C}$ cross at most *four* times, then the situation completely deteriorates. Consider a collection of n pairwise crossing line segments, no three of which pass through the same point, and enclose each of them in a very narrow triangle whose width is at most $\varepsilon > 0$. If ε is small enough, then every pair of triangles intersect in precisely four points, and all $4\binom{n}{2}$ intersection points belong to the boundary of their union (Fig. 4).

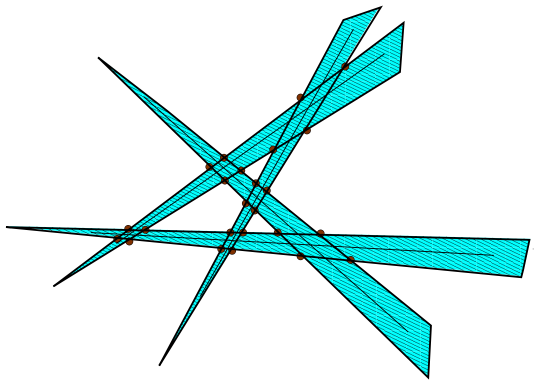

Figure 4: n pairwise crossing triangles with $\Omega(n^2)$ intersections on the boundary of their union.

As Whitesides and Zhao [45] discovered, if we exclude certain types of crossings between the members of $\mathcal{C}$, it is possible to give a linear upper bound on the complexity of $\cup\mathcal{C}$, even if two members of $\mathcal{C}$ may intersect in more than two points. A family $\mathcal{C}$ of simply connected regions bounded by simple closed curves in general position in the plane is called *k-admissible*, if any two members $C_1, C_2 \in \mathcal{C}$ have at most k boundary points in common and $C_1 \setminus C_2$ is connected (see Fig. 5). Clearly, we can restrict our attention to the case when k is even, because the members of $\mathcal{C}$ are in general position, i.e., their boundary curves cannot touch each other, so any two of them intersect in an even number of points.

Theorem 4.1. ([45]) *Let* $\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$ *be a* k*-admissible family of* $n \geq 3$ *simply connected regions in general position in the plane.*

Then the boundary of $\cup_{i=1}^{n} C_i$ *consists of at most* $k(3n-6)$ *elementary arcs, and this bound cannot be improved.*

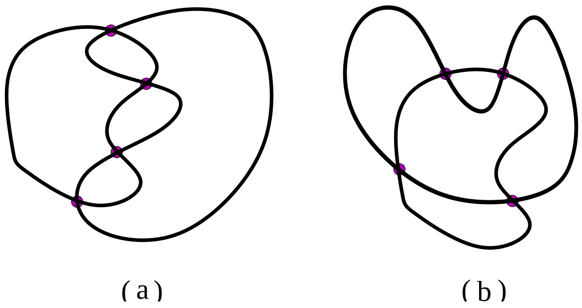

Figure 5: A pair of regions belonging to some (a) 4-admissible, (b) non-admissible family.

In case $k = 2$ Theorem 4.1 reduces to Theorem 2.2.

It was pointed out in [35] that Theorem 4.1 can be easily deduced from the following remarkable result of Chojnacki (alias Hanani) [9] (see also [44, 28]).

Lemma 4.2. ([9]) *Suppose that 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.*

This result can be regarded as a far-reaching generalization of the elementary fact that if two points of the plane can be connected by an arc crossing a fixed simple closed curve γ an even number of times, then they can also be connected by an arc which does not cross γ at all.

To see that Lemma 4.2 implies Theorem 4.1, it is enough to bound the number of points on $\mathrm{Bd}(\cup\mathcal{C})$, which belong to the boundary of more than one member of $\mathcal{C}$. For every C_i which contributes at least one arc to $\mathrm{Bd}(\cup\mathcal{C})$, fix a point p_i in the interior of such an arc. For any pair $C_i, C_j \in \mathcal{C}$ which has an intersection point $q \in \mathrm{Bd}(\cup\mathcal{C})$, draw an edge (but only one!) between p_i and p_j, as follows. Starting from p_i, follow $\mathrm{Bd}(C_i)$ to q in clockwise direction, and from there follow $\mathrm{Bd}(C_j)$ to p_j in counter-clockwise direction. It is not hard to verify that any two edges of this graph not incident to the same vertex cross an even number of times. Thus, the graph has at most $3n-6$ edges. That is, there are at most $3n-6$ pairs $\{C_i, C_j\}$ contributing intersection points to $\mathrm{Bd}(\cup\mathcal{C})$, and each of them can contribute at most k points.

5 Counting special intersections

Theorem 2.2 can be regarded as an upper bound on the number all intersection points between the curves $\mathrm{Bd}(C_i)$, which lie on the boundary of $\cup_{i=1}^{n} C_i$. Another way to generalize this result is to drop the condition that any pair of boundary curves intersect in at most two points, but to count only those intersections on $\mathrm{Bd}(\cup_{i=1}^{n} C_i)$, which belong to such pairs.

If two members of $\mathcal{C}$ have precisely two boundary points in common, then these points are called *regular intersection points*. All other intersection points between boundary curves are called *irregular*.

Theorem 5.1 ([35]) *Given a family $\mathcal{C}$ of $n \geq 3$ convex regions in general position in the plane, let R and I denote the number of regular and irregular intersection points of their boundaries, resp., which belong to* $\mathrm{Bd}(\cup\mathcal{C})$.

Then we have $R \leq 2I + 6n - 12$.

The last result is sharper than Theorem 2.2 in the sense that to obtain a $6n-12$ upper bound on the number of elementary arcs (the number of intersection points) on $\mathrm{Bd}(\cup\mathcal{C})$, we do not have to exclude all irregular intersections. It is sufficient to assume that no such intersection occurs on $\mathrm{Bd}(\cup\mathcal{C})$. However, for some technical reasons, we have been unable to establish Theorem 5.1 without the additional assumption that all members of $\mathcal{C}$ are *convex*. We conjecture that this assumption can be dropped.

It is not hard to show that the coefficient of I in Theorem 5.1 cannot be replaced by any constant smaller than 2.

If we want to get a non-trivial (i.e., subquadratic) upper bound on R, we have to limit the number of times two boundary curves are allowed to cross each other. But even under such an assumption we cannot expect a *linear* upper bound. There is a family of n disks and rectangles in general position in the plane satisfying $R = \Omega(n^{4/3})$. The best positive result in this direction is the following.

Theorem 5.2. ([3]) *Let $\mathcal{C}$ be a family of n simply connected regions in the plane. Suppose that they are bounded by simple closed curves in general position, any two of which intersect in at most s points, where s is a constant.*

Then there exists $\delta = \delta(s) > 0$ such that the number R of regular intersection points on $\mathrm{Bd}(\cup\mathcal{C})$ *satisfies*

(i) $R = O(n^{2-\delta})$;

(ii) $R = O(n^{1.5+\varepsilon})$ for any $\varepsilon > 0$, provided that every member of $\mathcal{C}$ is convex.

6 The union of fat triangles – Counting holes

The construction at the beginning of Section 4, showing that the union of n triangles may have *quadratic* complexity, uses extremely narrow triangles. It was proved in [30] that if we restrict how narrow the triangles can be, we can still establish a nearly linear upper bound on the complexity of their union. For any $\delta > 0$, a triangle is said to be *δ-fat* if each of its angles is at least δ.

Theorem 6.1. ([30]) *For any fixed $\delta > 0$, the boundary of the union of n δ-fat triangles in the plane consists of at most $O(n \log\log n)$ elementary arcs.*

By a slight modification of the construction of Wiernik and Sharir [46] cited after Theorem 3.2, one can easily give an example of n equilateral ($\pi/3$-fat) triangles, whose union has a slightly superlinear boundary complexity ($\Omega(n\alpha(n))$).

Given a family $\mathcal{C}$ of simply connected regions in the plane, a connected component of the complement of $\cup\mathcal{C}$ is called a *hole* determined by $\mathcal{C}$. The proof of Theorem 6.1 is based on the fact that every family of n δ-fat triangles in the plane determines at most a linear number of holes. The strongest known bound of this type is the following.

Theorem 6.2. ([36]) *Any family of n δ-fat triangles in the plane determines $O\left(\frac{n}{\delta}\log\frac{2}{\delta}\right)$ holes. This bound is tight up to the logarithmic factor.*

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 6.3. ([36]) *Let $\mathcal{C} = \{C_1, C_2, \ldots, C_n\}$ be a family of $n > 1$ triangles in the plane, and let α_i denote the smallest angle of C_i $(1 \le i \le n)$. Suppose $0 < \alpha_1 \le \alpha_2 \le \cdots \le \alpha_n$, and let $k \le n$ 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 $\mathcal{C}' = \{C_1', C_2', \ldots, C_n'\}$, where C_i' is isomorphic to C_i and $\mathcal{C}'$ determines $\Omega(nk)$ holes.

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 boundary of the union. The following result strengthens some earlier bounds in [14].

Theorem 6.4. ([36]) *Let $\mathcal{C}$ be a family of n wedges in the plane with angles $0 < \alpha_1 \le \alpha_2 \le \cdots \le \alpha_n < \pi$. Let $k \le n$ be the largest integer satisfying $\sum_{i=1}^{k} \alpha_i < \pi$.*

If $k \ge 2$, then the boundary complexity of $\cup\mathcal{C}$ is $O(nk \log k)$. Furthermore, there exists a family of n wedges with angles $\alpha_1, \alpha_2, \ldots, \alpha_n$, which determines $\Omega\left((\pi - \alpha_n)nk\right)$ holes.

The concept of δ-fatness, as well as Theorem 6.2, has been extended to arbitrary polygons by van Kreveld [26]. For other extensions and generalizations, see [38],[37],[15], [13], and [12].

7 Fat objects in space

Many of the theorems in previous sections have natural generalizations to higher dimensions. In this section, we mention only a few 3-dimensional results. Given a family $\mathcal{C}$ of 3-dimensional bodies, an *edge* of their union is defined as a maximal connected arc on $\mathrm{Bd}(\cup\mathcal{C})$, which belongs to two distinct members of $\mathcal{C}$. A point of $\mathrm{Bd}(\cup\mathcal{C})$, belonging to *three* distinct members of $\mathcal{C}$ is called a *vertex*. A maximal connected 2-dimensional piece of $\mathrm{Bd}(\cup\mathcal{C})$, which belongs to a single member of $\mathcal{C}$, is a *face*. The *complexity* of the boundary of $\cup\mathcal{C}$ is defined as the total number of vertices, edges, and faces of the boundary. These numbers are related to one another via Euler's Formula.

Let $\mathcal{C}$ be a collection of topological balls in $\mathbb{R}^3$ such that the intersection of any two of their surfaces is either empty or is a simple closed curve, and the intersection of any three surfaces consists of at most two points. Then $\mathcal{C}$ is called a family of *pseudo-balls*. Taking the intersection of the surface of each member of $\mathcal{C}$ with all the other members, and applying Theorem 2.2 to the resulting 2-dimensional arrangements, we obtain

Corollary 7.1. ([24]) *The complexity of the boundary of the union of n pseudo-balls in $\mathbb{R}^3$ is $O(n^2)$. This bound is asymptotically tight.*

Another generalization of Theorem 2.2 provides an upper bound on the complexity of the space of free placements of a convex polyhedral robot which is allowed to translate amidst polyhedral obstacles in 3-space.

Theorem 7.2. ([4]) *Let $\{C_1, C_2, \ldots, C_n\}$ be a family of pairwise disjoint convex polyhedral "obstacles" in $\mathbb{R}^3$ with a total of N faces, and let R be a convex polyhedral "robot," whose number of faces is a constant.*

Then the complexity of the union of the "expanded obstacles" $C_1 \oplus (-R)$, $C_2 \oplus (-R), \ldots, C_n \oplus (-R)$ is $O(nN \log n)$. In the worst case this bound cannot be improved, apart from the logarithmic factor.

For the case, when the robot is a ball, we have a similar result.

Theorem 7.3. ([2]) *Let $\{C_1, C_2, \ldots\}$ be a family of pairwise disjoint convex polyhedral "obstacles" in $\mathbb{R}^3$ with a total of N faces, and let R be a ball-shaped "robot."*

Then the complexity of the union of the "expanded obstacles" $C_1 \oplus (-R)$, $C_2 \oplus (-R), \ldots, C_n \oplus (-R)$ is $O(N^{2+\varepsilon})$ for every $\varepsilon > 0$.

In particular, the last result shows that the complexity of the union of n congruent infinite cylinders in $\mathbb{R}^3$ is only at most slightly superquadratic. No non-trivial (subcubic) upper bound is known for infinite cylinders with arbitrary radii.

Theorem 3.2 has the following analogue.

Theorem 7.4. ([34]) *The complexity of the upper envelope of n non-vertical triangles in $\mathbb{R}^3$ is $O(n^2\alpha(n))$. This bound is asymptotically tight.*

A somewhat weaker form of the last result is true in a much more general setting: the complexity of the upper envelope of n $(d-1)$-dimensional algebraic surface patches in $\mathbb{R}^d$, satisfying some natural conditions, is $O(n^{d-1+\varepsilon})$ for every $\varepsilon > 0$ (see [19, 42]).

It seems likely that the results of Section 6 also generalize to higher dimensions. The intersection of two (three) half-spaces in general position in 3-space is called a *dihedral (resp., trihedral) wedge.* A wedge (tetrahedron) is called *δ-fat* if its dihedral angle (resp., each of its solid angles) is at least δ.

We conjecture that the complexity of the union of any family of δ-fat tetrahedra in $\mathbb{R}^3$ is at most slightly superquadratic in n. To prove this conjecture for *congruent* tetrahedra, it would be sufficient to show that the same statement holds for δ-fat trihedral wedges. We can prove only a weaker result.

Theorem 7.5. ([33]) *The complexity of the boundary of the union of n δ-fat dihedral wedges in $\mathbb{R}^3$ is $O(n^{2+\varepsilon})$ for every $\varepsilon > 0$.*

We do not know the answer to the following simple question: is it true that the complexity of the union of n cubes in $\mathbb{R}^3$ is $o(n^3)$? In fact, we do not even know whether n cubes in $\mathbb{R}^3$ always determine at most $o(n^3)$ holes.

It is not hard to show that the complexity of the union of *axis-parallel* cubes is $O(n^2)$ and that this bound is asymptotically tight. In the case of *congruent* axis-parallel cubes, this bound can be improved to linear [8]. For congruent, but not necessarily axis-parallel cubes, we have the following recent result.

Theorem 7.6. ([33]) *The complexity of the boundary of the union of n congruent cubes in $\mathbb{R}^3$ is $O(n^{2+\varepsilon})$ for every $\varepsilon > 0$.*

8 Fat wedges – Extremal hypergraph theory

It was shown by Katona and Kovalev [21, 23] that for any family $\mathcal{C}$ of convex sets in $\mathbb{R}^d$, the number of *holes*, i.e., connected components of $\mathbb{R}^d \setminus \cup\mathcal{C}$, is at most $\sum_{i=0}^{d} \binom{n}{i}$, with equality only if $\mathcal{C}$ consists of hyperplanes or parallel strips in general position. In fact, if any d members of $\mathcal{C}$ have only a bounded number, s, of boundary points in common, then the *complexity* of $\cup\mathcal{C}$ is also $O(n^d)$, because each vertex p of the union is determined by d members, whose boundaries pass through p.

The aim of this section is to sketch a proof of the following weak (but non-trivial) version of Theorem 7.5.

Proposition 8.1. *There is an $\varepsilon > 0$ such that the complexity of the boundary of the union of n δ-fat wedges in $\mathbb{R}^d$ is $O(n^{3-\varepsilon})$.*

Let $K^{(3)}(m)$ denote a complete 3-uniform hypergraph with 3 disjoint m-element vertex classes, consisting of all triples containing exactly one element from each class. As in [34], our basic tool is Erdös's result from extremal hypergraph theory.

Lemma 8.2. ([16]) *Let H be a 3-uniform hypergraph on n vertices containing no subhypergraph isomorphic to $K^{(3)}(m)$. Then H has at most n^{3-1/m^2} triples.*

Three l-membered families of half-spaces, $\mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_3 \subset \mathbb{R}^3$ are said to *meet regularly* if their arrangement, restricted to the convex hull of the points

$$\{\mathrm{Bd}(h_1) \cap \mathrm{Bd}(h_2) \cap \mathrm{Bd}(h_3) \mid h_1 \in \mathcal{H}_1, h_2 \in \mathcal{H}_2, h_3 \in \mathcal{H}_3\},$$

is combinatorially isomorphic to the arrangement

$$\mathcal{H}_1^0 = \{\{(x,y,z) \in \mathbb{R}^3 : x \le i\} \mid 1 \le i \le l\},$$

$$\mathcal{H}_2^0 = \{\{(x,y,z) \in \mathbb{R}^3 : y \le j\} \mid 1 \le j \le l\},$$

$$\mathcal{H}_3^0 = \{\{(x,y,z) \in \mathbb{R}^3 : z \leq k\}, | \; 1 \leq k \leq l\},$$

restricted to the cube $[1,l]^3$. (See Fig. 6.)

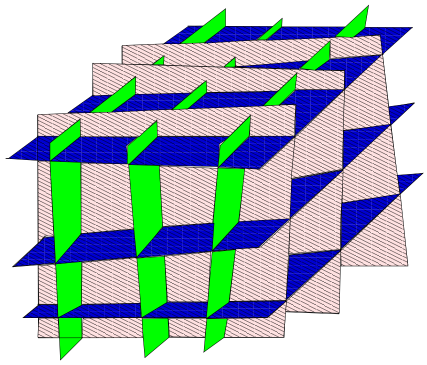

Figure 6: The boundary planes of three 3-membered families of regularly meeting half-spaces.

Lemma 8.3. ([34]) *For any l, there exists $L = L(l)$ such that any three L-membered families of half-spaces in $\mathbb{R}^3$ contain three l-membered subfamilies which meet regularly.*

Obviously, it is sufficient to prove Proposition 8.1 in the special case when the angle of every wedge is precisely δ, because every wedge can be obtained as the union of at most $\lceil \pi/\delta \rceil$ such wedges. Color the wedges with a constant number of colors so that the directions of the (infinite) edges and the directions of the corresponding faces of any two wedges of the same color differ from each other by less than 1 degree.

Assume, in order to obtain a contradiction, that there is a family $\mathcal{W}$ of n dihedral wedges whose union has at least $n^{3-\varepsilon}$ vertices on its boundary, for some $\varepsilon > 0$. Combine the last two lemmas. We obtain that if $\varepsilon = \varepsilon(l) > 0$ is small enough, there exist three pairwise disjoint l-membered monochromatic subfamilies $\mathcal{W}_1, \mathcal{W}_2, \mathcal{W}_3 \subset \mathcal{W}$ such that, expressing every $w_{st} \in \mathcal{W}_s$ as the intersection of two half-spaces, $h_{st} \cap h'_{st}$ $(1 \leq s \leq 3,\ 1 \leq t \leq l)$, the families of half-spaces

$$\mathcal{H}_1 := \{h_{1i} \mid 1 \leq i \leq l\},$$
$$\mathcal{H}_2 := \{h_{2j} \mid 1 \leq j \leq l\},$$
$$\mathcal{H}_3 := \{h_{3k} \mid 1 \leq k \leq l\}$$

meet regularly, and every point of the set

$$S := \{\mathrm{Bd}(h_{1i}) \cap \mathrm{Bd}(h_{2j}) \cap \mathrm{Bd}(h_{3k}) \mid (1 \leq i,j,k \leq l)\}$$

belongs to the boundary of $\cup\mathcal{W}$. Note that any two half-spaces belonging to the same family $\mathcal{H}_s$ are almost parallel. We can also assume without loss of generality that the isomorphism between $\mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_3$ and $\mathcal{H}_1^0, \mathcal{H}_2^0, \mathcal{H}_3^0$ takes each half-space h_{1i} (resp., h_{2j}, h_{3k}) to $\{(x,y,z) \in \mathbb{R}^3 : x \leq i\}$ (resp., $\{(x,y,z) \in \mathbb{R}^3 : y \leq j\}$, $\{(x,y,z) \in \mathbb{R}^3 : z \leq k\}$).

It follows from the fact that any two wedges belonging to the same $\mathcal{W}_s$ have the same color, that their edges are almost parallel. Therefore, we can choose a plane P such that the angle between P and the edge of every wedge $w \in \mathcal{W}_1 \cup \mathcal{W}_2 \cup \mathcal{W}_3$ is larger than, say, 30 degrees. This implies, for example, that $w \cap P$ is $\delta/2$-fat (in the plane). Translating P parallel to itself, if necessary, we can assume without loss of generality that it does not pass through any vertex of $\mathrm{Bd}(\cup\mathcal{W})$, and that P cuts the set S into two parts as equally as possible. Let R (and B) denote the set of elements of S on one side of P (on the other, resp.). Color any point $(i,j,k) \in [1,l]^3$ *red* or *blue* according to whether $\mathrm{Bd}(h_{1i}) \cap \mathrm{Bd}(h_{2j}) \cap \mathrm{Bd}(h_{3k})$ belongs to R or B, and denote the sets of red and blue points by R' and B', resp.

A set X of points with integer coordinates is called *convex in a given direction* if, for every segment xy parallel to this direction, both of whose endpoints belong to X, all other integer points of xy also belong to X. We obviously have

Claim 8.4. *R' and B' are convex in the directions of all three coordinate axes.*

Connect a point $(i,j,k) \in R'$ to $(i',j',k') \in B'$ by a *directed edge*, whenever their distance is 1, i.e., when they differ only in one of their coordinates, and in this coordinate their difference is 1. Using the fact that $|R'|, |B'| \geq \lfloor l^3/2 \rfloor$, it follows by standard isoperimetric inequalities that the number of directed edges is at least $l^2/2$. We can assume, by symmetry, that at least $l^2/12$ of them are parallel to the z-axis and are pointed upwards. Let D denote the orthogonal projection of these edges to the (x,y)-plane, i.e.,

$$D := \{(i,j) \in [1,l]^2 \mid \exists k \text{ such that } (i,j,k) \in R' \text{ and } (i,j,k+1) \in B'\}.$$

In view of Claim 8.4, we have

Claim 8.5. *The set $D \subseteq [1,l]^2$ is convex in the directions of both coordinate axes, and $|D| \geq l^2/12$.*

The last claim easily implies that D contains all integer points within an axis-parallel square, whose side length is at least $l/50$.

Switching back to the original picture, this means that there are two subfamilies $\mathcal{H}_1' \subset \mathcal{H}_1, \mathcal{H}_2' \subset \mathcal{H}_2$, each of size $l' := \lceil l/50 \rceil$, whose cross-sections on the plane P "meet regularly," i.e., are combinatorially isomorphic to the arrangement

$$\{\{(x,y) \in \mathbb{R}^2 : x \leq i\} \mid 1 \leq i \leq l'\},$$
$$\{\{(x,y) \in \mathbb{R}^2 : y \leq j\} \mid 1 \leq j \leq l'\},$$

restricted to $[1,l']^2$. Consider now the l'-membered families of wedges, $\mathcal{W}_1' \subset \mathcal{W}_1$ and $\mathcal{W}_2' \subset \mathcal{W}_2$ corresponding to the members of $\mathcal{H}_1'$ and $\mathcal{H}_2'$, resp.

It follows from the definitions that, for every $h_1 \in \mathcal{H}_1'$, $h_2 \in \mathcal{H}_2'$, the point $\mathrm{Bd}(h_1) \cap \mathrm{Bd}(h_2) \cap P$ belongs to the boundary of

$$P \cap ((\cup\mathcal{W}_1') \cup (\cup\mathcal{W}_2'))\,.$$

Consequently, the complexity of the union of all planar wedges $w \cap P$ ($w \in (\cup \mathcal{W}_1') \cup (\cup \mathcal{W}_2')$) is at least $(l')^2$.

Recall that, by the choice of the direction of P, the intersection of every element of $\mathcal{W}_1' \cup \mathcal{W}_2'$ with P is $(\delta/2)$-fat planar wedge. Thus, according to Theorem 6.1, the complexity of their union cannot exceed $O(l' \log l')$. This contradicts the conclusion of the last paragraph, provided that $l' = \lceil l/50 \rceil$ is sufficiently large (and $\varepsilon = \varepsilon(l) > 0$ is sufficiently small). This completes the proof of Proposition 8.1.

References

1. P. Agarwal, M. Katz, and M. Sharir: Computing depth order and related problems, *Comput. Geom. Theory Appls.* **5** (1995), 187–206.
2. P. Agarwal and M. Sharir: Pipes, cigars, and kreplach: The union of Minkowski sums in three dimensions, *Discrete Comput. Geom.* 24 (2000), 645–685.
3. B. Aronov, A. Efrat, D. Halperin, and M. Sharir: On the number of regular vertices of the union of Jordan regions, in: *Algorithm Theory, SWAT'98 (Stockholm), Lecture Notes in Comput. Sci.* **1432**, Springer-Verlag, Berlin, 1998, 322–334.
4. B. Aronov and M. Sharir: On translational motion planning of a convex polyhedron in 3-space, *SIAM J. Comput.* **26** (1997), 1785–1803.
5. M. Atallah: Some dynamic computational geometry problems, *Computers and Mathematics with Applications* **11** (1985), 1171–1181.
6. M. de Berg, M. Katz, F. van der Stappen, and J. Vleugels: Realistic input models for geometric algorithms, in: *Proc. 13th Annual Symposium on Computational Geometry,* ACM Press, 1997, 294–303.
7. J. L. Bentley and T. A. Ottmann: Algorithms for reporting and counting geometric intersections, *IEEE Trans. Comput.* **C-28** (1979), 643–647.
8. J.-D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec: Voronoi diagrams in higher dimensions under certain polyhedral distance functions, *Discrete Comput. Geom.* **19** (1998), 485–519.
9. Ch. Chojnacki (A. Hanani): Über wesentlich unplättbare Kurven im dreidimensionalen Raume, *Fund. Math.* **23** (1934), 135–142.
10. K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir, E. Welzl: Combinatorial complexity bounds for arrangements of curves and surfaces, *Discrete and Computational Geometry* **5** (1990), 99-160.
11. 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.
12. A. Efrat: The complexity of the union of (α, β)-covered objects, *Proceedings of the 15th Annual Symposium on Computational Geometry,* ACM Press, 1999, 134–142.
13. A. Efrat and M. J. Katz: On the union of κ-curved objects. *Comput. Geom. Theory Appl.* **14** (1999), 241–254.
14. 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.
15. A. Efrat and M. Sharir: On the complexity of the union of fat convex objects in the plane, *Discrete Comput. Geom.* **23** (2000), 171–189.

16. P. Erdös: On extremal problems of graphs and hypergraphs, *Israel J. Math.* **2** (1964), 183–190.
17. L. Guibas and M. Sharir: Combinatorics and algorithms of arrangements, in: *New Trends in Discrete and Computational Geometry* (J. Pach, ed.), Springer-Verlag, Berlin, 1993, 9–36.
18. P. Gupta, R. Janardan, and M. Smid: A technique for adding range restrictions to generalized searching problems, *Inform. Process. Lett.* **64** (1997), 263–269.
19. 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.
20. S. Hart and M. Sharir: Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes, *Combinatorica* **6** (1986), 151–177.
21. G. O. H. Katona: On a problem of L. Fejes Toth, *Studia Sci. Math. Hungar.* **12** (1977), 77–80.
22. 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.
23. M. D. Kovalev: A property of convex sets and its application (Russian), *Mat. Zametki* **44** (1988), 89–99. English translation: *Math. Notes* 44 (1988), 537–543.
24. 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.
25. K. Kedem and M. Sharir: An efficient motion-planning algorithm for a convex polygonal object in two-dimensional polygonal space, *Discrete Comput. Geom.* **5** (1990), 43–75.
26. M. van Kreveld: On fat partitioning, fat covering and the union size of polygons, *Computational Geometry: Theory and Applications* **9** (1998), 197–210.
27. D. Leven and M. Sharir: Planning a purely translational motion for a convex object in two-dimensional space using generalized Voronoi diagrams, *Discrete Comput. Geom.* **2** (1987), 9–31.
28. L. Lovász, J. Pach, and M. Szegedy: On Conway's thrackle conjecture, *Discrete Comput. Geom.* **18** (1997), 369–376.
29. T. Lozano-Pérez and M. A. Wesley: An algorithm for planning collision-free paths among polyhedral obstacles, *Commun. ACM* **22** (1979), 560–570.
30. J. Matoušek, J. Pach, M. Sharir, S. Sifrony, and E. Welzl: Fat triangles determine linearly many holes, *SIAM Journal of Computing* **23** (1994), 154–169.
31. P. McMullen: On the upper-bound conjecture for convex polytopes, *J. Combinatorial Theory, Ser. B* **10** (1971), 187–200.
32. J. Pach and P.K. Agarwal: *Combinatorial Geometry*, J. Wiley and Sons, New York, 1995.
33. J. Pach, I. Safruti, and M. Sharir, The union of congruent cubes in three dimensions, 17th ACM Symposium on Computational Geometry, 2001, accepted.
34. J. Pach and M. Sharir: The upper envelope of piecewise linear functions and the boundary of a region enclosed by convex plates: combinatorial analysis, *Discrete Comput. Geom.* **4** (1989), 291–309.
35. J. Pach and M. Sharir: On the boundary of the union of planar convex sets, *Discrete Comput. Geom.* **21** (1999), 321–328.
36. J. Pach and G. Tardos: On the boundary complexity of the union of fat triangles, *Proceedings of 41st Annual Symposium on Foundations of Computer Science, Los Angeles, 2000.*

37. F. van der Stappen: *Motion Planning amidst Fat Obstacles (Ph. D. Thesis,* Faculteit Wiskunde & Informatica, Universiteit Utrecht, 1994.
38. F. van der Stappen, D. Halperin, and M. Overmars: The complexity of the free space for a robot moving amidst fat obstacles, *Computational Geometry: Theory and Applications* **3** (1993), 353–373.
39. J. T. Schwartz and M. Sharir: On the "piano movers" problem I,II, *Comm. Pure Applied Math.* **36** (1983), 345–398 and *Adv. Applied Math.* **4** (1983), 298–351.
40. J. T. Schwartz and M. Sharir: A survey of motion planning and related geometric algorithms, in: *Geometric Reasoning* (D. Kapur and J. Mundy, eds.), MIT Press, Cambridge, MA, 1989, 157–169.
41. J. T. Schwartz and M. Sharir: Algorithmic motion planning in robotics, in: *Handbook of Theoretical Computer Science* (J. van Leeuwen, ed.), Elsevier, Amsterdam, 1990, 391–430.
42. M. Sharir: Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* **12** (1994), 327–345.
43. M. Sharir and P.K. Agarwal: *Davenport-Schinzel Sequences and Their Geometric Applications,* Cambridge University Press, Cambridge, 1995.
44. W. T. Tutte: Toward a theory of crossing numbers, *Journal of Combinatorial Theory* **8** (1970), 45–53.
45. 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.
46. A. Wiernik and M. Sharir: Planar realizations of nonlinear Davenport-Schinzel sequences by segments, *Discrete Comput. Geom.* **3** (1988), 15–47.