Chapter 4

Polygons

Although we can think of a line $\ell \subset \mathbb{R}^2$ as an infinite point set that consists of all points in $\mathbb{R}^2$ that are on ℓ, there still exists a finite description for ℓ. Such a description is, for instance, provided by the three coefficients $a, b, c \in \mathbb{R}$ of an equation of the form $ax + by = c$, with $(a, b) \neq (0, 0)$. Actually this holds true for all of the fundamental geometric objects that were mentioned in Chapter 1: Each of them has constant *description complexity* (or, informally, just *size*), that is, it can be described by a constant[1] number of parameters.

In this course we will typically deal with objects that are not of constant size. Often these are formed by merely aggregating constant-size objects, for instance, points to form a finite set of points. But sometimes we also demand additional structure that goes beyond aggregation only. Probably the most fundamental geometric objects of this type are what we call *polygons*. You probably learned this term in school, but what *is* a polygon precisely? Consider the examples shown in Figure 4.1. Are these all polygons? If not, where would you draw the line?


(a)   (b)   (c)   (d)   (e)   (f)

Figure 4.1: *What is a polygon?*

4.1 Classes of Polygons

Obviously, there is not *the* right answer to such a question and certainly there are different types of polygons. Often the term polygon is used somewhat sloppily in place

---

[1]Unless specified differently, we will always assume that the dimension is (a small) constant. In a high-dimensional space $\mathbb{R}^d$, one has to account for a description complexity of $\Theta(d)$.

of what we call a *simple polygon*, defined below.

Definition 4.1. *A **simple polygon** is a compact region $P \subset \mathbb{R}^2$ that is bounded by a simple closed curve $\gamma : [0, 1] \to \mathbb{R}^2$ that consists of a finite number of line segments. A **curve** is a continuous map $\gamma : [0, 1] \to \mathbb{R}^2$. A curve γ is **closed**, if $\gamma(0) = \gamma(1)$ and it is **simple** if it is injective on $[0, 1)$, that is, the curve does not intersect itself.*

Out of the examples shown above only Polygon 4.1a is simple. For each of the remaining polygons it is impossible to combine the bounding segments into a simple closed curve.

The term *compact* for subsets of $\mathbb{R}^d$ means bounded and closed. A subset of $P \subset \mathbb{R}^d$ is *bounded*, if it is contained in the ball of radius r around the origin, for some finite $r > 0$. Being closed means that the boundary is considered to be part of the polygon. In order to formally define these terms, let us briefly review a few basic notions from topology.

The standard topology of $\mathbb{R}^d$ is defined in terms of the Euclidean metric. A point $p \in \mathbb{R}^d$ is *interior* to a set $P \subseteq \mathbb{R}^d$, if there exists an ε-ball $B_\varepsilon(p) = \{x \in \mathbb{R}^d : \|x - p\| < \varepsilon\}$ around p, for some $\varepsilon > 0$, that is completely contained in P. A set is *open*, if all of its points are interior; and it is *closed*, if its complement is open.

Exercise 4.2. *Determine for each of the following sets whether they are open or closed in $\mathbb{R}^2$. a) $B_1(0)$ b) $\{(1, 0)\}$ c) $\mathbb{R}^2$ d) $\mathbb{R}^2 \backslash \mathbb{Z}^2$ e) $\mathbb{R}^2 \backslash \mathbb{Q}^2$ f) $\{(x, y) : x \in \mathbb{R}, y \geqslant 0\}$*

Exercise 4.3. *Show that the union of countably many open sets in $\mathbb{R}^d$ is open. Show that the union of a finite number of closed sets in $\mathbb{R}^d$ is closed. (These are two of the axioms that define a topology. So the statements are needed to assert that the metric topology is a topology, indeed.) What follows for intersections of open and closed sets? Finally, show that the union of countably many closed sets in $\mathbb{R}^d$ is not necessarily closed.*

The *boundary* ∂P of a set $P \subset \mathbb{R}^d$ consists of all points that are neither interior to P nor to its complement $\mathbb{R}^d \setminus P$. By definition, for every $p \in \partial P$ every ball $B_\varepsilon(p)$ contains both points from P and from $\mathbb{R}^d \setminus P$. Sometimes one wants to consider a set $P \subset \mathbb{R}^d$ open although it is not. In that case one can resort to the *interior* P° of P that is formed by the subset of points interior to P. Similarly, the *closure* $\overline{P}$ of P is defined by $\overline{P} = P \cup \partial P$.

Lower-dimensional objects, such as line segments in $\mathbb{R}^2$ or triangles in $\mathbb{R}^3$, do not possess any interior point (because the ε-balls needed around any such point are full-dimensional). Whenever we want to talk about the interior of a lower-dimensional set S, we use the qualifier *relative* and write $\mathrm{relint}(S)$ to denote the interior of S relative to the smallest affine subspace that contains S.

For instance, the smallest affine subspace that contains a line segment is a line and so the relative interior of a line segment in $\mathbb{R}^2$ consists of all points except the endpoints, just like for an interval in $\mathbb{R}^1$. Similarly, for a triangle in $\mathbb{R}^3$ the smallest affine subspace that contains it is a plane. Hence its relative interior is just the interior of the triangle, considered as a two-dimensional object.

Exercise 4.4. *Show that for any* $P \subset \mathbb{R}^d$ *the interior* P° *is open. (Why is there something to show to begin with?) Show that for any* $P \subset \mathbb{R}^d$ *the closure* $\overline{P}$ *is closed.*

When describing a simple polygon P it is sufficient to describe only its boundary ∂P. As ∂P by definition is a simple closed curve γ that consists of finitely many line segments, we can efficiently describe it as a sequence $p_1, \ldots, p_n$ of points, such that γ is formed by the line segments $\overline{p_1 p_2}, \overline{p_2 p_3}, \ldots, \overline{p_{n-1} p_n}, \overline{p_n p_1}$. These points are referred to as the *vertices* of the polygon, and the segments connecting them are referred as the *edges* of the polygon. The *set of vertices* of a polygon P is denoted by $V(P)$, and the *set of edges* of P is denoted by $E(P)$.

Knowing the boundary, it is easy to tell apart the (bounded) interior from the (unbounded) exterior. This is asserted even for much more general curves by Theorem 2.1 (Jordan curve theorem). To prove this theorem in its full generality is surprisingly difficult. For simple polygons the situation is easier, though. The essential idea can be worked out algorithmically, which we leave as an exercise.

Exercise 4.5. *Describe an algorithm to decide whether a point lies inside or outside of a simple polygon. More precisely, given a simple polygon* $P \subset \mathbb{R}^2$ *as a list of its vertices* $(v_1, v_2, \ldots, v_n)$ *in counterclockwise order and a query point* $q \in \mathbb{R}^2$, *decide whether* q *is inside* P, *on the boundary of* P, *or outside. The runtime of your algorithm should be* $O(n)$.

There are good reasons to ask for the boundary of a polygon to form a simple curve: For instance, in the example depicted in Figure 4.1b there are several regions for which it is completely unclear whether they should belong to the interior or to the exterior of the polygon. A similar problem arises for the interior regions in Figure 4.1f. But there are more general classes of polygons that some of the remaining examples fall into. We will discuss only one such class here. It comprises polygons like the one from Figure 4.1d.

Definition 4.6. *A region* $P \subset \mathbb{R}^2$ *is a* **simple polygon with holes** *if it can be described as* $P = F \setminus \bigcup_{H \in \mathcal{H}} H^\circ$, *where* $\mathcal{H}$ *is a finite collection of pairwise disjoint simple polygons (called holes) and* F *is a simple polygon for which* $F^\circ \supset \bigcup_{H \in \mathcal{H}} H$.

The way this definition heavily depends on the notion of simple polygons makes it straightforward to derive a similar trichotomy as the Jordan Curve Theorem provides for simple polygons, that is, every point in the plane is either inside, or on the boundary, or outside of P (exactly one of these three).

4.2 Polygon Triangulation

From a topological point of view, a simple polygon is nothing but a disk and so it is a very elementary object. But geometrically a simple polygon can be—as if mocking the label we attached to it—a pretty complicated shape, see Figure 4.2 for an example. While

there is an easy and compact one-dimensional representation in terms of the boundary, as a sequence of vertices/points, it is often desirable to work with a more structured representation of the whole two-dimensional shape.


Figure 4.2: *A simple (?) polygon.*

For instance, it is not straightforward to compute the area of a general simple polygon. In order to do so, one usually describes the polygon in terms of simpler geometric objects, for which computing the area is easy. Good candidates for such shapes are triangles, rectangles, and trapezoids. Indeed, it is not hard to show that every simple polygon admits a "nice" partition into triangles, which we call a triangulation.

Definition 4.7. *A **triangulation** of a simple polygon* P *with vertex set* $V(P)$ *is a collection* $\mathcal{T}$ *of triangles, such that*

(1) $P = \bigcup_{T \in \mathcal{T}} T$;

(2) $V(P) = \bigcup_{T \in \mathcal{T}} V(T)$; *and*

(3) for every distinct pair $T, U \in \mathcal{T}$, *the intersection* $T \cap U$ *is either a common vertex, or a common edge, or empty.*

Exercise 4.8. *Show that each condition in Definition 4.7 is necessary in the following sense: Give an example of a non-triangulation that would form a triangulation if the condition was omitted. Is the definition equivalent if (3) is replaced by* $T^\circ \cap U^\circ = \emptyset$, *for every distinct pair* $T, U \in \mathcal{T}$?

If we are given a triangulation of a simple polygon P it is easy to compute the area of P by simply summing up the area of all triangles from $\mathcal{T}$. Triangulations are an incredibly useful tool in planar geometry, and one reason for their importance is that every simple polygon admits one.

Theorem 4.9. *Every simple polygon has a triangulation.*

Proof. Let P be a simple polygon on n vertices. We prove the statement by induction on n. For $n = 3$ we face a triangle P that is a triangulation by itself. For $n > 3$ consider the lexicographically smallest vertex v of P, that is, among all vertices of P with a smallest x-coordinate the one with smallest y-coordinate. Denote the neighbors of v (next vertices) along ∂P by u and w. Consider the line segment $\overline{uw}$. We distinguish two cases.

Case 1: except for its endpoints u and w, the segment $\overline{uw}$ lies completely in P°. Then $\overline{uw}$ splits P into two smaller polygons, the triangle uvw and a simple polygon P' on $n-1$ vertices (Figure 4.3a). By the inductive hypothesis, P' has a triangulation that together with T yields a triangulation of P.


(a) Case 1.   **(b)** Case 2.

Figure 4.3: *Cases in the proof of Theorem 4.9.*

Case 2: $\mathrm{relint}(\overline{uw}) \not\subset P^\circ$ (Figure 4.3b). By choice of v, the polygon P is contained in the closed halfplane to the right of the vertical line through v. Therefore, as the segments $\overline{uv}$ and $\overline{vw}$ are part of a simple closed curve defining ∂P, every point sufficiently close to v and between the rays vu and vw must be in P°.

On the other hand, since $\mathrm{relint}(\overline{uw}) \not\subset P^\circ$, there is some point from ∂P in the interior of the triangle $T = uvw$ (by the choice of v the points u, v, w are not collinear and so T is a triangle, indeed) or on the line segment $\overline{uw}$. In particular, as ∂P is composed of line segments, there is a vertex of P in T° or on $\overline{uw}$ (otherwise, a line segment would have to intersect the line segment $\overline{uw}$ twice, which is impossible). Among all such vertices select p to be one that is furthest from the line uw. Then the open line segment $\mathrm{relint}(\overline{vp})$ is contained in T° and, thus, it splits P into two polygons P_1 and P_2 on less than n vertices each (in one of them, u does not appear as a vertex, whereas w does not appear as a vertex in the other). By the inductive hypothesis, both P_1 and P_2 have triangulations and their union yields a triangulation of P. □

Exercise 4.10. *In the proof of Theorem 4.9, would the argument in Case 2 also work if the point p was chosen to be a vertex of P in T° that is closest to v (in Euclidean distance)?*

The configuration from Case 1 above is called an *ear*: three consecutive vertices u, v, w of a simple polygon P such that the relative interior of $\overline{uw}$ lies in P°. In fact, we could have skipped the analysis for Case 2 by referring to the following theorem.

Theorem 4.11 (Meisters [13, 14]). *Every simple polygon that is not a triangle has two non-overlapping ears, that is, two ears A and B such that $A^\circ \cap B^\circ = \emptyset$.*

But knowing Theorem 4.9 we can obtain Theorem 4.11 as a direct consequence of the following

Theorem 4.12. *Every triangulation of a simple polygon on $n \geqslant 4$ vertices contains at least two (triangles that are) ears.*

Exercise 4.13. *Prove Theorem 4.12.*

Exercise 4.14. *Let P be a simple polygon with vertices $v_1, v_2, \ldots, v_n$ (in counterclockwise order), where v_i has coordinates (x_i, y_i). Show that the area of P is*

$$\frac{1}{2} \sum_{i=1}^{n} x_i y_{i+1} - x_{i+1} y_i,$$

where $(x_{n+1}, y_{n+1}) = (x_1, y_1)$.

The number of edges and triangles in a triangulation of a simple polygon are completely determined by the number of vertices, as the following simple lemma shows.

Lemma 4.15. *Every triangulation of a simple polygon on $n \geqslant 3$ vertices consists of $n - 2$ triangles and $2n - 3$ edges.*

Proof. Proof by induction on n. The statement is true for $n = 3$. For $n > 3$ consider a simple polygon P on n vertices and an arbitrary triangulation T of P. Any edge uv in T that is not an edge of P (and there must be such an edge because P is not a triangle) partitions P into two polygons P_1 and P_2 with n_1 and n_2 vertices, respectively. Since $n_1, n_2 < n$ we conclude by the inductive hypothesis that T partitions P_1 into $n_1 - 2$ triangles and P_2 into $n_2 - 2$ triangles, using $2n_1 - 3$ and $2n_2 - 3$ edges, respectively.

All vertices of P appear in exactly one of P_1 or P_2, except for u and v, which appear in both. Therefore $n_1 + n_2 = n + 2$ and so the number of triangles in T is $(n_1 - 2) + (n_2 - 2) = (n_1 + n_2) - 4 = n + 2 - 4 = n - 2$. Similarly, all edges of T appear in exactly one of P_1 or P_2, except for the edge uv, which appears in both. Therefore the number of edges in T is $(2n_1 - 3) + (2n_2 - 3) - 1 = 2(n_1 + n_2) - 7 = 2(n + 2) - 7 = 2n - 3$. $\square$

The universal presence of triangulations is something particular about the plane: The natural generalization of Theorem 4.9 to dimension three and higher does not hold. What is this generalization, anyway?

Tetrahedralizations in $\mathbb{R}^3$.   A simple polygon is a planar object that is a topological disk that is locally bounded by patches of lines. The corresponding term in $\mathbb{R}^3$ is a *polyhedron*, and although we will not formally define it here yet, a literal translation of the previous sentence yields an object that topologically is a ball and is locally bounded by patches of planes. A triangle in $\mathbb{R}^2$ corresponds to a tetrahedron in $\mathbb{R}^3$ and a *tetrahedralization* is a nice partition into tetrahedra, where "nice" means that the union of the tetrahedra covers the object, the vertices of the tetrahedra are vertices of the polyhedron, and any two distinct tetrahedra intersect in either a common triangular face, or a common edge, or a common vertex, or not at all.[2]

---

[2] These "nice" subdivisions can be defined in an abstract combinatorial setting, where they are called *simplicial complices*.

Unfortunately, there are polyhedra in $\mathbb{R}^3$ that do not admit a tetrahedralization. The following construction is due to Schönhardt [18]. It is based on a triangular prism, that is, two congruent triangles placed in parallel planes where the corresponding sides of both triangles are connected by a rectangle (Figure 4.4a). Then one triangle is twisted/rotated slightly within its plane. As a consequence, the rectangular faces are not plane anymore, but they obtain an inward dent along their diagonal in direction of the rotation (Figure 4.4b). The other (former) diagonals of the rectangular faces—labeled ab', bc', and

(a)   (b)

Figure 4.4: *The Schönhardt polyhedron cannot be subdivided into tetrahedra without adding new vertices.*

ca' in Figure 4.4b—are now epigonals, that is, they lie in the exterior of the polyhedron. Since these epigonals are the only edges between vertices that are not part of the polyhedron, there is no way to add edges to form a tetrahedron for a subdivision. Clearly the polyhedron is not a tetrahedron by itself, and so we conclude that it does not admit a subdivision into tetrahedra without adding new vertices. Actually, it is NP-complete to decide whether a non-convex polyhedron has a tetrahedralization [15]. If adding new vertices—so-called Steiner vertices—is allowed, then there is no problem to construct a tetrahedralization, and this holds true in general. Even if a tetrahedralization of a polyhedron exists, there is another significant difference to polygons in $\mathbb{R}^2$. While the number of triangles in a triangulation of a polygon depends only on the number of vertices, the number of tetrahedra in two different tetrahedralization of the same polyhedron may be different. See Figure 4.5 for a simple example of a polyhedron that has tetrahedralization with two or three tetrahedra. Deciding whether a convex polyhedron has a tetrahedralization with at most a given number of tetrahedra is NP-complete [6].

Exercise 4.16. *Characterize all possible tetrahedralizations of the three-dimensional cube.*

Algorithms.   Knowing that a triangulation exists is nice, but it is much better to know that it can also be constructed efficiently.

Figure 4.5: *Two tetrahedralizations of the same polyhedron, a triangular bipyramid.
The left partition uses two polyhedra; both the top vertex t and the bottom
vertex b belong to only one tetrahedron. The right partition uses three
polyhedra that all share the dashed diagonal bt.*

Exercise 4.17. *Convert Theorem 4.9 into an $O(n^2)$ time algorithm to construct a
triangulation for a given simple polygon on n vertices.*

The runtime achieved by the straightforward application of Theorem 4.9 is not op-
timal. We will revisit this question at several times during this course[3] and discuss
improved algorithms for the problem of triangulating a simple polygon.

The best (in terms of worst-case runtime) algorithm known due to Chazelle [7] com-
putes a triangulation in linear time. But this algorithm is very complicated and we will
not discuss it here. There is also a somewhat simpler randomized algorithm to compute
a triangulation in expected linear time [4], which we will not discuss in detail, either.
The question of whether there exists a simple (which is not really a well-defined term,
of course, except that Chazelle's Algorithm does not qualify) deterministic linear time
algorithm to triangulate a simple polygon remains open [10].

Polygons with holes. It is interesting to note that the complexity of the triangulation
problem changes to $\Theta(n \log n)$, if the polygon may contain holes [5]. This means that
there is an algorithm to construct a triangulation for a given simple polygon with holes
on a total of n vertices (counting both the vertices on the outer boundary and those of
holes) in $O(n \log n)$ time. But there is also a lower bound of $\Omega(n \log n)$ operations that
holds in all models of computation in which there exists a corresponding lower bound
for comparison-based sorting. This difference in complexity is a very common pattern:
There are many problems that are (sometimes much) harder for simple polygons with
holes than for simple polygons. So maybe the term "simple" has some justification, after
all. . .

---

[3]This is actually not true in this iteration of the course. But in the full version of the lecture notes you
can find the corresponding material in the appendix, in chapters A and ??.

General triangle covers. What if we drop the "niceness" conditions required for triangulations and just want to describe a given simple polygon as a union of triangles? It turns out this is a rather drastic change and, for instance, it is unlikely that we can efficiently find an optimal/minimal description of this type: Christ has shown [8] that it is NP-hard to decide whether for a simple polygon P on n vertices and a positive integer k, there exists a set of at most k triangles whose union is P. In fact, the problem is not even known to be in NP, because it is not clear whether the coordinates of solutions can always be encoded compactly.

4.3 The Art Gallery Problem

In 1973 Victor Klee posed the following question: "How many guards are necessary, and how many are sufficient to patrol the paintings and works of art in an art gallery with n walls?" From a geometric point of view, we may think of an "art gallery with n walls" as a simple polygon bounded by n edges, that is, a simple polygon P with n vertices. And a guard can be modeled as a point where we imagine the guard to stand and observe everything that is in sight. In sight, finally, refers to the walls of the gallery (edges of the polygon) that are opaque and, thus, prevent a guard to see what is behind. In other words, a guard (point) g can watch over every point $p \in P$, for which the line segment $\overline{gp}$ lies completely in $P°$, see Figure 4.6.


**Figure 4.6**: *The region that a guard g can observe.*

It is not hard to see that $\lfloor n/3 \rfloor$ guards are necessary in general.

Exercise 4.18. *Describe a family $(P_n)_{n \geqslant 3}$ of simple polygons such that P_n has n vertices and at least $\lfloor n/3 \rfloor$ guards are needed to guard it.*

What is more surprising: $\lfloor n/3 \rfloor$ guards are always sufficient as well. Chvátal [9] was the first to prove that, but then Fisk [11] gave a much simpler proof using—you may have guessed it—triangulations. Fisk's proof was considered so beautiful that it was included into "Proofs from THE BOOK" [3], a collection inspired by Paul Erdős' belief in "a place where God keeps aesthetically perfect proofs". The proof is based on the following lemma.

Lemma 4.19. *Every triangulation of a simple polygon is 3-colorable. That is, each vertex can be assigned one of three colors in such a way that adjacent vertices receive different colors.*

Proof. Induction on n. For $n = 3$ the statement is obvious. For $n > 3$, by Theorem 4.12 the triangulation contains an ear uvw. Cutting off the ear creates a triangulation of a polygon on $n - 1$ vertices, which by the inductive hypothesis admits a 3-coloring. Now whichever two colors the vertices u and w receive in this coloring, there remains a third color to be used for v. □


Figure 4.7: *A triangulation of a simple polygon on 17 vertices and a 3-coloring of it. The vertices shown solid orange form the smallest color class and guard the polygon using $\lfloor 17/3 \rfloor = 5$ guards.*

Theorem 4.20 (Fisk [11]). *Every simple polygon on n vertices can be guarded using at most $\lfloor n/3 \rfloor$ guards.*

Proof. Consider a triangulation of the polygon and a 3-coloring of the vertices as ensured by Lemma 4.19. Take the smallest color class, which clearly consists of at most $\lfloor n/3 \rfloor$ vertices, and put a guard at each vertex. As every point of the polygon is contained in at least one triangle and every triangle has exactly one vertex in the guarding set, the whole polygon is guarded. □

4.4 Optimal Guarding

While Exercise 4.18 shows that the bound in Theorem 4.20 is tight in general, it is easy to see that Fisk's method does not necessarily minimize the number of guards. Also, it is natural to lift the restriction that guards can be placed at vertices only, and allow guards to be placed anywhere on the boundary or even anywhere in the interior of the polygon. In all these variants, we can ask for the minimum number of guards required to guard a given polygon P. These problems have been shown to be NP-hard by Lee and Lin [12] already in the 1980s. However, if the guards are not constrained to lie on vertices, it is not clear whether the corresponding decision problem actually is in NP. In

Figure 4.8: *To guard this polygon with three guards, there must be one guard on each of the green dashed segments. The middle guard g_m must be to the left of the blue curve, to the right of the red curve, and on the dashed green line. The intersection point of these three curves has irrational coordinates.*

fact, recent results by Abrahamsen et al. suggest that this is unlikely to be the case. In the remainder of this section we will briefly discuss some of these results.

Recall that to show that a problem is in NP, one usually describes a *certificate* that allows to verify a solution for any problem instance in polynomial time. If we restrict the guards to be on vertices, a natural certificate for a solution is the set of vertices on which we place guards. In the general problem, a natural candidate for a certificate are the coordinates of the guards. Since no more than $\lfloor n/3 \rfloor$ guards are required, this seems a reasonable certificate. But what if the number of bits needed to explicitly represent these coordinates are exponential in n? One might be tempted to think that any reasonable guard can be placed at an intersection point of some lines that are defined by polygon vertices. Alas, in general this is not correct: some guards with irrational coordinates may be required, even if all vertices of P have integral coordinates. This surprising result has been presented in 2017 and we will sketch its main ideas, referring to the paper by Abrahamsen, Adamaszek, and Miltzow [1] for more details and the exact construction.

Consider the polygon shown in Figure 4.8, which consists of a main rectangular region with triangular, rectangular, and trapezoidal regions attached. We will argue that, if this polygon is guarded with less than four guards, at least one of the guards has an irrational coordinate. The polygon contains three pairs of triangular regions with the following structure. Each pair is connected by a green dashed segment in the figure. This segment contains one edge of each of the two triangles and separates their interiors. Hence, a single guard that sees both of these triangles has to be placed on this separating segment. Further, there is no other point that can guard two of these six triangles. Therefore, if we have only three guards, each of them must be placed on one of these three disjoint segments. The small rectangular regions to the left, top, and bottom outside the main rectangular region further constrain the positions of the guards along these segments.

Let the guards be g_ℓ, g_m, and g_r, as in the figure. The guard g_ℓ cannot see all the points inside the left two trapezoidal regions, and thus g_m has to be placed appropriately.

66

For each position of g_ℓ on its segment, we get a unique rightmost position on which a second guard can be placed to guard the two trapezoids. The union of these points defines an arc that is a segment of a quadratic curve (the roots of a quadratic polynomial). We get an analogous curve for g_r and the two trapezoids attached to the right. By a careful choice of the vertex coordinates, these two curves cross at a point p that also lies on the segment for the guard g_m and has irrational coordinates. It then follows from a detailed argument (see [1]) that p is the only feasible placement of g_m. Let us point out that the choice of the vertex coordinates to achieve this is far from trivial. For example, there can only be a single line defined by two points with rational coordinates that passes through p, and this is the line on which the guard g_m is constrained to lie on.

Exercise 4.21. *Let* P *be a polygon with vertices on the integer grid, and let* g *be a point inside that polygon with at least one irrational coordinate. Show that there can be at most one diagonal of* P *passing through* g*.*

Nevertheless, the sketched construction leads to the following result.

Theorem 4.22 (Abrahamsen et al. [1]). *For any* k*, there is a simple polygon* P *with integer vertex coordinates such that* P *can be guarded by* 3k *guards, while a guard set having only rational coordinates requires* 4k *guards.*

Abrahamsen, Adamaszek, and Miltzow [2] showed recently that the art gallery problem is actually complete in a complexity class called $\exists\mathbb{R}$. The *existential theory of the reals* (see [16, 17] for details) is the set of true sentences of the form $\exists x_1, \ldots, x_n \in \mathbb{R}: \phi(x_1, \ldots, x_n)$ for a quantifier-free Boolean formula ϕ without negation[4] that can use the constants 0 and 1, as well as the operators $+$, $*$, and $<$. For example, $\exists x, y: (x < y) \wedge (x * y < 1 + 1)$ is such a formula. A problem is in the complexity class $\exists\mathbb{R}$ if it allows for a polynomial-time reduction to the problem of deciding such formulas, and it is complete if in addition every problem in $\exists\mathbb{R}$ can be reduced to it in polynomial time.

Exercise 4.23. *Show that* NP $\subseteq \exists\mathbb{R}$*.*

For the art gallery problem, the result by Abrahamsen et al. [2] implies that the coordinates of an optimal guard set may be doubly-exponential in the input size. This statement does not exclude the possibility of a more concise, implicit way to express the existence of an optimal solution. However, if we found such a way, then this would imply that the art gallery problem is in NP, which, in turn, would imply NP $= \exists\mathbb{R}$.

Questions

11. *What is a simple polygon/a simple polygon with holes?* Explain the definitions and provide some examples of members and non-members of the respective classes.

---

[4]If we also allowed negation (and hence also the relations $\leqslant$ and $=$), then it would be possible to express statements like $\exists x: x * x = 1 + 1$ that have only irrational solutions, which is impossible to achieve by using only strict inequalities $<$. Interestingly, however, the complexity class $\exists\mathbb{R}$, when defined in terms of these more expressive formulas, would remain the same, see [17].

For a given polygon you should be able to tell which of these classes it belongs to or does not belong to and argue why this is the case.

12. *What is a closed/open/bounded set in $\mathbb{R}^d$? What is the interior/closure of a point set?* Explain the definitions and provide some illustrative examples. For a given set you should be able to argue which of the properties mentioned it possesses.

13. *What is a triangulation of a simple polygon? Does it always exist?* Explain the definition and provide some illustrative examples. Present the proof of Theorem 4.9 in detail.

14. *How about higher dimensional generalizations? Can every polyhedron in $\mathbb{R}^3$ be nicely subdivided into tetrahedra?* Explain Schönhardt's construction.

15. *How many points are needed to guard a simple polygon?* Present the proofs of Theorem 4.12, Lemma 4.19, and Theorem 4.20 in detail.

16. *Is there a compact representation for optimal guard placements?* State Theorem 4.22 and sketch the construction.

References

[1] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow, Irrational guards are sometimes needed. In *Proc. 33rd Internat. Sympos. Comput. Geom.*, pp. 3:1–3:15, 2017.

[2] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow, The art gallery problem is $\exists\mathbb{R}$-complete. In *Proc. 50th Annu. ACM Sympos. Theory Comput.*, pp. 65–73, 2018.

[3] Martin Aigner and Günter M. Ziegler, *Proofs from THE BOOK*, Springer-Verlag, Berlin, 3rd edn., 2003.

[4] Nancy M. Amato, Michael T. Goodrich, and Edgar A. Ramos, A randomized algorithm for triangulating a simple polygon in linear time. *Discrete Comput. Geom.*, 26/2, (2001), 245–265.

[5] Takao Asano, Tetsuo Asano, and Ron Y. Pinter, Polygon triangulation: efficiency and minimality. *J. Algorithms*, 7/2, (1986), 221–231.

[6] Alexander Below, Jesús A. De Loera, and Jürgen Richter-Gebert, The complexity of finding small triangulations of convex 3-polytopes. *J. Algorithms*, 50/2, (2004), 134–167.

[7] Bernard Chazelle, Triangulating a simple polygon in linear time. *Discrete Comput. Geom.*, 6/5, (1991), 485–524.

[8] Tobias Christ, Beyond triangulation: covering polygons with triangles. In *Proc. 12th Algorithms and Data Struct. Sympos.*, vol. 6844 of *Lecture Notes Comput. Sci.*, pp. 231–242, Springer-Verlag, 2011.

[9] Václav Chvátal, A combinatorial theorem in plane geometry. *J. Combin. Theory Ser. B*, 18/1, (1975), 39–41.

[10] Erik D. Demaine, Joseph S. B. Mitchell, and Joseph O'Rourke, The Open Problems Project, Problem #10. http://cs.smith.edu/~orourke/TOPP/P10.html.

[11] Steve Fisk, A short proof of Chvátal's watchman theorem. *J. Combin. Theory Ser. B*, 24/3, (1978), 374.

[12] Der-Tsai Lee and Arthur K. Lin, Computational complexity of art gallery problems. *IEEE Trans. Inform. Theory*, 32/2, (1986), 276–282.

[13] Gary H. Meisters, Polygons have ears. *Amer. Math. Monthly*, 82/6, (1975), 648–651.

[14] Gary H. Meisters, Principal vertices, exposed points, and ears. *Amer. Math. Monthly*, 87/4, (1980), 284–285.

[15] Jim Ruppert and Raimund Seidel, On the difficulty of triangulating three-dimensional nonconvex polyhedra. *Discrete Comput. Geom.*, 7, (1992), 227–253.

[16] Marcus Schaefer, Complexity of some geometric and topological problems. In *Proc. 17th Int. Sympos. Graph Drawing (GD 2009)*, pp. 334–344, 2009.

[17] Marcus Schaefer and Daniel Štefankovič, Fixed Points, Nash Equilibria, and the Existential Theory of the Reals. *Theory Comput. Syst.*, 60, (2017), 172–193.

[18] Erich Schönhardt, Über die Zerlegung von Dreieckspolyedern in Tetraeder. *Math. Ann.*, 98, (1928), 309–312.