POLYGONS 1

Polygons are to planar geometry as integers are to numerical mathematics: a discrete subset of the full universe of possibilities that lends itself to efficient computations. And triangulations are the prime factorizations of polygons, alas without the benefit of the "Fundamental Theorem of Arithmetic" guaranteeing unique factorization. This chapter introduces triangulations (Section 1.1) and their combinatorics (Section 1.2), and then applies these concepts to the alluring art gallery theorem (Section 1.3), a topic at the roots of computational geometry which remains an active area of research today. Here we encounter a surprising difference between 2D triangulations and 3D tetrahedralizations.

Triangulations are highly constrained decompositions of polygons. Dissections are less constrained partitions, and engender the fascinating question of which pairs of polygons can be dissected and reassembled into each other. This so-called "scissors congruence" (Section 1.4) again highlights the fundamental difference between 2D and 3D (Section 1.5), a theme throughout the book.

1.1 DIAGONALS AND TRIANGULATIONS

Computational geometry is fundamentally *discrete* as opposed to continuous. Computation with curves and smooth surfaces are generally considered part of another field, often called "geometric modeling." The emphasis on computation leads to a focus on representations of geometric objects that are simple and easily manipulated. Fundamental building blocks are the *point* and the line *segment*, the portion of a line between two points. From these are built more complex structures. Among the most important of these structures are 2D polygons and their 3D generalization, polyhedra.

A *polygon*[1] P is the closed region of the plane bounded by a finite collection of line segments forming a closed curve that does not intersect itself. The line segments are called *edges* and the points where adjacent edges meet are called *vertices*. In general, we insist that vertices be true corners at which there is a bend between the adjacent edges, but in some

[1] Often the term *simple polygon* is used, to indicate that it is "simply connected," a concept we explore in Chapter 5.

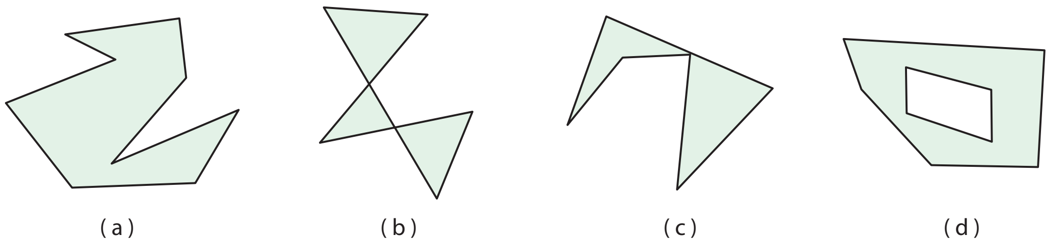

Figure 1.1. (a) A polygon. (b)–(d) Objects that are not polygons.

circumstances (such as in Chapter 2) it will be useful to recognize "flat vertices." The set of vertices and edges of P is called the *boundary* of the polygon, denoted as ∂P. Figure 1.1(a) shows a polygon with nine edges joined at nine vertices. Diagrams (b)–(d) show objects that fail to be polygons.

The fundamental "Jordan curve theorem," formulated and proved by Camille Jordan in 1882, is notorious for being both obvious and difficult to prove in its full generality. For polygons, however, the proof is easier, and we sketch the main idea.

Theorem 1.1 (Polygonal Jordan Curve). *The boundary ∂P of a polygon P partitions the plane into two parts. In particular, the two components of $\mathbb{R}^2 \setminus \partial P$ are the bounded interior and the unbounded exterior.*[2]

Sketch of Proof. Let P be a polygon in the plane. We first choose a fixed direction in the plane that is not parallel to any edge of P. This is always possible because P has a finite number of edges. Then any point x in the plane not on ∂P falls into one of two sets:

1. The ray through x in the fixed direction crosses ∂P an even number of times: x is exterior. Here a ray through a vertex is not counted as crossing ∂P.
2. The ray through x in the fixed direction crosses ∂P an odd number of times: x is interior.

Notice that all points on a line segment that do not intersect ∂P must lie in the same set. Thus the even sets and the odd sets are connected. And moreover, if there is a path between points in different sets, then this path must intersect ∂P. $\square$

This proof sketch is the basis for an algorithm for deciding whether a given point is inside a polygon, a low-level task that is encountered every time a user clicks inside some region in a computer game, and in many other applications.

[2] The symbol '$\setminus$' indicates set subtraction: $A \setminus B$ is the set of points in A but not in B.

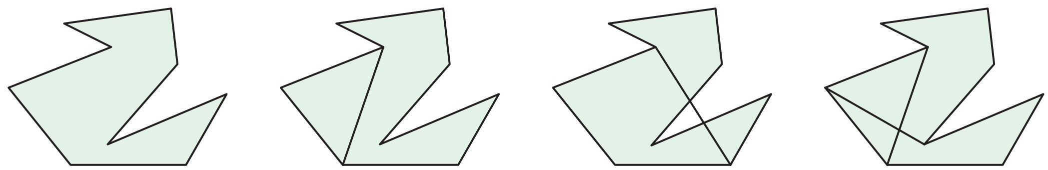

Figure 1.2. (a) A polygon with (b) a diagonal; (c) a line segment; (d) crossing diagonals.

Exercise 1.2. *Flesh out the proof of Theorem 1.1 by supplying arguments to (a) justify the claim that if there is a path between the even- and odd-crossings sets, the path must cross ∂P; and (b) establish that for two points in the same set, there is a path connecting them that does not cross ∂P.*

Algorithms often need to break polygons into pieces for processing. A natural decomposition of a polygon P into simpler pieces is achieved by drawing diagonals. A *diagonal* of a polygon is a line segment connecting two vertices of P and lying in the interior of P, not touching ∂P except at its endpoints. Two diagonals are *noncrossing* if they share no interior points. Figure 1.2 shows (a) a polygon, (b) a diagonal, (c) a line segment that is not a diagonal, and (d) two crossing diagonals.

Definition. A *triangulation* of a polygon P is a decomposition of P into triangles by a maximal set of noncrossing diagonals.

Here *maximal* means that no further diagonal may be added to the set without crossing (sharing an interior point with) one already in the set. Figure 1.3 shows a polygon with three different triangulations. Triangulations lead to several natural questions. How many different triangulations does a given polygon have? How many triangles are in each triangulation of a given polygon? Is it even true that every polygon always *has* a triangulation? Must every polygon have at least one diagonal? We start with the last question.

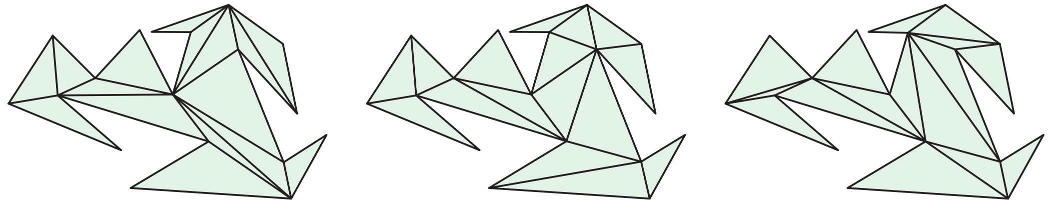

Figure 1.3. A polygon and three possible triangulations.

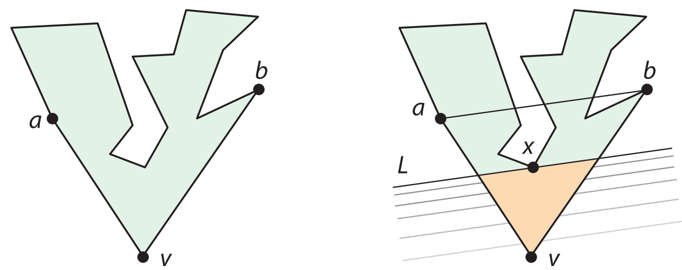

Figure 1.4. Finding a diagonal of a polygon through sweeping.

Lemma 1.3. *Every polygon with more than three vertices has a diagonal.*

Proof. Let v be the lowest vertex of P; if there are several, let v be the rightmost. Let a and b be the two neighboring vertices to v. If the segment ab lies in P and does not otherwise touch ∂P, it is a diagonal. Otherwise, since P has more than three vertices, the closed triangle formed by a, b, and v contains at least one vertex of P. Let L be a line parallel to segment ab passing through v. Sweep this line from v parallel to itself upward toward ab; see Figure 1.4. Let x be the first vertex in the closed triangle abv, different from a, b, or v, that L meets along this sweep. The (shaded) triangular region of the polygon below line L and above v is empty of vertices of P. Because vx cannot intersect ∂P except at v and x, we see that vx is a diagonal. □

Since we can decompose any polygon (with more than three vertices) into two smaller polygons using a diagonal, induction leads to the existence of a triangulation.

Theorem 1.4. *Every polygon has a triangulation.*

Proof. We prove this by induction on the number of vertices n of the polygon P. If $n = 3$, then P is a triangle and we are finished. Let $n > 3$ and assume the theorem is true for all polygons with fewer than n vertices. Using Lemma 1.3, find a diagonal cutting P into polygons P_1 and P_2. Because both P_1 and P_2 have fewer vertices than n, P_1 and P_2 can be triangulated by the induction hypothesis. By the Jordan curve theorem (Theorem 1.1), the interior of P_1 is in the exterior of P_2, and so no triangles of P_1 will overlap with those of P_2. A similar statement holds for the triangles of P_2. Thus P has a triangulation as well. □

Exercise 1.5. *Prove that every polygonal region with polygonal holes, such as Figure 1.1(d), admits a triangulation of its interior.*

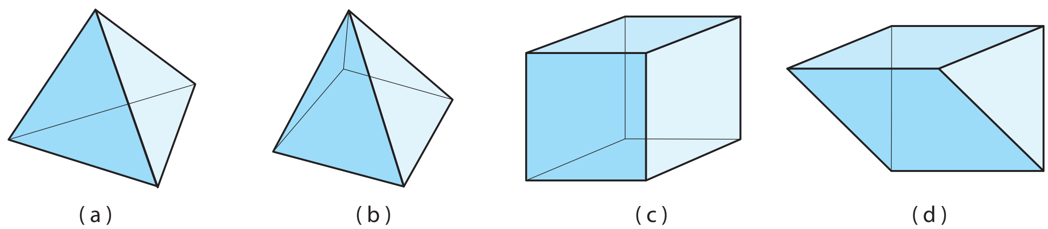

Figure 1.5. Polyhedra: (a) tetrahedron, (b) pyramid with square base, (c) cube, and (d) triangular prism.

That every polygon has a triangulation is a fundamental property that pervades discrete geometry and will be used over and over again in this book. It is remarkable that this notion does not generalize smoothly to three dimensions. A *polyhedron* is the 3D version of a polygon, a 3D solid bounded by finitely many polygons. Chapter 6 will define polyhedra more precisely and explore them more thoroughly. Here we rely on intuition. Figure 1.5 gives examples of polyhedra.

Just as the simplest polygon is the triangle, the simplest polyhedron is the *tetrahedron*: a pyramid with a triangular base. We can generalize the 2D notion of polygon triangulation to 3D: a *tetrahedralization* of a polyhedron is a partition of its interior into tetrahedra whose edges are diagonals of the polyhedron. Figure 1.6 shows examples of tetrahedralizations of the polyhedra just illustrated.

Exercise 1.6. *Find a tetrahedralization of the cube into five tetrahedra.*

We proved in Theorem 1.4 that all polygons can be triangulated. Does the analogous claim hold for polyhedra: can all polyhedra be

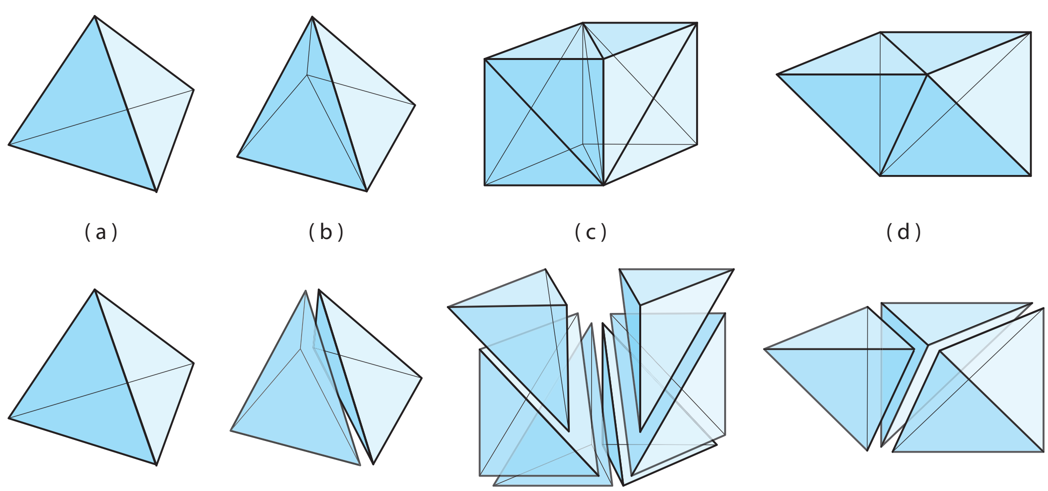

Figure 1.6. Tetrahedralizations of the polyhedra from Figure 1.5.

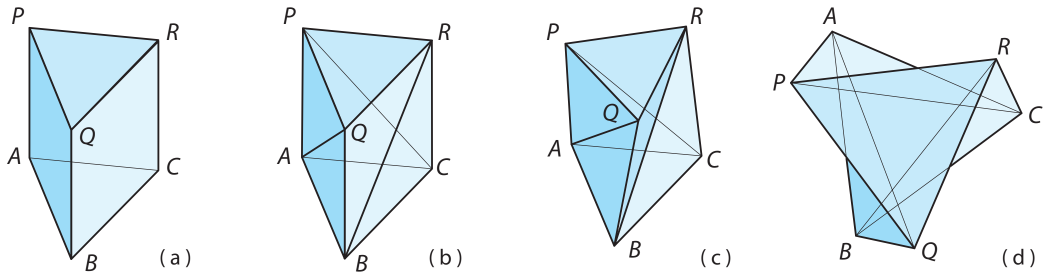

Figure 1.7. Construction of the Schönhardt polyhedron from a triangular prism, where (d) is the overhead view.

tetrahedralized? In 1911, Nels Lennes proved the surprising theorem that this is not so. We construct an example of a polyhedron, based on the 1928 model by Erich Schönhardt, which cannot be tetrahedralized. Let A, B, C be vertices of an equilateral triangle (labeled counterclockwise) in the xy-plane. Translating this triangle vertically along the z-axis reaching $z = 1$ traces out a triangular prism, as shown in Figure 1.7(a). Part (b) shows the prism with the faces partitioned by the diagonal edges AQ, BR, and CP. Now twist the top PQR triangle $\pi/6$ degrees in the $(z = 1)$-plane, rotating and stretching the diagonal edges. The result is the Schönhardt polyhedron, shown in (c) and in an overhead view in (d) of the figure. Schönhardt proved that this is the smallest example of an untetrahedralizable polyhedron.

Exercise 1.7. *Prove that the Schönhardt polyhedron cannot be tetrahedralized.*

> **UNSOLVED PROBLEM 1** — **Tetrahedralizable Polyhedra**
>
> Find characteristics that determine whether or not a polyhedron is tetrahedralizable. Even identifying a large natural class of tetrahedralizable polyhedra would be interesting.

This is indeed a difficult problem. It was proved by Jim Ruppert and Raimund Seidel in 1992 that it is NP-complete to determine whether a polyhedron is tetrahedralizable. *NP-complete* is a technical term from complexity theory that means, roughly, an intractable algorithmic problem. (See the Appendix for a more thorough explanation.) It suggests in this case that there is almost certainly no succinct characterization of tetrahedralizability.

1.2 BASIC COMBINATORICS

We know that every polygon has at least one triangulation. Next we show that the number of triangles in any triangulation of a fixed polygon is the same. The proof is essentially the same as that of Theorem 1.4, with more quantitative detail.

Theorem 1.8. *Every triangulation of a polygon P with n vertices has $n-2$ triangles and $n-3$ diagonals.*

Proof. We prove this by induction on n. When $n = 3$, the statement is trivially true. Let $n > 3$ and assume the statement is true for all polygons with fewer than n vertices. Choose a diagonal d joining vertices a and b, cutting P into polygons P_1 and P_2 having n_1 and n_2 vertices, respectively. Because a and b appear in both P_1 and P_2, we know $n_1 + n_2 = n + 2$. The induction hypothesis implies that there are $n_1 - 2$ and $n_2 - 2$ triangles in P_1 and P_2, respectively. Hence P has

$$(n_1 - 2) + (n_2 - 2) = (n_1 + n_2) - 4 = (n + 2) - 4 = n - 2$$

triangles. Similarly, P has $(n_1 - 3) + (n_2 - 3) + 1 = n - 3$ diagonals, with the $+1$ term counting d. □

Many proofs and algorithms that involve triangulations need a special triangle in the triangulation to initiate induction or start recursion. "Ears" often serve as special triangles. Three consecutive vertices a, b, c form an *ear* of a polygon if ac is a diagonal of the polygon. The vertex b is called the ear *tip*.

Corollary 1.9. *Every polygon with more than three vertices has at least two ears.*

Proof. Consider any triangulation of a polygon P with $n > 3$ vertices, which by Theorem 1.8 partitions P into $n-2$ triangles. Each triangle covers at most two edges of ∂P. Because there are n edges on the boundary of P but only $n-2$ triangles, by the pigeonhole principle at least two triangles must contain two edges of P. These are the ears. □

Exercise 1.10. *Prove Corollary 1.9 using induction.*

Exercise 1.11. *Show that the sum of the interior angles of any polygon with n vertices is $\pi(n-2)$.*

Exercise 1.12. *Using the previous exercise, show that the total turn angle around the boundary of a polygon is 2π. Here the* turn angle *at a vertex v is π minus the internal angle at v.*

Exercise 1.13. *Three consecutive vertices a, b, c form a* mouth *of a polygon if ac is an* external diagonal *of the polygon, a segment wholly outside. Formulate and prove a theorem about the existence of mouths.*

Exercise 1.14. *Let a polygon P with h holes have n total vertices (including hole vertices). Find a formula for the number of triangles in any triangulation of P.*

★ Exercise 1.15. *Let P be a polygon with vertices (x_i, y_i) in the plane. Prove that the area of P is*

$$\frac{1}{2}\left|\sum(x_i y_{i-1} - x_{i-1} y_i)\right|.$$

Although the number of triangles in any triangulation of a polygon is the same, it is natural to explore the number of different triangulations of a given polygon. For instance, Figure 1.3 shows a polygon with three different triangulations.

Exercise 1.16. *For each polygon in Figure 1.8, find the number of distinct triangulations.*

Exercise 1.17. *For each $n > 3$, find a polygon with n vertices that has a unique triangulation.*

The number of triangulations of a fixed polygon P has much to do with the "shape" of the polygon. One crucial measure of shape is the internal angles at the vertices. A vertex of P is called *reflex* if its angle is greater than π, and *convex* if its angle is less than or equal to π. Sometimes it is useful to distinguish a *flat* vertex, whose angle is exactly π, from a *strictly convex* vertex, whose angle is strictly less than π. A polygon P is a *convex* polygon if all vertices of P are convex. In general we exclude flat vertices, so unless otherwise indicated, the vertices of a convex polygon

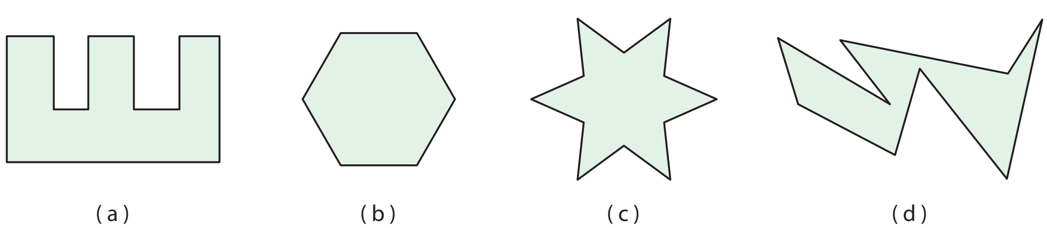

Figure 1.8. Find the number of distinct triangulations for each of the polygons given.

are strictly convex. With this understanding, a convex polygon has the following special property.

Lemma 1.18. *A diagonal exists between any two nonadjacent vertices of a polygon P if and only if P is a convex polygon.*

Proof. The proof is in two parts, both established by contradiction. First assume P is not convex. We need to find two vertices of P that do not form a diagonal. Because P is not convex, there exists a sequence of three vertices a, b, c, with b reflex. Then the segment ac lies (at least partially) exterior to P and so is not a diagonal.

Now assume P is convex but there are a pair of vertices a and b in P that do not form a diagonal. We identify a reflex vertex of P to establish the contradiction. Let σ be the shortest path connecting a to b entirely within P. It cannot be that σ is a straight segment contained inside P, for then ab is a diagonal. Instead, σ must be a chain of line segments. Each corner of this polygonal chain turns at a reflex vertex — if it turned at a convex vertex or at a point interior to P, it would not be the shortest. $\square$

For a convex polygon P, where every pair of nonadjacent vertices determines a diagonal, it is possible to count the number of triangulations of P based solely on the number of vertices. The result is the *Catalan number*, named after the nineteenth-century Belgian mathematician Eugène Catalan.

Theorem 1.19. *The number of triangulations of a convex polygon with $n+2$ vertices is the Catalan number*

$$C_n = \frac{1}{n+1}\binom{2n}{n}. \tag{1.1}$$

Proof. Let P_{n+2} be a convex polygon with vertices labeled from 1 to $n+2$ counterclockwise. Let $\mathcal{T}_{n+2}$ be the set of triangulations of P_{n+2}, where $\mathcal{T}_{n+2}$ has t_{n+2} elements. We wish to show that t_{n+2} is the Catalan number C_n.

Let ϕ be the map from $\mathcal{T}_{n+2}$ to $\mathcal{T}_{n+1}$ given by contracting the edge $\{1, n+2\}$ of P_{n+2}. To *contract* an edge ab is to shrink it to a point c so that c becomes incident to all the edges and diagonals that were incident to either a or b. Let T be an element of $\mathcal{T}_{n+1}$. What is important to note is the number of triangulations of $\mathcal{T}_{n+2}$ that map to T (i.e., the number of elements of $\phi^{-1}(T)$) equals the degree of vertex 1 in T. Figure 1.9 gives an example where (a) five triangulations of the octagon all map to (b) the same triangulation of the heptagon, where the vertex labeled 1 has degree five. This is evident since each edge incident to 1 can *open up* into a triangle in $\phi^{-1}(T)$, shown by the shaded triangles

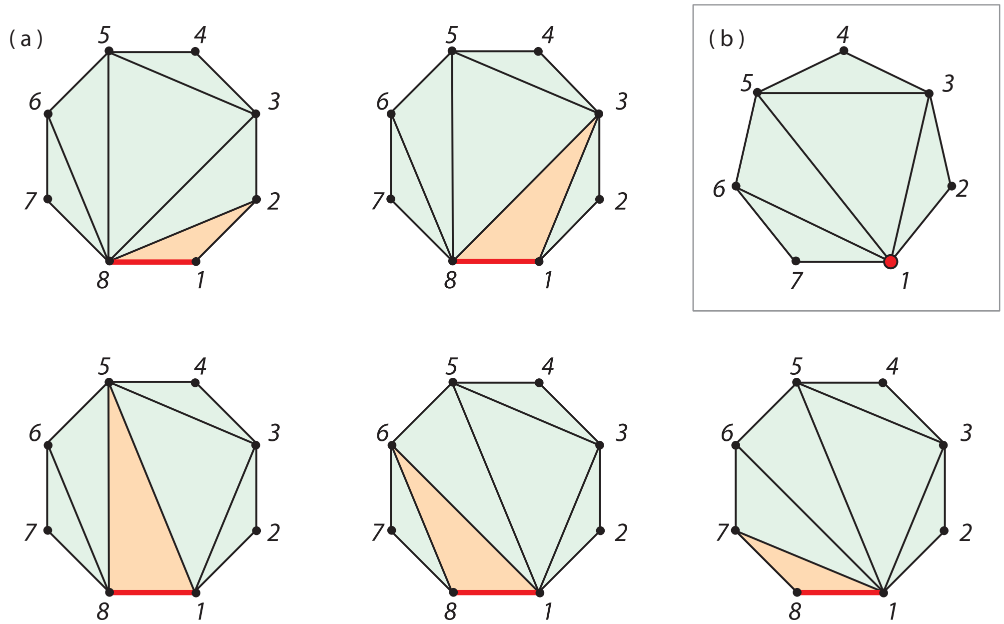

Figure 1.9. The five polygons in (a) all map to the same polygon in (b) under contraction of edge $\{1, 8\}$.

in (a). So we see that

$$t_{n+2} = \sum_{T \in \mathcal{T}_{n+1}} \text{degree of vertex 1 of } T.$$

Because this polygon is convex, this is true for all vertices of T. Therefore we can sum over all vertices of T, obtaining

$$\begin{aligned}(n+1) \cdot t_{n+2} &= \sum_{i=1}^{n+1} \sum_{T \in \mathcal{T}_{n+1}} \text{degree of vertex } i \text{ of } T \\ &= \sum_{T \in \mathcal{T}_{n+1}} \sum_{i=1}^{n+1} \text{degree of vertex } i \text{ of } T \\ &= 2\,(2n-1) \cdot t_{n+1}.\end{aligned}$$

The last equation follows because the sum of the degrees of all vertices of T double-counts the number of edges of T and the number of diagonals of T. Because T is in $\mathcal{T}_{n+1}$, it has $n+1$ edges, and by Theorem 1.8, it has $n-2$ diagonals. Solving for t_{n+2}, we get

$$\begin{aligned}t_{n+2} &= \frac{2(2n-1)}{n+1} \cdot t_{n+1} = 2^n \cdot \frac{2n-1}{n+1} \cdot \frac{2n-3}{n} \cdots \frac{3}{3} \cdot \frac{1}{2} \\ &= \frac{(2n)!}{(n+1)!\, n!} = \frac{1}{n+1}\binom{2n}{n}.\end{aligned}$$

This is the Catalan number C_n, completing the proof. □

For the octagon in Figure 1.9, the formula shows there are $C_6 = 132$ distinct triangulations. Is it possible to find a closed formula for the number of triangulations for nonconvex polygons P with n vertices? The answer, unfortunately, is NO, because small changes in the position of vertices can lead to vastly different triangulations of the polygon. What we do know is that convex polygons achieve the maximum number of triangulations.

Theorem 1.20. *Let P be a polygon with $n+2$ vertices. The number of triangulations of P is between 1 and C_n.*

Proof. Exercise 1.17 shows there are polygons with exactly one triangulation, demonstrating that the lower bound is realizable. For the upper bound, let P be any polygon with n labeled, ordered vertices, and let Q be a convex polygon also with n vertices, labeled similarly. Each diagonal of P corresponds to a diagonal of Q, and if two diagonals of P do not cross, neither do they cross in Q. So every triangulation of P (having $n-1$ diagonals by Theorem 1.8) determines a triangulation of Q (again with $n-1$ diagonals). Therefore P can have no more triangulations than Q, which by Theorem 1.19 is C_n. □

Thus we see that convex polygons yield the most triangulations. Because convex polygons have no reflex vertices (by definition), there might possibly be a relationship between the number of triangulations and the number of reflex vertices of a polygon. Sadly, this is not the case. Let P be a polygon with five vertices. By Theorem 1.19, if P has no reflex vertices, it must have 5 triangulations. Figure 1.10(a) shows P with one reflex vertex and only one triangulation, whereas parts (b) and (c) show P with two reflex vertices and two triangulations. So the number of triangulations does not necessarily decrease with the number of reflex vertices. In fact, the number of triangulations does not depend on the number of reflex vertices at all. Figure 1.10(d) shows a polygon with a unique triangulation with three reflex vertices. This example can be generalized to polygons with unique triangulations that contain arbitrarily many reflex vertices.

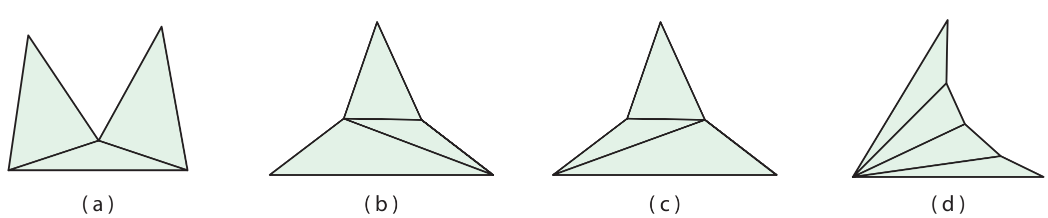

Figure 1.10. Triangulations of special polygons.

Exercise 1.21. *For each $n > 3$, find a polygon with n vertices with exactly two triangulations.*

Exercise 1.22. *For any $n \geq 3$, show there is no polygon with $n + 2$ vertices with exactly $C_n - 1$ triangulations.*

UNSOLVED PROBLEM 2 — Counting Triangulations

Identify features of polygons P that lead to a closed formula for the number of triangulations of P in terms of those features.

We learned earlier that properties can be lost in the move from 2D polygons to 3D polyhedra. For example, all polygons can be triangulated but not all polyhedra can be tetrahedralized. Moreover, by Theorem 1.8 above, we know that *every* polygon with n vertices must have the same number of triangles in *any* of its triangulation. For polyhedra, this is far from true. In fact, two different tetrahedralizations of the *same* polyhedron can result in a different number of tetrahedra! Consider Figure 1.11, which shows a polyhedron partitioned into two tetrahedra (a) and also into three (b).

(a)

(b)

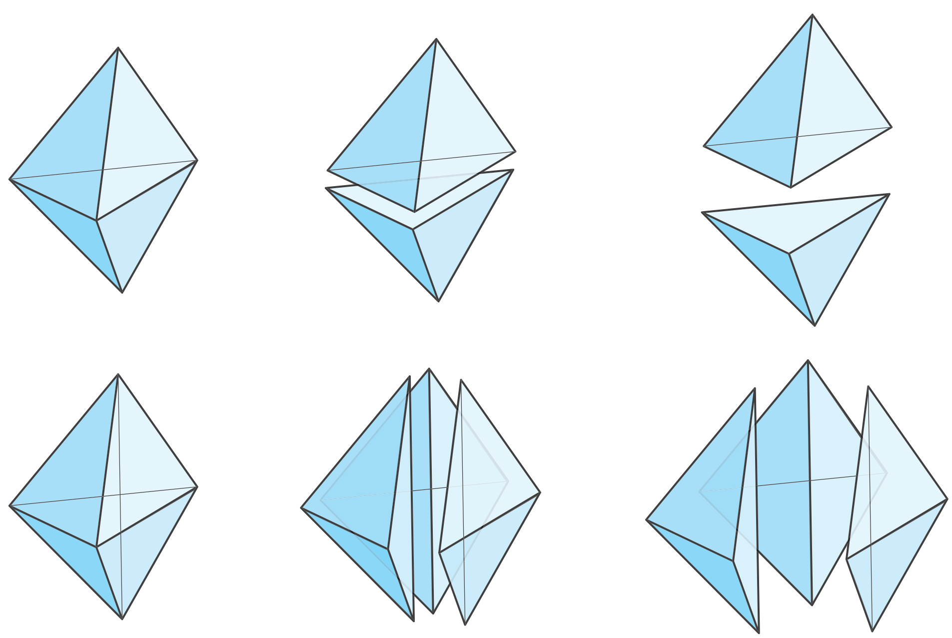

Figure 1.11. A polyhedron partitioned into (a) two and (b) three tetrahedra.

Even for a polyhedron as simple as the cube, the number of tetrahedra is not the same for all tetrahedralizations. It turns out that up to rotation and reflection, there are six different tetrahedralizations of the cube, one of which was shown earlier in Figure 1.6(c). Five of the six partition the cube into six tetrahedra, but one cuts it into only five tetrahedra.

Exercise 1.23. *Is it possible to partition a cube into six* congruent *tetrahedra? Defend your answer.*

Exercise 1.24. *Find the six different tetrahedralizations of the cube up to rotation and reflection.*

★ Exercise 1.25. *Classify the set of triangulations on the boundary of the cube that "induce" tetrahedralizations of the cube, where each such tetrahedralization matches the triangulation on the cube surface.*

As is common in geometry, concepts that apply to 2D and to 3D generalize to arbitrary dimensions. The n-dimensional generalization of the triangle/tetrahedron is the *n-simplex* of $n + 1$ vertices. Counting n-dimensional "triangulations" is largely unsolved:

UNSOLVED PROBLEM 3 — Simplices and Cubes

Find the smallest triangulation of the n-dimensional cube into n-simplices. It is known, for example, that the 4D cube (the *hypercube*) may be partitioned into 16 4-simplices, and this is minimal. But the minimum number is unknown except for the few small values of n that have yielded to exhaustive computer searches.

Exercise 1.26. *Show that the n-dimensional cube can be triangulated into exactly $n!$ simplices.*

1.3 THE ART GALLERY THEOREM

A beautiful problem posed by Victor Klee in 1973 engages several of the concepts we have discussed: Imagine an art gallery whose floor plan is modeled by a polygon. A guard of the gallery corresponds to a point on our polygonal floor plan. Guards can see in every direction, with a full 360° range of visibility. Klee asked to find the fewest number of (stationary) guards needed to protect the gallery. Before tackling this problem, we need to define what it means to "see something" mathematically.

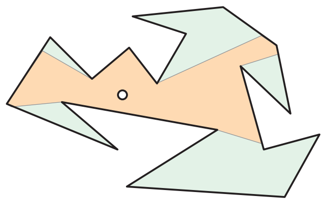

Figure 1.12. Examples of the range of visibility available to certain placement of guards.

A point x in polygon P is *visible* to point y in P if the line segment xy lies in P. This definition allows the line of sight to have a grazing contact with the boundary ∂P (unlike the definition for diagonal). A set of guards *covers* a polygon if every point in the polygon is visible to some guard. Figure 1.12 gives three examples of the range of visibility available to single guards in polygons.

A natural question is to ask for the *minimum* number of guards needed to cover polygons. Of course, this minimum number depends on the "complexity" of the polygon in some way. We choose to measure complexity in terms of the number of vertices of the polygon. But two polygons with n vertices can require different numbers of guards to cover them. Thus we look for a bound that is good for *any* polygon with n vertices.[3]

Exercise 1.27. *For each polygon in Figure 1.8, find the minimum number of guards needed to cover it.*

Exercise 1.28. *Suppose that guards themselves block visibility so that a line of sight from one guard cannot pass through the position of another. Are there are polygons for which the minimum of our more powerful guards needed is strictly less than the minimum needed for these weaker guards?*

Let's start by looking at some examples for small values of n. Figure 1.13 shows examples of covering guard placements for polygons with a small number of vertices. Clearly, any triangle only needs one guard to cover it. A little experimentation shows that the first time two guards are needed is for certain kinds of hexagons.

Exercise 1.29. *Prove that any quadrilateral needs only one guard to cover it. Then prove that any pentagon needs only one guard to cover it.*

[3] To find the minimum number of guards for a particular polygon turns out to be, in general, an intractable algorithmic task. This is an instance of another NP-complete problem; see the Appendix.

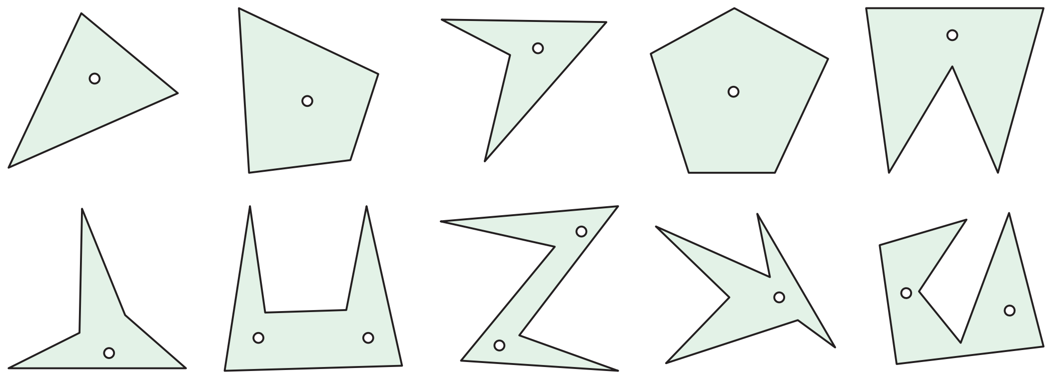

Figure 1.13. Examples of guard placements for different polygons.

Exercise 1.30. *Modify Lemma 1.18 to show that one guard placed anywhere in a convex polygon can cover it.*

By the previous exercise, convex polygons need only one guard for coverage. The converse of this statement is not true, however. There are polygons that need only one guard but which are not convex. These polygons are called *star* polygons. Figure 1.8(c) is an example of a star polygon.

While correct placement avoids the need for a second guard for quadrilaterals and pentagons, one can begin to see how reflex vertices will cause problems in polygons with large numbers of vertices. Because there can exist only so many reflex angles in a polygon, we can construct a useful example, based on prongs. Figure 1.14 illustrates the comb-shaped design made of 5 prongs and 15 vertices. We can see that a comb of n prongs has $3n$ vertices, and since each prong needs its own guard, then at least $\lfloor n/3 \rfloor$ guards are needed. Here the symbols $\lfloor\ \rfloor$ indicate the *floor* function: the largest integer less than or equal to the enclosed argument.[4] Thus we have a lower bound on Klee's problem: $\lfloor n/3 \rfloor$ guards are sometimes necessary.

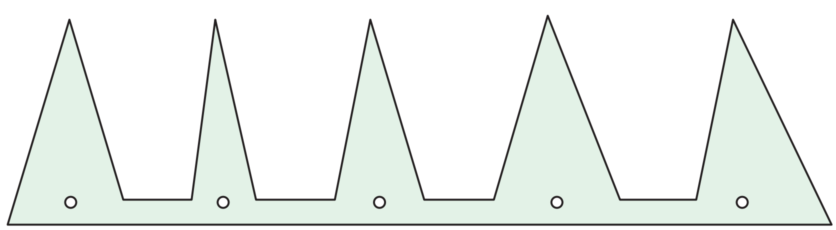

Figure 1.14. A comb-shaped example.

[4] Later we will use its cousin, the *ceiling* function $\lceil\ \rceil$, the smallest integer greater than or equal to the argument.

Exercise 1.31. *Construct a polygon P and a placement of guards such that the guards see every point of ∂P but P is not covered.*

UNSOLVED PROBLEM 4 Visibility Graphs

The *visibility graph* of a polygon P is the graph with a node for each vertex of P and an arc connecting two nodes when the corresponding vertices of P can see one another. Find necessary and sufficient conditions that determine when a graph is the visibility graph of some polygon.

Now that we have a lower bound of $\lfloor n/3 \rfloor$, the next question is whether this number always suffices, that is, is it also an upper bound for all polygons? Other than proceeding case by case, how can we attack the problem from a general framework? The answer lies in triangulating the polygon. Theorem 1.4 implies that every polygon with n vertices can be covered with $n - 2$ guards by placing a guard in each triangle, providing a crude upper bound. But we have been able to do better than this already for quadrilaterals and pentagons. By placing guards not *in* each triangle but on the *vertices*, we can possibly cover more triangles by fewer guards. In 1975, Vašek Chvátal found a proof for the minimum number of guards needed to cover any polygon with n vertices. His proof is based on induction, with some delicate case analysis. A few years later, Steve Fisk found another, beautiful inductive proof, which follows below.

Theorem 1.32 (Art Gallery). *To cover a polygon with n vertices, $\lfloor n/3 \rfloor$ guards are needed for some polygons, and sufficient for all of them.*

Proof. We already saw in Figure 1.14 that $\lfloor n/3 \rfloor$ guards can be necessary. We now need to show this number also suffices.

Consider a triangulation of a polygon P. We use induction to prove that each vertex of P can be assigned one of three colors (i.e., the triangulation can be *3-colored*), so that any pair of vertices connected by an edge of P or a diagonal of the triangulation must have different colors. This is certainly true for a triangle. For $n > 3$, Corollary 1.9 guarantees that P has an ear abc, with vertex b as the ear tip. Removing the ear produces a polygon P' with $n - 1$ vertices, where b has been removed. By the induction hypothesis, the vertices of P' can be 3-colored. Replacing the ear, coloring b with the color not used by a or c, provides a coloring for P.

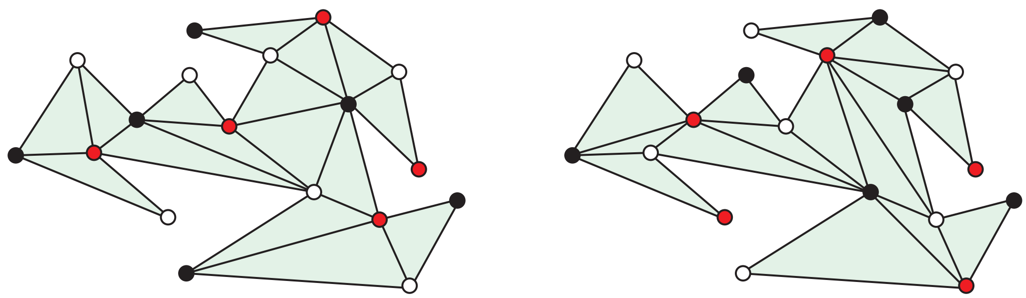

Figure 1.15. Triangulations and colorings of vertices of a polygon with $n = 18$ vertices. In both figures, red is the least frequently used color, occurring five times.

Since there are n vertices, by the pigeonhole principle, the least frequently used color appears on at most $\lfloor n/3 \rfloor$ vertices. Place guards at these vertices. Figure 1.15 shows two examples of triangulations of a polygon along with colorings of the vertices as described. Because every triangle has one corner a vertex of this color, and this guard covers the triangle, the museum is completely covered. □

Exercise 1.33. *For each polygon in Figure 1.16, find a minimal set of guards that cover it.*

Exercise 1.34. *Construct a polygon with $n = 3k$ vertices such that placing a guard at every third vertex fails to protect the gallery.*

The classical art gallery problem as presented has been generalized in several directions. Some of these generalizations have elegant solutions, some have difficult solutions, and several remain unsolved problems. For instance, the shape of the polygons can be restricted (to polygons with right-angled corners) or enlarged (to include polygons with holes), or the mobility of the guards can be altered (permitting guards to walk along edges, or along diagonals).

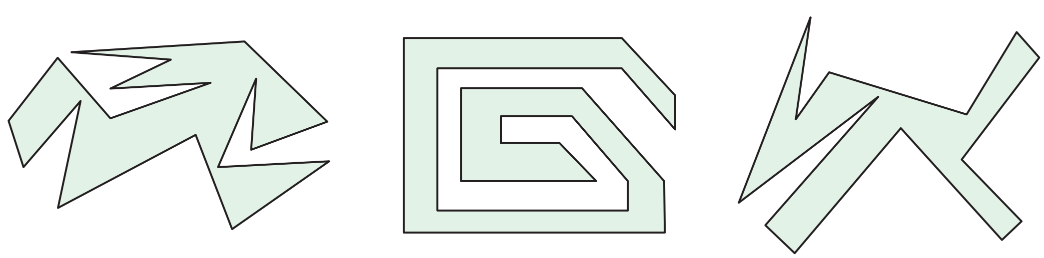

Figure 1.16. Find a set of minimal guards that cover the polygons.

Exercise 1.35. *Why is it not possible to easily extend Fisk's proof above to the case of polygons with holes?*

Exercise 1.36. *Using Exercise 1.14, derive an upper bound on the number of guards needed to cover a polygon with h holes and n total vertices. (Obtaining a tight upper bound is extremely difficult, and only recently settled.)*

When all edges of the polygon meet at right angles (an *orthogonal* polygon), fewer guards are needed, as established by Jeff Kahn, Maria Klawe, and Daniel Kleitman in 1980. In contrast, covering the exterior rather than the interior of a polygon requires (in general) more guards, established by Joseph O'Rourke and Derick Wood in 1983.

Theorem 1.37 (Orthogonal Gallery). *To cover polygons with n vertices with only right-angled corners, $\lfloor n/4 \rfloor$ guards are needed for some polygons, and sufficient for all of them.*

Theorem 1.38 (Fortress). *To cover the* exterior *of polygons with n vertices, $\lceil n/2 \rceil$ guards are needed for some polygons, and sufficient for all of them.*

Exercise 1.39. *Prove the Fortress theorem.*

Exercise 1.40. *For any $n > 3$, construct a polygon P with n vertices such that $\lceil n/3 \rceil$ guards, placed anywhere on the plane, are sometimes necessary to cover the exterior of P.*

UNSOLVED PROBLEM 5 — Edge Guards

An *edge guard* along edge e of polygon P sees a point y in P if there exists x in e such that x is visible to y. Find the number of edge guards that suffice to cover a polygon with n vertices. Equivalently, how many edges, lit as fluorescent bulbs, suffice to illuminate the polygon? Godfried Toussaint conjectured that $\lfloor n/4 \rfloor$ edge guards suffice except for a few small values of n.

UNSOLVED PROBLEM 6 Mirror Walls

For any polygon P whose edges are perfect mirrors, prove (or disprove) that only one guard is needed to cover P. (This problem is often stated in the language of the theory of billiards.) In one variant of the problem, any light ray that directly hits a vertex is absorbed.

The art gallery theorem shows that placing a guard at every vertex of the polygon is three times more than needed to cover it. But what about for a polyhedron in three dimensions? It seems almost obvious that guards at every vertex of any polyhedron should cover the interior of the polyhedron. It is remarkable that this is not so.

The reason the art gallery theorem succeeds in two dimensions is the fundamental property that all polygons can be triangulated. Indeed, Theorem 1.4 is not available to us in three dimensions: not all polyhedra are tetrahedralizable, as demonstrated earlier in Figure 1.7(c). If our polyhedron indeed was tetrahedralizable, then every tetrahedron would have guards in the corners, and all the tetrahedra would then cover the interior.

Exercise 1.41. *Let P be a polyhedron with a tetrahedralization where all edges and diagonals of the tetrahedralization are on the boundary of P. Make a conjecture about the number of guards needed to cover P.*

Exercise 1.42. *Show that even though the Schönhardt polyhedron (Figure 1.7) is not tetrahedralizable, it is still covered by guards at every vertex.*

Because not all polyhedra are tetrahedralizable, the "obviousness" of coverage by guards at vertices is less clear. In 1992, Raimund Seidel constructed a polyhedron such that guards placed at every vertex do *not* cover the interior. Figure 1.17 illustrates a version of the polyhedron. It can be constructed as follows. Start with a large cube and let $\varepsilon \ll 1$ be a very small positive number. On the front side of the cube, create an $n \times n$ array of 1×1 squares, with a separation of $1 + \varepsilon$ between their rows and columns. Create a tunnel into the cube at each square that does not quite go all the way through to the back face of the cube, but instead stops ε short of that back face. The result is a deep dent at each square of the front face. Repeat this procedure for the top face and the right face, staggering the squares so their respective dents do not intersect. Now imagine standing deep in the interior, surrounded by dent faces above and below, left and right, fore and aft. From a sufficiently central point, no vertex can be seen!

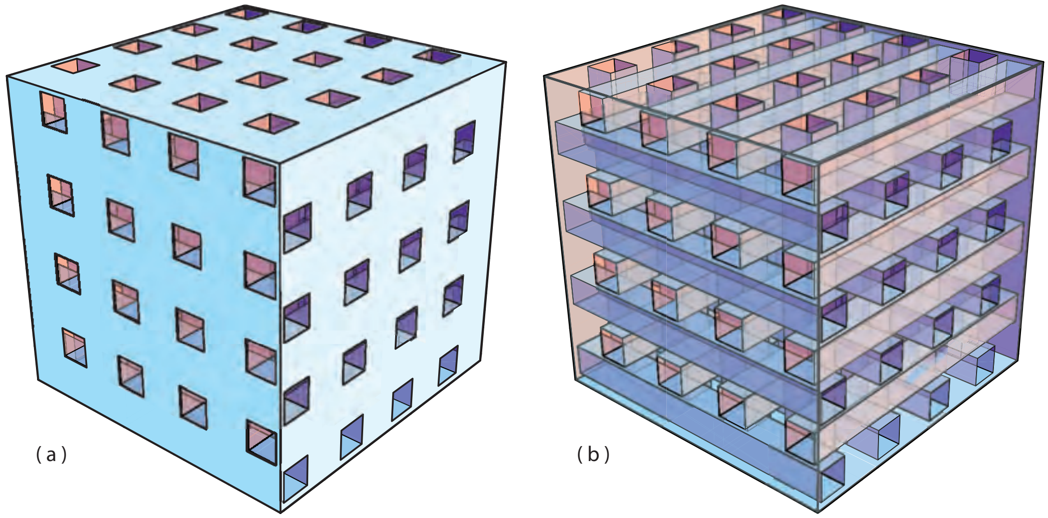

Figure 1.17. (a) The Seidel polyhedron with (b) three faces removed to reveal the interior.

Exercise 1.43. *Prove the above claim, which implies that guards at every vertex of the Seidel polyhedron do not cover the entire interior. Notice that this implies the Seidel polyhedron is not tetrahedralizable.*

Exercise 1.44. *Let n be the number of vertices of the Seidel polyhedron. What order of magnitude, as a function of n, is the number of guards needed to cover the entire interior of the polyhedron? (See the Appendix for the Ω notation that captures this notion of "order of magnitude" precisely.)*

1.4 SCISSORS CONGRUENCE IN 2D

The crucial tool we have employed so far is the triangulation of a polygon P by its diagonals. The quantities that have interested us have been combinatorial: the number of edges of P and the number of triangles in a triangulation of P. Now we loosen the restriction of only cutting P along diagonals, permitting arbitrary straight cuts. The focus will move from combinatorial regularity to simply preserving the area.

A *dissection* of a polygon P cuts P into a finite number of smaller polygons. Triangulation can be viewed as an especially constrained form of dissection. The first three diagrams in Figure 1.18 show dissections of a square. Part (d) is not a dissection because one of the partition pieces is not a polygon.

Given a dissection of a polygon P, we can rearrange its smaller polygonal pieces to create a new polygon Q of the same area. We say two

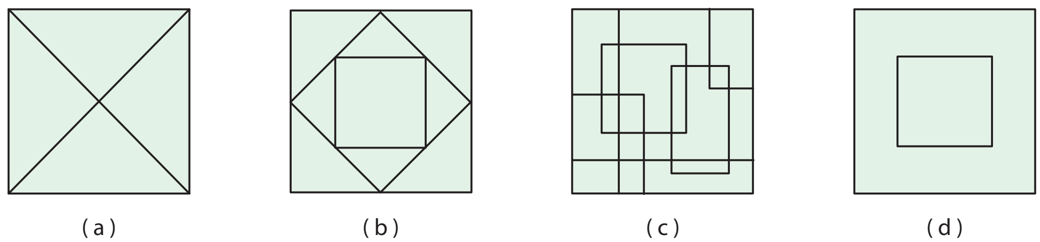

Figure 1.18. Three dissections (a)–(c) of a square, and (d) one that is not a dissection.

polygons P and Q are *scissors congruent* if P can be cut into polygons $P_1, \ldots, P_n$ which then can be reassembled by rotations and translations to obtain Q. Figure 1.19 shows a sequence of steps that dissect the *Greek cross* and rearrange the pieces to form a square, detailed by Henry Dudeney in 1917. However, the idea behind the dissection appears much earlier, in the work of the Persian mathematician and astronomer Mohammad Abu'l-Wafa Al-Buzjani of the tenth century.

The delight of dissections is seeing one familiar shape surprisingly transformed to another, revealing that the second shape is somehow hidden within the first. The novelty and beauty of dissections have attracted puzzle enthusiasts for centuries. Another dissection of the Greek cross, this time rearranged to form an equilateral triangle, discovered by Harry Lindgren in 1961, is shown in Figure 1.20.

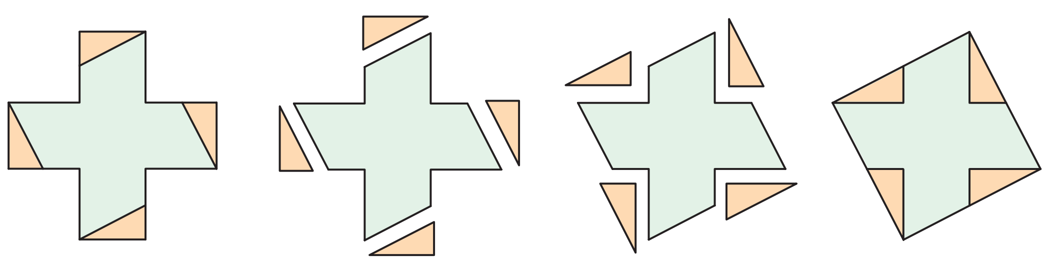

Figure 1.19. The Greek cross is scissors congruent to a square.

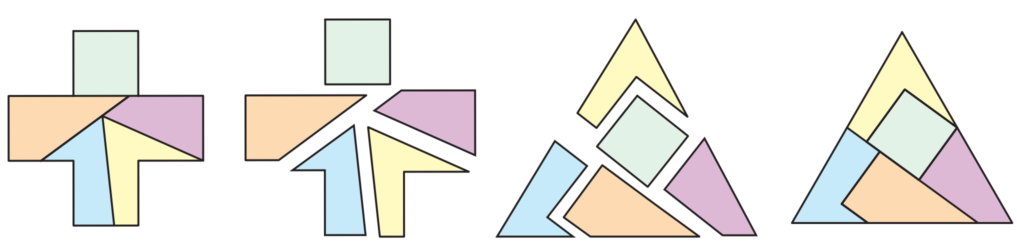

Figure 1.20. Lindgren's dissection of a Greek cross to an equilateral triangle.

Exercise 1.45. *Find another dissection of the Greek cross, something quite different from that of Figure 1.19, that rearranges to form a square.*

Exercise 1.46. *Find a dissection of two Greek crosses whose combined pieces form one square.*

Exercise 1.47. *Show that any triangle can be dissected using at most three cuts and reassembled to form its mirror image. As usual, rotation and translation of the pieces are permitted, but not reflection.*

Exercise 1.48. *Assume no three vertices of a polygon P are collinear. Prove that out of all possible dissections of P into triangles, a triangulation of P will always result in the fewest number of triangles.*

If we are given two polygons P and Q, how do we know whether they are scissors congruent? It is obvious that they must have the same area. What other properties or characteristics must they share? Let's look at some special cases.

Lemma 1.49. *Every triangle is scissors congruent with some rectangle.*

Figure 1.21 illustrates a proof of this lemma. Given any triangle, choose its longest side as its base, of length b. Make a horizontal cut halfway up from the base. From the top vertex, make another cut along the perpendicular from the apex. The pieces can then be rearranged to form a rectangle with half the altitude a of the triangle and the same base b. Note this dissection could serve as a proof that the area of a triangle is $ab/2$.

Lemma 1.50. *Any two rectangles of the same area are scissors congruent.*

Proof. Let R_1 be an $(l_1 \times h_1)$-rectangle and let R_2 be an $(l_2 \times h_2)$-rectangle, where $l_1 \cdot h_1 = l_2 \cdot h_2$. We may assume that the rectangles are not identical, so that $h_1 \neq h_2$. Without loss of generality, assume $h_2 < h_1 \leq l_1 < l_2$.

We know from $l_1 < l_2$ that rectangle R_2 is longer than R_1. However, for this construction, we do not want it to be *too* long. If $2l_1 < l_2$,

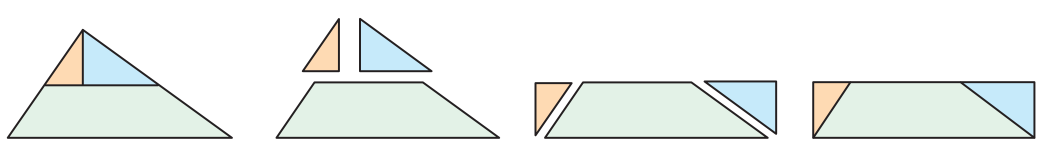

Figure 1.21. Every triangle is scissors congruent with a rectangle.

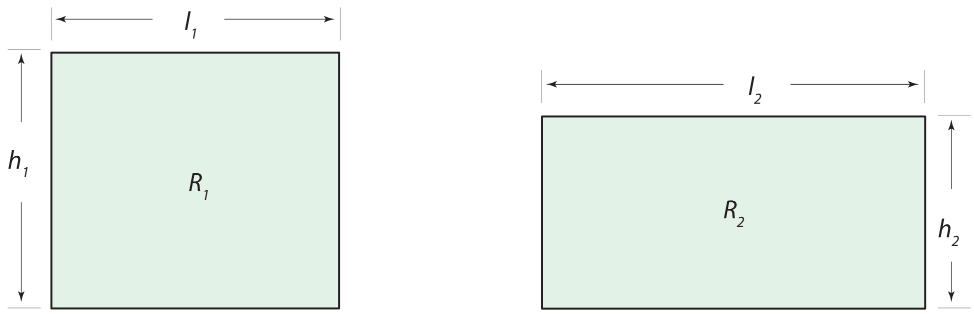

Figure 1.22. Two rectangles satisfying $h_2 < h_1 \leq l_1 < l_2 < 2l_1$.

then cut R_2 in half (with a vertical cut) and stack the two smaller rectangles on one another. This stacking will reduce the length of R_2 by half but will double its height. However, because $l_1 \cdot h_1 = l_2 \cdot h_2$, the height of the stacked rectangles $2h_2$ will still be less than h_1. Repeat this process of cutting and stacking until we have two rectangles with $h_2 < h_1 \leq l_1 < l_2 < 2l_1$, as shown in Figure 1.22.

After placing R_1 and R_2 so that their lower left corners coincide and they are flush along their left and base sides, draw the diagonal from x, the top left corner of R_1, to y, the bottom right corner of R_2. The resulting overlay of lines, as shown in Figure 1.23(a), dissects each rectangle into a small triangle, a large triangle, and a pentagon. We claim that these dissections result in congruent pieces, as depicted in Figure 1.23(b). It is clear the pentagons C are identical. In order to see that the small triangles A_1 and A_2 are congruent, first notice that they are similar to each other as well as similar to the large triangle xoy, as labeled in Figure 1.23(a). Using $l_1 \cdot h_1 = l_2 \cdot h_2$, the equation

$$\frac{h_1 - h_2}{l_2 - l_1} = \frac{h_1}{l_2} \tag{1.2}$$

can be seen to hold by cross-multiplying. Because A_1 is similar to xoy, whose altitude/base ratio is h_1/l_2, and the height of A_1 is $h_1 - h_2$, equation (1.2) shows that the base of A_1 is $l_2 - l_1$. But since the base

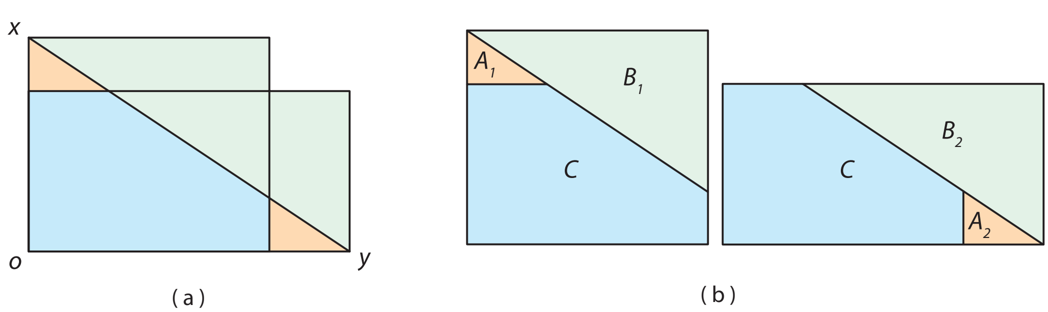

Figure 1.23. Any two rectangles of the same area are scissors congruent.

length of A_2 is $l_2 - l_1$, it follows that A_1 and A_2 are congruent. A nearly identical argument shows that the large triangles B_1 and B_2 are congruent. The theorem follows immediately. $\square$

Exercise 1.51. *Let polygon P_1 be scissors congruent to polygon P_2, and let polygon P_2 be scissors congruent to polygon P_3. Show that polygon P_1 is scissors congruent to polygon P_3. In other words, show that scissors congruence is* transitive.

Exercise 1.52. *Dissect a 2×1 rectangle into three pieces and rearrange them to form a $\sqrt[3]{4} \times \sqrt[3]{2}$ rectangle.*

It is immediate that scissors congruence implies equal area, but the converse is by no means obvious. This fundamental result was proved by Farkas Bolyai in 1832 and independently by Paul Gerwien in 1833.

Theorem 1.53 (Bolyai-Gerwein). *Any two polygons of the same area are scissors congruent.*

Proof. Let P and Q be two polygons of the same area α. Using Theorem 1.4, dissect P into n triangles. By Lemma 1.49, each of these triangles is scissors congruent to a rectangle, which yields n rectangles. From Lemma 1.50, these n rectangles are scissors congruent to n other rectangles with base length 1. Stacking these n rectangles on top of one another yields a rectangle R with base length 1 and height α. Using the same method, we see that Q is scissors congruent with R as well. Since P is scissors congruent with R, and R with Q, we know from Exercise 1.51 that P is scissors congruent with Q. $\square$

Example 1.54. The Bolyai-Gerwein theorem not only proves the existence of a dissection, it gives an algorithm for constructing a dissection. Consider the Greek cross of Figure 1.19, say with total area 5/2. We give a visual sketch of the dissection implied by the proof of the theorem to show scissors congruence with a square of the same area. The first step is a triangulation, as shown in Figure 1.24, converting the cross into 10 triangles, each of area 1/4 and base length 1. Second, each triangle is dissected to a rectangle of width 1 and height 1/4. Finally these are stacked to form a large rectangle of area 5/2.

Now starting from the square of area 5/2, a triangulation yields two triangles of base length $\sqrt{5}$, as shown in Figure 1.25. Each triangle is then transformed into a $\sqrt{5}/4 \times \sqrt{5}$ rectangle. Each rectangle needs to be transformed into another rectangle of base length 1 (and height 5/4). Since this rectangle is too long (as described in the proof of Lemma 1.50), it needs to be cut into two pieces and stacked. Then, the (stacked) rectangle is cut and rearranged to form two rectangles of base length 1.

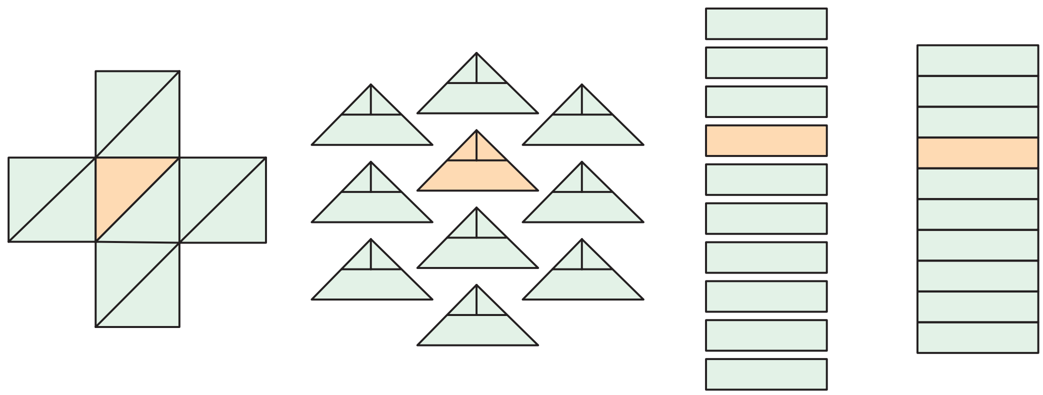

Figure 1.24. Cutting the Greek cross into a rectangle of base length 1 using the Bolyai-Gerwein proof. The transformations to the colored triangle are tracked through the stages.

Although the Bolyai-Gerwein proof is constructive, it is far from optimal in terms of the number of pieces in the dissection. Indeed, we saw in Figure 1.19 that a 5-piece dissection suffices to transform the Greek cross to a square.

Exercise 1.55. *Following the proof of the Bolyai-Gerwein theorem, what is the actual number of polygonal pieces that results from transforming the Greek cross into a square? Assume the total area of the square is 5/2 and use Figures 1.24 and 1.25 for guidance.*

★ Exercise 1.56. *Show that a square and a circle are not scissors congruent, even permitting curved cuts.*

It is interesting to note that the Bolyai-Gerwien theorem is true for polygons not only in the Euclidean plane, but in hyperbolic and elliptic geometry as well.

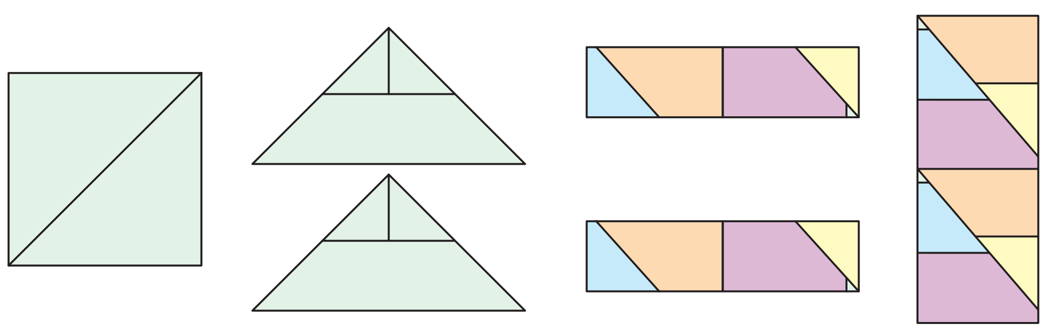

Figure 1.25. Cutting the square into a rectangle of base length 1 using the Bolyai-Gerwein proof. The last transformation is color-coded to show the fit of the pieces.

UNSOLVED PROBLEM 7 — Fair Partitions

For each positive integer n, is it always possible to partition a given convex polygon into n convex pieces such that each piece has the same area and the same perimeter? This has been established only for $n = 2$ and $n = 3$.

1.5 SCISSORS CONGRUENCE IN 3D

From the discussion above, we see that equal area suffices to guarantee a dissection of one polygon to another. Is this true in higher dimensions? That is, if we are given two polyhedra of the same volume, can we make them scissors congruent? In 1900, in his famous address to the International Congress of Mathematicians, the renowned mathematician David Hilbert asked the same question: Are any two polyhedra of the same volume scissors congruent? This problem was solved in the negative by Hilbert's student Max Dehn a few years later. Indeed, Dehn constructed two tetrahedra with congruent bases and the same height which are *not* scissors congruent. In order to understand Dehn's results, we need to take a closer look at polyhedra.

Unlike polygons, where angles only appear at the vertices, polyhedra have angles along edges as well. The angle along each edge of a polyhedron, formed by its two adjacent faces, is called the dihedral angle.

Definition. The *dihedral angle* θ at the edge e of a polyhedron shared between two faces f_1 and f_2 is the angle between two unit normal vectors n_1 and n_2 to f_1 and f_2, respectively. Thus $n_1 \cdot n_2 = \cos\theta$. By convention, the normal vectors point to the exterior of the polyhedron, and the dihedral angle at e is the interior angle.

For example, the dihedral angle along each edge of a cube is $\pi/2$. For further examples of dihedral angles, we will use Figure 1.26, which shows four tetrahedra embedded inside the cube.

Example 1.57. The tetrahedron on the left in Figure 1.27 repeats Figure 1.26(a) with labels. The dihedral angle along the edges AD, BC, and BD is $\pi/2$, and the edges AB and CD have dihedral angles of $\pi/4$. To find the dihedral angle along edge AC, we look back at the decomposition of the cube in Figure 1.6(c). Here the cube is tetrahedralized into six polyhedra congruent to the polyhedron on the left. Thus the dihedral angle along AC is $\pi/3$.

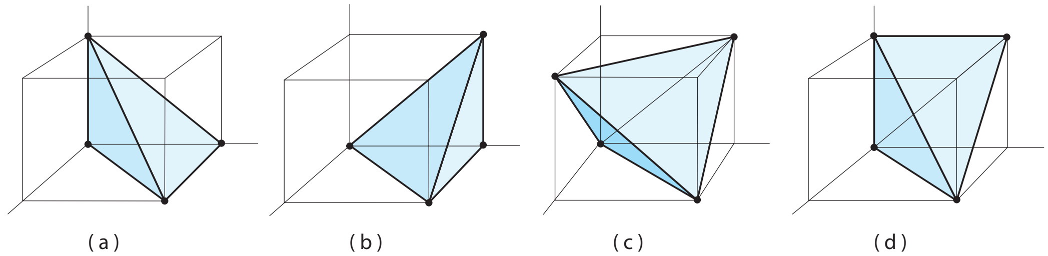

Figure 1.26. Four tetrahedra embedded inside the cube.

Example 1.58. The polyhedron on the right in Figure 1.27 repeats Figure 1.26(b). The dihedral angle along the edges AD, BD, and CD is $\pi/2$, because they are sides of the surrounding cube. Due to symmetry of the polyhedron, the edges AB, AC, and BC have the same dihedral angle. We draw the midpoint E of edge BC in order to calculate the dihedral angle AED along BC. If the cube of Figure 1.26(b) has side length x, then the length of DE is $x/\sqrt{2}$. Because the length of AD is x, the dihedral angle along BC is $\arctan \sqrt{2}$.

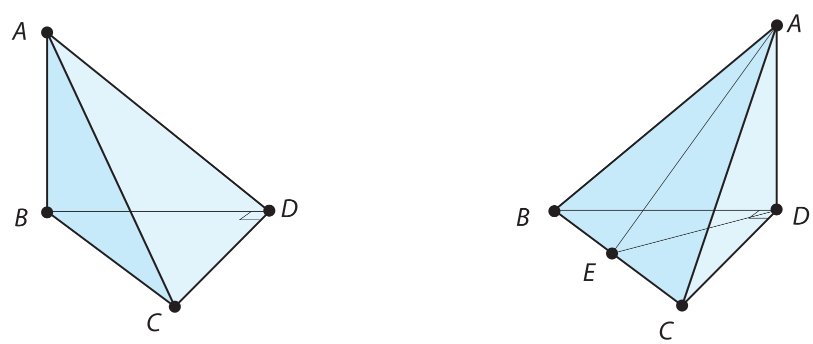

Figure 1.27. Two tetrahedra with congruent bases and the same height.

Exercise 1.59. *Find the dihedral angles of the tetrahedra in Figure 1.26(c) and (d).*

Exercise 1.60. *Find the dihedral angles of a regular dodecahedron.*

The dihedral angles are the key ingredient in understanding the ideas of Dehn. Instead of looking at these angles themselves, Dehn was interested in using them *up to rational multiples of* π. More precisely, let $f : \mathbb{R} \to \mathbb{Q}$ be a function from the real numbers to the rationals that satisfies three properties:

1. $f(v_1 + v_2) = f(v_1) + f(v_2)$ for all $v_1, v_2 \in \mathbb{R}$;
2. $f(qv) = qf(v)$ for all $q \in \mathbb{Q}$ and $v \in \mathbb{R}$;
3. $f(\pi) = 0$.

We call any such function a *d-function* (d for dihedral). For instance, for any d-function f, we see that

$$f\left(\frac{5\pi}{2}\right) = \frac{5}{2} \cdot f(\pi) = \frac{5}{2} \cdot 0 = 0.$$

Similarly, f maps any rational multiple of π to 0. We define a *rational angle* as an angle that is a rational multiple of π, and an *irrational angle* as one that is not.

For an edge e of a polyhedron, let $l(e)$ denote the length of e and let $\phi(e)$ denote the dihedral angle of e. For any choice of d-function f, Dehn's idea is to associate the value

$$l(e) \cdot f(\phi(e))$$

to each edge e, which he called its *mass*. Thus the mass of any edge is 0 when its dihedral angle is rational. We define the *Dehn invariant* of a polyhedron P to be the sum of the masses along the edges of P:

$$D_f(P) := \sum_{e \in P} l(e) \cdot f(\phi(e)).$$

Notice that the Dehn invariant depends on the choice of a d-function: For different choices of f, we get different Dehn invariants. The beauty of the Dehn invariant is that it is truly invariant under dissections. This is captured in the following theorem, which we prove using techniques invented by Hugo Hadwiger much later, in 1949.

Theorem 1.61 (Dehn-Hadwiger). *Let f be any d-function. If P is a polyhedron dissected into polyhedra $P_1, P_2, \ldots, P_n$, then*

$$D_f(P) = D_f(P_1) + D_f(P_2) + \cdots + D_f(P_n).$$

Proof. Let f be any d-function. The Dehn invariant of P sums the masses of the edges of P. After the dissection of P into several polyhedra, many new edges are introduced. Let e be an edge in the decomposition of P. There are three possible ways for e to appear in P, only one of which contributes to the mass sum.

1. Edge e is contained in an edge of P; see Figure 1.28(a). Let $\phi(e)$ be the dihedral angle of P along e and let $\{\phi_1(e), \phi_2(e), \ldots, \phi_k(e)\}$ be the set of dihedral angles of the polyhedral pieces along e in the decomposition. Then $l(e) \cdot f(\phi_i(e))$ is the mass contributed by the polyhedron P_i along edge e. The sum of the masses along e in the decomposition is

$$l(e) \cdot f(\phi_1(e)) + l(e) \cdot f(\phi_2(e)) + \cdots + l(e) \cdot f(\phi_m(e)),$$

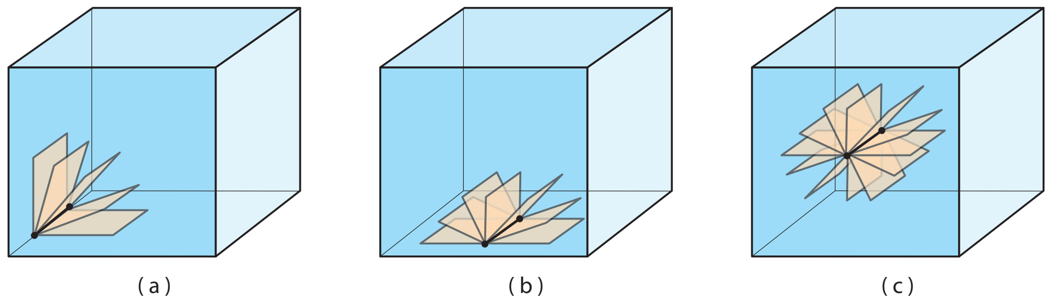

Figure 1.28. An edge of the dissection contained in (a) an original edge, (b) an interior of a face, or (c) the interior of the polyhedron.

which becomes

$$l(e) \cdot f(\phi_1(e) + \phi_2(e) + \cdots + \phi_m(e))$$

by Property 1 of a d-function. Since this is equal to $l(e) \cdot f(\phi(e))$, the masses add up in the required manner.

2. Edge e is contained in the interior of a face of P; see Figure 1.28(b). In this case, the sum of the masses becomes $l(e) \cdot f(\phi(e)) = l(e) \cdot f(\pi) = 0$. So a new edge created from a dissection that appears in the interior of a face of P has no mass.
3. Edge e is contained in the interior of P; see Figure 1.28(c). By a similar argument as before, $l(e) \cdot f(2\pi) = 0$, again contributing no new mass.

Thus the mass sum under the dissection depends only on the edges of P. As each edge e is covered exactly once by dissection edges, whose lengths sum to $l(e)$, the mass sum for any dissection is exactly the same as the mass sum for the original P. □

Corollary 1.62. *Let P and Q be polyhedra and let f be any d-function. If $D_f(P)$ does not equal $D_f(Q)$, then P and Q are not scissors congruent.*

Proof. We prove this by contradiction. Let P and Q be scissors congruent. Then there is a dissection of P into polyhedra $P_1, P_2, \ldots, P_n$. By the Dehn-Hadwiger theorem,

$$D_f(P) = D_f(P_1) + D_f(P_2) + \cdots + D_f(P_n) = D_f(Q),$$

the last equality following because the rearrangement of the polyhedral pieces forms Q. □

Example 1.63. Consider the tetrahedron on the left of Figure 1.27, calling it T_1. Example 1.57 shows that its set of dihedral angles is $\{\pi/2, \pi/3, \pi/4\}$. Because all the dihedral angles are rational, the mass for all the edges is 0. Thus $D_f(T_1) = 0$ for any d-function f.

Example 1.64. Consider the tetrahedron on the right of Figure 1.27, calling it T_2. If the cube of Figure 1.26(b) has side length 1, Example 1.58 shows three edges of length 1 with dihedral angle $\pi/2$ and three edges of length $\sqrt{2}$ with dihedral angle $\arctan\sqrt{2}$. Hence,

$$D_f(T_2) = 3\ f\left(\frac{\pi}{2}\right) + 3\sqrt{2}\ f(\arctan\sqrt{2}).$$

The first term is zero, but the second term need not be. For example, f could be the identity function on irrational multiples of π, in which case $D_f(T_2) = 3\sqrt{2}\arctan\sqrt{2} \neq 0$.

These two examples show the tetrahedra T_1 and T_2 in Figure 1.27, both having congruent bases and the same height and so the same volume, have different values for their Dehn invariant. By the Dehn-Hadwiger theorem, T_1 and T_2 are not scissors congruent. Generalizing this example, any polyhedron with all rational dihedral angles can never be dissected to a polyhedron with at least one irrational dihedral angle.

Exercise 1.65. *Show that a regular tetrahedron cannot be scissors congruent with a cube.*

★ Exercise 1.66. *Show that no Platonic solid is scissors congruent to any other Platonic solid.*

Although the Dehn-Hadwiger theorem can be used to show that two polyhedra are *not* scissors congruent, it does not directly tell us anything about the converse. In 1965, Jean-Pierre Sydler showed the following to be true, although its proof is quite involved.

Theorem 1.67 (Sydler). *Polyhedra P and Q are scissors congruent if $D_f(P) = D_f(Q)$ for every d-function f.*

This, along with the Dehn-Hadwiger theorem, show that Dehn invariants are a *complete* set of scissors-congruent invariants for polyhedra. To understand the power of this result, consider the tetrahedron in Figure 1.26(a). By Example 1.63, this tetrahedron as well as any cube have D_f equal to zero for any d-function f. Thus, by Sydler's theorem, a dissection of a cube exists whose rearrangement yields this tetrahedron! Indeed, Sydler demonstrated a beautiful construction showing how this dissection works. Figure 1.29 shows the tetrahedron being transformed into a prism whose base is an isosceles right triangle. It is then not hard to imagine how this prism can be made into a rectangular block, which then can be made into a cube.

Exercise 1.68. *Complete the construction above, dissecting the tetrahedron of Figure 1.29 into a cube.*

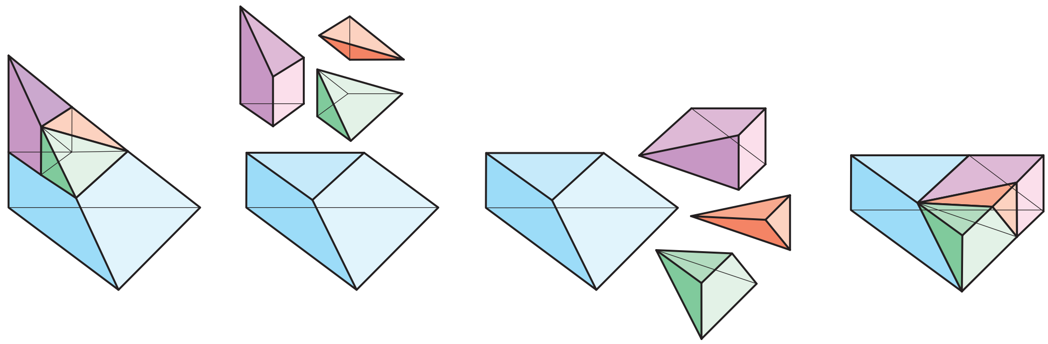

Figure 1.29. Sydler's dissection of a tetrahedron into a prism.

UNSOLVED PROBLEM 8 — Dehn Construction

Given two equal-volume polyhedra with all rational dihedral angles, construct an efficient algorithm for finding a dissection, as guaranteed by Sydler's theorem.

★ Exercise 1.69. *Let P be a $2 \times 1 \times 1$ rectangular prism. Cut P into eight or fewer pieces and rearrange the pieces to form a cube.*

UNSOLVED PROBLEM 9 — Five-Piece Puzzle

Can a $2 \times 1 \times 1$ rectangular prism be cut into five or fewer pieces, which can then be rearranged to form a cube?

SUGGESTED READINGS

Joseph O'Rourke. *Computational Geometry in C*. Cambridge University Press, 2nd edition, 1998.

Chapter 1 of this text covers the first three sections of this chapter with a more algorithmic slant.

Richard Stanley. *Enumerative Combinatorics*, Volume I and II. Cambridge University Press, 1997, 1999.

This monumental two-volume work covers the art of counting combinatorial objects. The classical text in graduate school. In particular, more than 100 objects are described that are all counted by the Catalan number.

Joseph O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, 1987.

A monograph on art gallery theorems, now more than 20 years old, but still (we think) the first source to consult. Out of print but available at `http://cs.smith.edu/~orourke//books/ArtGalleryTheorems/art.html`. For a more recent survey, see Jorge Urrutia, "Art gallery and illumination problems" (in Jörg-Rüdiger Sack and Jorge Urrutia, editors, *Handbook of Computational Geometry*, chapter 22, pages 973–1027, Elsevier, 2000).

Greg Frederickson. *Dissections: Plane & Fancy*. Cambridge University Press, 1997.

A delightfully readable book on dissections of polygons and polyhedra. Frederickson also wrote two other books on more specialized dissections: *Hinged Dissections: Swinging & Twisting* (Cambridge University Press, 2002) and *Piano-Hinged Dissections: Time to Fold!* (A K Peters, 2006).

Vladimir Boltyanskii, *Hilbert's Third Problem*. V. H. Winston & Sons, 1978.

A gem of a book, covering a broad scope of problems and proofs related to scissors congruence in 2D and 3D. Includes a reworking of Sylder's proof of the converse of the Dehn-Hadwiger theorem.