

Chapter 1

Fundamentals


1.1 Models of Computation

Real RAM Model. A memory cell stores a real number. Any single arithmetic operation or comparison can be computed in constant time. In addition, sometimes also roots, logarithms, other analytic functions, indirect addressing (integral), or floor and ceiling are used.

This is a quite powerful (and somewhat unrealistic) model of computation, as a single real number in principle can encode an arbitrary amount of information. Therefore we have to ensure that we do not abuse the power of this model.

Algebraic Computation Trees (Ben-Or [1]). A computation is regarded as a binary tree.

- The leaves contain the (possible) results of the computation.
- Every node v with one child has an operation of the form $+$, $-$, $*$, $/$, $\sqrt{}$, \dots associated to it. The operands of this operation are constant input values, or among the ancestors of v in the tree.
- Every node v with two children has associated to it a branching of the form > 0 , ≥ 0 , or $= 0$. The branch is with respect to the result of v 's parent node. If the expression yields true, the computation continues with the left child of v ; otherwise, it continues with the right child of v .



If every branch is based on a linear function in the input values, we face a *linear computation tree*. Analogously one can define, say, quadratic computation trees. The term *decision tree* is used if all of the results are either true or false.

The complexity of a computation or decision tree is the maximum number of vertices along any root-to-leaf path. It is well known that $\Omega(n \log n)$ comparisons are required to sort n numbers. But also for some problems that appear easier than sorting at first glance, the same lower bound holds. Consider, for instance, the following problem.

Element Uniqueness

Input: $\{x_1, \dots, x_n\} \subset \mathbb{R}$, $n \in \mathbb{N}$.


Output: Is $x_i = x_j$, for some $i, j \in \{1, \dots, n\}$ with $i \neq j$?

Ben-Or [1] has shown that any algebraic decision tree to solve Element Uniqueness for n elements has complexity $\Omega(n \log n)$.


1.2 Basic Geometric Objects

We will mostly be concerned with the d -dimensional Euclidean space \mathbb{R}^d , for small $d \in \mathbb{N}$; typically, $d = 2$ or $d = 3$. The basic objects of interest in \mathbb{R}^d are the following.

Points. A point p , typically described by its d Cartesian coordinates $p = (x_1, \dots, x_d)$.




Directions. A vector $v \in \mathcal{S}^{d-1}$ (the $(d-1)$ -dimensional unit sphere), typically described by its d Cartesian coordinates $v = (x_1, \dots, x_d)$, with $\|v\| = \sqrt{\sum_{i=1}^d x_i^2} = 1$.



Lines. A line is a one-dimensional affine subspace. It can be described by two distinct points p and q as the set of all points r that satisfy $r = p + \lambda(q - p)$, for some $\lambda \in \mathbb{R}$.


While any pair of distinct points defines a unique line, a line in \mathbb{R}^2 contains infinitely many points and so it may happen that a collection of three or more points lie on a line. Such a collection of points is termed *collinear*¹.

Rays. If we remove a single point from a line and take the closure of one of the connected components, then we obtain a ray. It can be described by two distinct points p and q as the set of all points r that satisfy $r = p + \lambda(q - p)$, for some $\lambda \geq 0$. The *orientation* of a ray is the direction $(q - p)/\|q - p\|$.



¹Not *colinear*, which refers to a notion in the theory of coalgebras.

Line segment. A line segment is the intersection of two collinear rays of opposite orientation. It can be described by two points p and q as the set of all points r that satisfy $r = p + \lambda(q - p)$, for some $\lambda \in [0, 1]$. We will denote the line segment through p and q by \overline{pq} . Depending on the context we may allow or disallow *degenerate* line segments consisting of a single point only ($p = q$ in the above equation).



Hyperplanes. A hyperplane \mathcal{H} is a $(d-1)$ -dimensional affine subspace. It can be described algebraically by $d+1$ coefficients $\lambda_1, \dots, \lambda_{d+1} \in \mathbb{R}$, where $\|(\lambda_1, \dots, \lambda_{d+1})\| = 1$, as the set of all points (x_1, \dots, x_d) that satisfy the linear equation $\mathcal{H} : \sum_{i=1}^d \lambda_i x_i = \lambda_{d+1}$.

If the above equation is converted into an inequality, we obtain the algebraic description of a *halfspace* (in \mathbb{R}^2 : halfplane).

Spheres and balls. A sphere is the set of all points that are equidistant to a fixed point. It can be described by a point c (center) and a number $\rho \in \mathbb{R}$ (radius) as the set of all points p that satisfy $\|p - c\| = \rho$. The *ball* of radius ρ around p consists of all points p that satisfy $\|p - c\| \leq \rho$.

References

- [1] Michael Ben-Or, Lower bounds for algebraic computation trees. In *Proc. 15th Annu. ACM Sympos. Theory Comput.*, pp. 80–86, 1983, URL <http://dx.doi.org/10.1145/800061.808735>.

Chapter 2

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 the previous section: Each of them has constant *description complexity* (or, informally, just *size*), that is, it can be described by a constant¹ 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 2.1. Are all of these polygons? If not, where would you draw the line?




Figure 2.1: *What is a polygon?*

2.1 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

¹Unless 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 2.1 A **simple polygon** is a compact region $P \subset \mathbb{R}^2$ that is bounded by a simple closed curve $\gamma : [0, 1] \rightarrow \mathbb{R}^2$ that consists of a finite number of line segments. A curve is a continuous map $\gamma : [0, 1] \rightarrow \mathbb{R}^2$. 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 2.1a 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 2.2 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) $\mathbb{R}^2 \setminus \mathbb{Z}^2$ e) $\mathbb{R}^2 \setminus \mathbb{Q}^2$ f) $\{(x, y) : x \in \mathbb{R}, y \geq 0\}$

Exercise 2.3 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* \bar{P} of P is defined by $\bar{P} = P \cup \partial P$.

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 object, we use the qualifier *relative* and consider it relative to the smallest affine subspace that contains the object.

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 2.4 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 \bar{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, \dots, p_n of points, such that γ is formed by the line segments $\overline{p_1 p_2}, \overline{p_2 p_3}, \dots, \overline{p_{n-1} p_n}, \overline{p_n p_1}$. These points are referred to as the *vertices* of the polygon, and the segments connecting them are referred as the *edges* of the polygon.

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 the well-known Jordan-Curve Theorem.

Theorem 2.5 (Jordan 1887) Any simple closed curve $\gamma : [0, 1] \rightarrow \mathbb{R}^2$ divides the plane into exactly two connected components whose common boundary is formed by γ .

In full generality, the proof of the deceptively obvious claim is surprisingly difficult. We will not prove it here, the interested reader can find a proof, for instance, in the book of Mohar and Thomassen [11]. There exist different generalizations of the theorem and there also has been some debate about to which degree the original proof of Jordan is actually correct. For simple polygons the situation is easier, though. The essential idea can be worked out algorithmically, which we leave as an exercise.

Exercise 2.6 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, \dots, 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 2.1b 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 2.1f. But there are more general classes of polygons that some of the remaining examples fall into. We will discuss two such classes here. The first comprises polygons like the one from Figure 2.1d.

Definition 2.7 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).

The second class describes polygons that are “almost-simple” in the sense that they are arbitrarily close to a simple polygon. In many algorithmic scenarios such polygons can be treated very similarly to simple polygons.

Definition 2.8 A weakly simple polygon is a bounded region $P \subset \mathbb{R}^2$ such that

1. P is connected, ∂P is connected, and \bar{P} is simply-connected;
2. ∂P consists of a finite number of line segments, no two of which intersect except at a common endpoint; and
3. there exists a $k \in \mathbb{N}$ such that for every $\varepsilon > 0$ there exists a simple polygon Q_ε on at most k vertices for which the symmetric difference $(Q_\varepsilon \cup P) \setminus (Q_\varepsilon \cap P)$ has area less than ε .

It remains to define the terms connected and simply-connected. A set $P \subseteq \mathbb{R}^d$ is *connected*² if for every pair $p, q \in P$ there is a curve within P that connects p and q . A set $P \subseteq \mathbb{R}^d$ is *simply-connected* if it is connected and if for every simple closed curve $\gamma : [0, 1] \rightarrow P$ the bounded region enclosed by γ (well-defined by Theorem 2.5) is completely contained in P . For instance, the simple polygon with one hole depicted in Figure 2.2a is not simply-connected, because the red curve around the hole encloses points that do not belong to the polygon. The polygon P shown in Figure 2.2b is not simple but weakly simple; the simple polygon shown in orange provides a pretty good approximation that can be made arbitrarily close (but soon would be indistinguishable from P).




Figure 2.2: Weakly simple or not?

Exercise 2.9 Which of the shapes depicted in Figure 2.1 are weakly simple polygons?

Exercise 2.10 Show that the class of weakly simple polygons would change if we demanded a weakly simple polygon to be closed. Would it change the definition if we asked for a weakly simple polygon to be open?

²In general, a topological space is connected if it cannot be described as a disjoint union of open subsets. The property defined here is called path-connected. But as for \mathbb{R}^d both notions are equivalent, we stick to the more intuitive and geometric one.

Exercise 2.11 *Show that in Definition 2.8 the condition that the closure is simply-connected is necessary in the following sense: If dropped then there exist two “weakly simple” polygons that have the same boundary but different interior.*

2.2 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 2.3 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 2.3: *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 2.12 *A triangulation of a simple polygon P is a collection \mathcal{T} of triangles, such that*

- (1) $P = \bigcup_{T \in \mathcal{T}} T$;
- (2) *the vertices of all triangles in \mathcal{T} are also vertices of P ; 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.*

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 2.13 *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 2.4a). By the inductive hypothesis, P' has a triangulation that together with T yields a triangulation of P .




Figure 2.4: Cases in the proof of Theorem 2.13.

Case 2: the relative interior of \overline{uw} does not lie completely in P° (Figure 2.4b). 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 $\overline{uw} \not\subset P^\circ$, 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). Let p denote a leftmost such vertex. Then the open line segment \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 . \square

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 2.14 (Meisters [9, 10]) *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 2.13 we can obtain Theorem 2.14 as a direct consequence of the following

Theorem 2.15 *Every triangulation of a simple polygon on $n \geq 4$ vertices contains at least two (triangles that are) ears.*

Exercise 2.16 *Prove Theorem 2.15.*

Exercise 2.17 *Let P be a simple polygon with vertices v_1, v_2, \dots, v_n (in counterclockwise order), where v_i has coordinates (x_i, y_i) . Show that the area of P is*

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

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 2.18 *Every triangulation of a simple polygon on $n \geq 3$ 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 2.13 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.³

Unfortunately, there are polyhedra in \mathbb{R}^3 that do not admit a tetrahedralization. The following construction is due to Schönhardt [12]. 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 2.5a). 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 2.5b). The other (former) diagonals of the rectangular faces—labeled ab' , bc' , and




Figure 2.5: *The Schönhardt polyhedron cannot be subdivided into tetrahedra without adding new vertices.*

ca' in Figure 2.5b—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. 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.

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

Exercise 2.19 *Convert Theorem 2.13 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 2.13 is not optimal. We will revisit this question at several times during this course and discuss improved algorithms for the problem of triangulating a simple polygon.

³These “nice” subdivisions can be defined in an abstract combinatorial setting, where they are called *simplicial complices*.

The best (in terms of worst-case runtime) algorithm known due to Chazelle [4] computes 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 [2], which we will not discuss in detail, either. Instead you will later see a much simpler algorithm with a pretty-close-to linear runtime bound. 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 [7].

Polygons with holes. It is interesting to note that the complexity of the problem changes to $\Theta(n \log n)$, if the polygon may contain holes [3]. 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 the 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. . .

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 [5] 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.

2.3 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° .




Figure 2.6: *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 2.20 Describe a family $(P_n)_{n \geq 3}$ 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 [6] was the first to prove that, but then Fisk [8] 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” [1], 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 2.21 *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 2.15 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 . \square

Theorem 2.22 (Fisk [8]) *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 2.21. 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. \square




Figure 2.7: 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.

Questions

1. *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. 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.
2. *What is an open/closed/bounded/connected/simply-connected 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.
3. *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 2.13 in detail.
4. *How about higher dimensional generalizations? Can every polyhedron in \mathbb{R}^3 be nicely subdivided into tetrahedra?* Explain Schönhardt's construction.
5. *How many points are needed to guard a simple polygon?* Present the proofs of Theorem 2.15, Lemma 2.21, and Theorem 2.22 in detail.

References

- [1] Martin Aigner and Günter M. Ziegler, *Proofs from THE BOOK*. Springer, Berlin, 3rd edn., 2003.
- [2] 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, URL <http://dx.doi.org/10.1007/s00454-001-0027-x>.
- [3] Takao Asano, Tetsuo Asano, and Ron Y. Pinter, Polygon Triangulation: Efficiency and Minimality. *J. Algorithms*, **7**, 2, (1986), 221–231, URL [http://dx.doi.org/10.1016/0196-6774\(86\)90005-2](http://dx.doi.org/10.1016/0196-6774(86)90005-2).
- [4] Bernard Chazelle, Triangulating a simple polygon in linear time. *Discrete Comput. Geom.*, **6**, 5, (1991), 485–524, URL <http://dx.doi.org/10.1007/BF02574703>.
- [5] 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, URL http://dx.doi.org/10.1007/978-3-642-22300-6_20.

- [6] Václav Chvátal, A Combinatorial Theorem in Plane Geometry. *J. Combin. Theory Ser. B*, **18**, 1, (1975), 39–41, URL [http://dx.doi.org/10.1016/0095-8956\(75\)90061-1](http://dx.doi.org/10.1016/0095-8956(75)90061-1).
- [7] 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>.
- [8] Steve Fisk, A short proof of Chvátal's watchman theorem. *J. Combin. Theory Ser. B*, **24**, 3, (1978), 374, URL [http://dx.doi.org/10.1016/0095-8956\(78\)90059-X](http://dx.doi.org/10.1016/0095-8956(78)90059-X).
- [9] Gary H. Meisters, Polygons have ears. *Amer. Math. Monthly*, **82**, 6, (1975), 648–651, URL <http://www.jstor.org/stable/2319703>.
- [10] Gary H. Meisters, Principal Vertices, Exposed Points, and Ears. *Amer. Math. Monthly*, **87**, 4, (1980), 284–285, URL <http://www.jstor.org/stable/2321563>.
- [11] Bojan Mohar and Carsten Thomassen, *Graphs on Surfaces*. Johns Hopkins University Press, Baltimore, 2001.
- [12] Erich Schönhardt, Über die Zerlegung von Dreieckspolyedern in Tetraeder. *Math. Ann.*, **98**, (1928), 309–312, URL <http://dx.doi.org/10.1007/BF01451597>.