

Tiling a Polygon with Parallelograms¹

Richard Kenyon²

Abstract. Given a simple polygon in the plane we devise a quadratic algorithm for determining the existence of, and constructing, a tiling of the polygon with parallelograms. We also show that any two parallelogram tilings can be obtained from one another by a sequence of “rotations,” and give a condition for the uniqueness of such a tiling. Three generalizations of this problem, that of tiling by a fixed set of triangles, a fixed set of trapezoids, or parallelogram tiling for polygonal regions with holes, are shown to be NP-complete.

Key Words. Tiling, NP-complete.

1. Introduction. A **tiling** of a simple polygon (or polygonal region with holes) P by parallelograms is a set $\{T_i\}_{i=1}^k$ of parallelograms whose union is the closed region bounded by P and for which any two T_i, T_j are either disjoint, have an entire edge in common, or have a vertex in common. We also assume that if a parallelogram T_i touches the polygon P , then it touches along an entire edge (of both T_i and P) or at a vertex of both. (T_i and P may, of course, have several vertices or edges in common in this way.)

Thus, for example, the polygon in Figure 1 does not have a tiling by parallelograms in this sense. We discuss the more general definition of tiling (allowing subdivisions of edges of P) in Section 4. In particular we show that the problems we consider are essentially no more difficult (although they may take longer to solve).

Given some simple closed polygon P in the plane, under what conditions can we tile P by parallelograms? Given the existence of such a tiling, how can we construct it? How can we describe the set of all such tilings for a given polygon? These are the questions we discuss in this paper.

Tiling the plane with a fixed set of shapes has been considered by many authors. In 1966 Berger [B] proved that the problem of tiling the whole plane with a fixed set of shapes is undecidable.

A first instance of the problem of tiling a finite region with a fixed set of tiles was shown to be NP-complete in 1977 by Garey *et al.* [GJP].

These results have not prevented anybody from continuing to work on these problems, however. Closer to the paper at hand is the paper by Thurston [T]

¹ Parts of this paper originally appeared in a technical report for the conference “Journées de Géométrie Algorithmique” at Sophia-Antipolis, France, under the same title. The author was supported by a Chateaubriand Fellowship and this work was partially completed while the author was at Princeton University, supported by an IBM graduate fellowship.

² Institut Fourier, B.P. 74, 38402 Saint-Martin d’Hères, France.




Fig. 1. A pentagon can have no parallelogram tiling.

based on work with Conway which discusses tilings of polygonal regions by simple shapes such as dominoes and hexominoes, using some combinatorial group theory.

This author in his thesis [Ke] discusses tiling a generalized parallelogram with parallelograms.

Section 2 of this paper discusses tiling a simple polygon, and gives an algorithm for determining existence of and constructing such a tiling. The results in this section were found independently and simultaneously by Kannan and Soroker [KS]. In particular the proof of Theorem 3 is essentially theirs. Since the proof is quite natural I have reproduced it here. The proof of Lemma 2, which is essential to the algorithm, is quite different than that appearing in [KS], and has I believe an easier implementation.

Section 3 defines “rotations,” and shows that any two tilings of a simple polygon can be obtained from one another by rotations. This allows us to discuss the possible uniqueness of a parallelogram tiling. Section 4 discusses the problem of paving versus tiling.

Section 5 contains a proof that parallelogram-tiling a nonsimply connected polygonal region is NP-complete, and a proof that the problems of tiling by a fixed set of triangles or trapezoids can be reduced to tiling by general polygons (and hence NP-complete).

2. Simple Polygons. Let P be a simple polygon and let $\{v_1, v_2, \dots, v_n\}$ be the set of oriented edges of P , in counterclockwise order from some starting vertex. We call the **edge type** of an edge v the vector $[v]$ at the origin which is parallel to v and has the same length. (Thus the edge type records the length and direction of an edge.) The **absolute** edge types record edge type only up to sign. If $[v] = -[w]$, we say that v and w have the same absolute type but *opposite orientation*.

If v is an edge of P we define $v^\perp = i[v]$ to be the direction perpendicular to v , with orientation pointing from v into the interior of P .

If a parallelogram has edges $a, b, -a, -b$ in counterclockwise (cclw) order, we often denote that parallelogram by $[a, b]$.

LEMMA 1. *If a parallelogram tiling of P exists, then for each edge in P there must be an edge of the same absolute type but opposite orientation in P .*

PROOF. For each edge $v = v_0$ in polygon P , consider the parallelogram t_1 for which v_0 is an edge. Let v_1 be the edge opposite v_0 in t_1 . Let t_2 be the parallelogram which is adjacent to the edge v_1 , and inductively define t_n to be the parallelogram

adjacent to t_{n-1} along the edge v_{n-1} . This chain of parallelograms (t_0, t_1, \dots) always increases in the direction v^\perp (perpendicular to v pointing into P), so eventually the chain hits the boundary of P again. The boundary edge has the same length as v but is oppositely oriented. \square

For edges c, C which are paired in this way in a tiling, we denote by $\langle c, C \rangle$ the parallelogram chain from c to C .

To find a tiling, we first determine a **matching** of oppositely oriented edges (of the same length), which tells us which edges will be paired (as in Lemma 1) in the eventual tiling. This matching is subject to two rather obvious conditions:

1. Two matched pairs of edges of the same absolute edge type cannot cross each other with respect to the cyclic ordering of edges in P . (Two matched pairs a, A and b, B **cross** if between a and A in cyclic order around P exactly one of b, B occurs.) This follows from the proof of Lemma 1: the two chains of parallelograms connecting the matched pairs cannot cross each other in the usual sense.
2. Two matched edges must each "see" the other in the interior of P , in the sense that there is a monotone increasing path in the interior of P from one edge to the other. This is again because there must be a chain of parallelograms in the interior of P leading from one edge to the other.

A matching satisfying the above two conditions is called **good**, otherwise **bad**. See Figure 2.

LEMMA 2. *There is at most one matching satisfying the above two conditions.*

The proof is contained in Section 2.1.




Fig. 2. A good matching of edges in a polygon.

Once a good matching is found, there is another condition for the existence of a tiling; this is called **peripheral monotonicity**: given two matched pairs of edges a, A and b, B in P which cross, the edge b or B which is (in cyclic order) between a and A must be increasing in the a^\perp direction. That is, if the edge between a and A is, say, b , we must have $a^\perp \cdot [b] > 0$.

We define a **peripheral pair** to be a matched pair a, A which, in cyclic order, contains no other matched pair between a and A . In other words, the path $v_{a,A}$ along P from the front vertex of edge a to the back vertex of A contains no matched pairs. The monotonicity condition tells us that the “peripheral path” $v_{a,A}$ of P must be monotonically increasing in the a^\perp direction.

We shall see from Theorem 3 below that peripheral monotonicity of a matching actually implies that it is good (and hence, from Lemma 2, unique). It is clear that peripheral monotonicity implies condition 1 of a good matching.

THEOREM 3. *A polygon P has a tiling by parallelograms iff there is a matching of the edges of P which is peripherally monotonic.*

PROOF. Let Q be a closed oriented polygonal curve in the plane (not necessarily simple). For each point in the plane not on Q there is associated a **winding number** $w(x)$ of Q about x . It is an integer which describes how many times the curve Q winds around x . The winding number can be defined as follows: Take any half-line l from x to ∞ , oriented away from x , which avoids the vertices of Q and intersects edges of Q transversally if at all. Then $w(x)$ is the number of right-to-left crossings of Q with l , minus the number of left-to-right crossings. It is an exercise to show that $w(x)$ is independent of the choice of direction of the half-line l from x (and in fact any path from x to ∞ will do as long as it intersects Q transversally, and a finite number of times).

If Q is a simple polygon, $w(x) = 1$ for all points x in the region enclosed by Q , and 0 for points outside. We do not define $w(x)$ for points lying on the curve Q itself.

Let us suppose Q is a closed polygonal curve with a peripherally monotonic matching of its (oppositely oriented) edges as before (this condition on the matching does not require Q to be simple). Find a peripheral pair a, A of edges of Q .

We define $v'_{a,A}$ to be the path obtained by rigidly translating $v_{a,A}$ by the vector $-a$. Now $v'_{a,A}$ starts at the back of a and ends at the head of A . The region $R_{a,A}$ bounded by $v'_{a,A}$, a , $v_{a,A}$, and A can be tiled by parallelograms of the form $[a, v_i]$ for edges v_i of $v_{a,A}$. The polygonal curve Q' which is determined from Q by replacing the path $a, v_{a,A}, A$ by the path $v'_{a,A}$ now has two fewer edges, and the same matching as Q except that the pair a, A is gone. Because the edges were translated parallel to themselves and the cyclic order was preserved, the matching on Q' is still peripherally monotonic. See Figure 3.

In addition, we see that except for those points in the region $R_{a,A}$, the winding number of points with respect to Q' is unchanged from what it was for Q . For those points in $R_{a,A}$, the winding number is decreased by exactly 1.

Because the new curve Q' has again a peripherally monotonic matching, we can repeat the process, removing peripheral pairs of matched edges, until there is only




Fig. 3. A tiling obtained from the matching.

one pair left. At that point, the winding number of all points (not on the curve itself) is zero.

If Q was initially a simple polygon, we see *a posteriori* that during the intermediate stages $w(x)$ is always 1 or 0 when it is defined (because $w(x)$ was initially 0 or 1 and always decreases or remains fixed during removal of pairs). We can conclude from this that all the intermediate polygonal regions were noncrossing (at any crossing the winding number of one pair of regions opposite each other differs by two), and in fact the regions were disjoint unions of simple polygons, with Q winding exactly once around each subpolygon. So the parallelograms added were all disjoint.

Thus if P is a simple polygon this method creates a tiling of P by parallelograms.

For the other direction, suppose we are given a tiling of P . Then following parallelogram chains gives us a matching. Let a, A and b, B be two matched pairs whose matchings cross, i.e. (without loss of generality), we encounter the edges in the order a, b, A, B in the boundary P . Then somewhere in any tiling of P we have a parallelogram with edges $abAB$ in that order: this parallelogram occurs at the intersection of the two chains of parallelograms $\langle a, A \rangle$ and $\langle b, B \rangle$. However, $abAB$ is the boundary of a positively oriented parallelogram if and only if $a^\perp \cdot [b] > 0$. Thus peripheral monotonicity holds. \square

This proof is essentially the proof of [KS]; we thank D. Soroker for pointing out the errors in our original proof.

From these results we see that the problem of tiling by parallelograms and the problem of tiling with a *fixed* set of parallelograms are of the same difficulty. We must have parallelograms $[a, b]$ available for any two matched pairs of edges whose matchings cross, and these are the only parallelograms we actually need. This contrasts sharply with the case of triangulations and trapezoidizations: it is easy, for example, to triangulate a polygon, but not if we are only allowed to use triangles from a fixed set. See Section 5.

2.1. Finding the Matching. Theorem 3 gives a method for constructing the tiling, assuming we can find a matching satisfying the monotonicity condition. We give here an algorithm for the determination of such a matching.




Fig. 4. The edge coordinates in a polygon.

We match each absolute edge type separately. As before, let $\{v_i\}_{i=1}^n$ be the edges of P . Let $\{\pm[w_j]\}_{j=1}^t$ be the set of absolute edge types. We begin by separating each absolute edge type into **coordinate types** as follows.

Because of the noncrossing rule for good matchings, in cyclic order between any two matched edges w , W there must be an equal number of edges of each of the two orientations $\pm[w]$.

We assign coordinates to the edges of type $\pm[w]$ so that two edges of opposite orientation have the same coordinate if and only if there are an equal number of edges of type $+ [w]$ as $- [w]$ separating them (in cyclic order). One possibility is to assign some edge x_0 of type $[w]$ coordinate 0. Then assign each edge x of type $- [w]$ coordinate $n_+ - n_-$, where n_+ and n_- are respectively equal to the number of edges of type $[w]$ and the number of edges of type $- [w]$ between x_0 and x in cyclic order. Assign each edge y of type $[w]$ coordinate $1 + n_+ - n_-$, where n_+ and n_- correspond to the edges between x_0 and y . See Figure 4.

This procedure partitions an absolute edge type into different coordinate types; a matched pair of edges must have the same coordinate as well as opposite orientation.

It is important to note at this point that the edges having a fixed absolute type and coordinate *alternate in orientation* around P .

The matching among edges of the same coordinate and opposite orientation is now determined by the following lemma.

LEMMA 4. Suppose P has a good matching. Let V be a complete set of edges having the same absolute type and coordinate. Sort V by the perpendicular direction v^\perp for some $v \in V$; let u be a highest edge in V for which $u^\perp = v^\perp$. Let $V_u = \{u_1, u_2, \dots, u_k\}$ be the set of edges in V which are seen by u , in (cclw) cyclic order starting from u . Then u_1 and u_k are of different height and u must be matched to the lower of u_1 and u_k .

As an illustration of this lemma, see Figure 5. Recall that u “sees” v if there is a path from u to v in the interior of P which is monotonically increasing in the u^\perp direction.

PROOF. In this proof we use the phrase “ x is between a and b ,” or $x \in (a, b)$, to mean “ x is strictly between a and b in cclw cyclic order from a to b .”




Fig. 5. The highest \rightarrow is matched to the lowest \leftarrow which it sees.

First note that $V_u \neq \emptyset$ since otherwise u could not be matched; also, each $u_i \in V_u$ has the opposite orientation of u . Each edge of type $-[u]$ is matched with something *below* it, and each edge of type $[u]$ is matched to something *above* it. Suppose without loss of generality that u_1 is the lower of $\{u_1, u_k\}$.

Every edge in V is matched to something in V . Consider the edges to which u_1 could be matched. Firstly, u_1 cannot be matched with anything between u_k and u since then u would have to be matched with something between u and u_1 , which it could not by definition of u_1 . (Recall that matchings of edges of the same type cannot cross each other). See Figure 5(a).

Secondly, we show that u_1 cannot be matched to something in (u, u_1) . There are an equal number of edges $[u]$ as $-[u]$ between u and u_1 since u and u_1 have same coordinate type. Either all edges in V between u and u_1 are matched to each other, or one of these edges u' of type $-[u]$ is matched to an edge u'' of type $[u]$ *outside* (u, u_1) . However, in this latter case, there is a decreasing path leading from u' to its matched edge u'' ; this path must intersect the increasing path from u to u_1 , since u'' is outside (u, u_1) . If they intersect at a point y , then the path $u \rightarrow y \rightarrow u'$ is a monotone path from u to u' , contradicting the definition of u_1 . See Figure 5(b).

Thirdly, consider the possibility that u_1 is matched to something in (u_1, u_k) . Because u_1 is lower than u_k , there is no decreasing path from u_1 to anything between u_1 and u_k which is also lower than u (see Figure 5(c)), because such a path would necessarily obscure u_1 or u_k from the view of u . So u_1 cannot be matched to anything in (u_1, u_k) .

Thus u_1 is matched to u , the only remaining possibility. \square

2.2. Complexity of the Algorithm. The algorithm for parallelogram tilings of simple polygons in words is:

1. Separate edges into edge types and coordinate types.
2. Determine the good matching of oppositely oriented edges within the set of edges of the same absolute edge/coordinate type.
 - (a) Find the graph of "seeing" in a complete set of edges of the same absolute edge and coordinate type.
 - (b) Find the lower of u_1, u_k and match to u ; this divides the remaining edges $\pm[u]$ into two regions: apply the lemma in each region, and repeat to get the complete matching.

3. Find a peripheral matched pair and check monotonicity.
4. Slide the peripheral path across this pair and go back to step 3.

For a polygon of n edges, we have the following.

Step 1 is $O(n \log n)$, since separating edges into types (of which there may be $O(n)$) requires a certain amount of sorting.

Step 2(a). A simple method to construct the “seeing” graph for a particular edge/coordinate class $\pm[v]$ is to use a triangulation of the polygon. It is not hard to construct a triangulation in $O(n^2)$, in fact it can be done in linear time [C]. We only triangulate once, and then use the same triangulation for each edge/coordinate class.

From a given edge v in an edge/coordinate class, to find the edges which it sees, start at the triangle T_v containing edge v and move down in the tree of the triangulation (considering T_v to be the root), recording at each new triangle the set of points which are monotonically related to v . For each edge of a given triangle, either the entire edge or a connected subset of the edge (delineated by two points) is monotonically related to v , and this subset is determined by the subset in the parent triangle. Hence there is only a constant amount of work done in each triangle.

We see that for each class of size p_i we can find the “seeing” graph in time $O(np_i)$ (there are $n - 2$ triangles in a triangulation), and since the classes are disjoint, the total time for step 2(a) is $O(n^2)$.

Step 2(b) is $O(n^2)$, n^2 being the number of edges in the graph constructed.

Steps 3 and 4 are $O(k)$, where k is the number of parallelograms in the eventual tiling. Since $k = O(n^2)$ (in fact at most $n^2/4$, which can be seen by studying the process of tiling in steps 3 and 4) the algorithm in total is $O(n^2)$.

Determination of the nonexistence of a tiling using this algorithm is also $O(n^2)$, since we may not run into problems until step 3.

3. Rotations and Equivalence of Tilings. For a simple polygon there is only one edge matching (Lemma 2), but there can be many tilings (Figure 6).

If d is the number of absolute edge types, then we can think of a polygon P as the projection $\pi(\tilde{P})$, $\pi: \mathbf{R}^d \rightarrow \mathbf{R}^2$ of a polygon \tilde{P} sitting in the one-skeleton of a d -dimensional lattice in \mathbf{R}^d , so that the edge types are projections of the basis vectors of the lattice.

Then a tiling of P corresponds to a polygonal surface, lying on the 2-skeleton of this lattice (that is, composed of squares formed by pairs of basis vectors), spanning the polygon \tilde{P} , and which projects one-to-one to \mathbf{R}^2 under π .




Fig. 6. A rotation. Or, one matching, two tilings.

Each rotation in a tiling of P can be thought of as pushing the corresponding surface S in \mathbb{R}^d across a three-dimensional cube to get a new surface S' spanning the same boundary.

A **rotation** is defined as follows. Let x be a vertex in the tiling of P which lies in the interior of P and where exactly three parallelograms meet. Then the union of these three parallelograms forms a hexagon with opposite sides parallel. Such a hexagon has exactly two tilings by three parallelograms: replace the one present with the other one. This defines a **rotation from** x . Note that it preserves the number and type of parallelograms. (See Figure 6.)

We call two tilings **equivalent** if one can be obtained from the other by rotations. From the higher-dimensional picture it is easy to believe that any two tilings are equivalent in this sense. Our proof of this fact, however, is two dimensional.

THEOREM 5. *Any two parallelogram tilings of a simple polygon are equivalent.*

PROOF. We first prove that any two parallelogram tilings of a simple polygon P have the same number of tiles. Indeed, if there is any such tiling, there is a unique good matching (Lemma 2), and the number of parallelograms is the total number of crossings of pairs of matched edges. This number depends only on the cyclic ordering of the matched pairs around P , not on the tiling itself.

We now proceed by induction on the number of parallelograms $n(P)$ necessary to tile any polygon P . For $n(P) \leq 3$ the theorem is easy to check. Assume that two tilings are equivalent for any polygon which has a tiling of less than n tiles.

Let P be a polygon with a tiling of n parallelograms. Let c, C be a matched peripheral pair in a tiling of P (recall that this means the path $v_{c,C}$ from c to C around the boundary of P contains no other matched pairs). Let $\gamma = v_1, \dots, v_l$ be the monotonic path from c to C around the boundary of P . Let $\langle c, C \rangle$ be the chain of parallelograms from c to C .

For all v_i occurring in γ , there is a chain $\langle v_i, V_i \rangle$ which crosses the chain $\langle c, C \rangle$. This means that $\langle c, C \rangle$ contains exactly the parallelograms $[c, v_i]$, in some order.

If the order of parallelograms $[c, v_i]$ occurring in $\langle c, C \rangle$ is the same as the order of v_i occurring along γ , then the chain of parallelograms $\langle c, C \rangle$ lies exactly along the edges of γ on P , since c, C is a peripheral pair.

If the orders are different, then the chain does not follow γ exactly. (See for example chain $\langle a, a \rangle$ in Figure 7(b).) We show, however, that there is a sequence




Fig. 7. A "triangle" can have its vertices rotated by rotations within the region.

of rotations in the tiling that strictly decreases the number of parallelograms in the region between γ and $\langle c, C \rangle$. It then follows that there is a sequence of rotations which move the chain $\langle c, C \rangle$ so it lies along γ . (See $\langle a, a \rangle$ in Figure 7(a).) By induction, in the remaining region of P any two tilings are equivalent. Therefore any two tilings of P are equivalent, since we can move them both to tilings which are known to be equivalent.

Assume that the path $\langle c, C \rangle$ does not follow γ exactly. Then there are two edges v_i, v_j whose order is reversed between $\langle c, C \rangle$ and γ . Therefore the chain of parallelograms $\langle v_i, V_i \rangle$ crosses the chain $\langle v_j, V_j \rangle$ in the region between $\langle c, C \rangle$ and γ . They cross at a parallelogram $[v_i, v_j]$.

There is a triangular region T_{c, v_i, v_j} of parallelogram chains formed by the three chains between the three parallelograms $[c, v_i]$, $[c, v_j]$, and $[v_i, v_j]$.

By the lemma below (Lemma 6), by using rotations within T_{c, v_i, v_j} we can move the path $\langle c, C \rangle$ so that the parallelogram $[v_i, v_j]$ is *outside* the region between γ and the new chain $\langle c, C \rangle$ (and nothing else moves in). This completes the proof. \square

LEMMA 6. *Given a triangular region T (without holes) bounded by parallelogram chains between three parallelograms $[a, b]$, $[a, c]$, and $[b, c]$, there is a sequence of rotations within T which yields a tiling containing a rotatable hexagon formed by $[a, b]$, $[a, c]$, and $[b, c]$ and lying on the three parallelogram chains $\langle a, a \rangle$, $\langle b, b \rangle$, $\langle c, c \rangle$.*

As an illustration see Figure 7.

PROOF. We proceed by induction on the size of the triangular region T . If T is a hexagon, a single rotation will suffice. If T is larger than a hexagon, two of the “vertex” parallelograms are not adjacent (without loss of generality $[a, b]$ and $[a, c]$); take some parallelogram $[a, d]$ on the chain between these two. The chain from $[a, d]$ leading into T crosses one of the two other chains bounding T when it exits T , at say $[b, d]$. Then there is a smaller triangular region T' bounded by $[a, b]$, $[a, d]$, and $[b, d]$. By induction, through a sequence of rotations we arrive at a tiling containing a rotatable hexagon formed by $[a, b]$, $[a, d]$, and $[b, d]$ which lies on the chains $\langle a, a \rangle$, $\langle b, b \rangle$, and $\langle d, d \rangle$. By rotating this hexagon, $[a, b]$ moves closer to the $[a, c]$ and $[b, c]$ endpoints of chains $\langle a, a \rangle$ and $\langle b, b \rangle$, respectively.

We repeat this process, moving the three original vertices of T strictly closer at each step. Eventually they must be mutually adjacent; at that point they form a hexagon lying on all three chains and can be rotated themselves. \square

The same theorem is not true for nonsimply connected polygonal regions: see, for example, Figure 8.

COROLLARY 7. *Given a polygon P with a parallelogram tiling, the tiling is unique iff there are no triples of matched pairs of edges $\{a, A\}$, $\{b, B\}$, $\{c, C\}$ such that each pair crosses the other two.*




Fig. 8. Two inequivalent tilings of a nonsimply connected region.

PROOF. If there are no such 3-crossings, then in particular there are no interior vertices where exactly three parallelograms meet, so no rotations are possible and hence the tiling is unique.

On the other hand, let $\{a, A\}$, $\{b, B\}$, $\{c, C\}$ be such a 3-crossing. There is a triangular region formed by the chains between the parallelograms $[a, b]$, $[a, c]$, and $[b, c]$. By Lemma 6, there is a rotation possible. \square

4. Pavings Versus Tilings. We define a **paving** of a polygonal region P by parallelograms to be any exact covering of the region, $P = \bigcup T_i$, with a finite number of parallelograms T_i , such that the interiors of any two parallelograms are disjoint.

We show how to reduce a parallelogram-paving problem of a simple polygon P to a parallelogram-tiling problem, by adding vertices to P .

It is clear by a generalization of Lemma 1 that a polygon has a paving by parallelograms only if, for each edge direction, the total signed length (counting orientation) of edges having that direction is zero.

Let V be a complete set of edges of P all having the same direction (disregarding orientation). We assign coordinates to the *vertices* of edges in V in a generalization of the assignment of coordinate types in a tiling: after choosing an arbitrary vertex v_0 as having coordinate $c(v_0) = 0$, the coordinate $c(v)$ of a vertex v (of an edge in V) is the total signed length of edges of V between v_0 and v in cyclic order.

We can similarly assign to any point x in an edge of V (x is not necessarily a vertex) a coordinate $c(x)$ giving the signed length of edges in V from v_0 to x .

Let V' be the set of coordinates of the vertices of edges in V assigned in this manner. V' is finite. For each point x in an edge of V , if $c(x)$ is in V' , we subdivide the edge of P containing x , at the point x (in effect putting in a vertex at x .)

This adds a finite number of vertices to P , creating a new polygon P' . As a simple example, in Figure 1, the corresponding P' has a vertex in the center of the upper edge.

PROPOSITION 8. P has a paving iff P' has a tiling.

PROOF. Suppose we have a paving of P . There is a corresponding "matching" of the edges of P , which is obtained as in the case of a tiling, by following parallelogram chains. In this case, however, a parallelogram chain may split into two or more chains in the middle of P , or may merge with other chains.

Nevertheless, the coordinates of edge points $c(x)$ defined above are preserved under this matching, thus we have a matching induced on P' . It is in fact a good, peripherally monotonic matching by the corresponding properties of the paving of P . Thus there is a tiling of P' .

The other direction of the proposition is trivial. \square

If a quick algorithm for constructing pavings is being sought, this is not the best way; adding these dummy vertices may square the total number of vertices. A better way is to consider the problem as a flow problem, with the length of an edge of a particular direction being proportional to the flow through it. This point of view is taken in [KS], to which we refer the reader for an $O(n^2)$ algorithm.

5. NP-Complete Tiling Problems. Since we now know about how to tile a polygon with a fixed set of parallelograms, it is natural to ask about the same problem for more general shapes. The problem of characterizing sets of shapes which allow the solution of these tiling problems is likely to be a difficult one. There are certainly sets of shapes, though, for which the problem is hard; in particular we have:

THEOREM 9 [GJP]. *Given set of 1×1 squares with colored edges and an integer n , with $n >$ the number of colors, tiling an $n \times n$ square with these tiles so that the edges of adjacent squares have the same color is NP-complete.*

COROLLARY 10. *Given a set of polygonal tiles, each with $O(n)$ edges and area ≥ 1 , and a polygon P with $O(n)$ edges and area $O(n^k)$, tiling P is NP-complete.*

The area considerations are needed so that P is tiled by at most polynomially many tiles. We did not have this problem in the case of parallelogram tilings since we saw that $n^2/4$ parallelograms suffice to tile a polygon with n edges. Note also that by “tiling P ” here we mean any tiling in the sense of the introduction: a placement of translated copies of the tiles, covering P so that any two are disjoint or meet on a set of vertices and whole edges. Similarly for the intersection of a tile with the boundary polygon P .

We could also define tiling to include rotations of the tiles; the proof below can be adapted to show NP-completeness for this case too.

PROOF OF COROLLARY 10. Note that tiling P is in NP; there are at most n^k tiles needed to tile P , and given a potential tiling we can check that it is legitimate as follows: subdivide P into connected regions cut out by the tile boundaries and the boundary of P ; for each such region check that it is covered by exactly one tile. Also, we check that adjacent tiles meet along whole edges (i.e., it is a tiling not a paving). All this checking involves a polynomial amount of work.

It is also easy to convert a problem about colored tiles to a problem about polygonal tiles: it is only necessary to replace each edge of a given color by a polygonal “key” that only fits with edges of the inverse key; see, for example, [R].

The only difference is that Garey *et al.* [GJP] do not assume any particular color for the boundary of the enclosing square: any color of tile can touch the boundary. This will not be the case if we replace colors by keys.

To this end we define some more tiles which fit only at the boundary of P . We define special keys k_t, k_l, k_b, k_r for the top, left, bottom, and right edges of P , and another key k_i . We increase our set of tiles by adding tiles which have edge keys in cclw order $k_a, k_i, k_c, -k_i$ where k_c ranges over all color keys, and k_a ranges over the set $\{k_t, k_l, k_b, k_r\}$ (and with the appropriate orientation). We also add four corner tiles, e.g., for the upper right corner, k_t, k_i, k_i, k_r .

Now let P be the polygon with edges

$$\underbrace{k_b, \dots, k_b}_{n+1} \underbrace{k_r, \dots, k_r}_{n+1} \underbrace{k_t, \dots, k_t}_{n+1} \underbrace{k_l, \dots, k_l}_{n+1}$$

There is a tiling of P with these tiles if and only if there is a tiling of an $n \times n$ square with the colored squares given. This is because a tiling of P can only use the extra tiles at the boundary of P , and must use exactly those tiles.

The interior $n \times n$ square of P is tiled using only the original color-keyed tiles.

The polygon P we defined has at most $O(nm)$ edges, where m is the maximum number of edges in a key. We can certainly arrange it so that $m \leq O(\log n)$, which shows that the problem with n edges is also NP-complete. \square

COROLLARY 11. *Given a set of triangles of area ≥ 1 and polygon P with n edges and area $O(n^k)$, the problem of tiling P with the triangles is NP-complete.*

PROOF. Let T be a set of polygonal shapes, and let P be a polygon. For each tile $t_i \in T$, we triangulate t_i in a special way, such that no two edges in the triangulation have the same length, unless they are edges in the boundary of t_i . This requires in general adding vertices in the interior of t_i .

For each $t_i \in T$, we repeat the procedure, assuring that no two edges in the entire set of edges of all the triangulations are the same length, unless they are on the boundary of some tile t_i .

Then we have the property that in any tiling by these triangles, the triangles are forced to fit together to make copies of the tiles t_i .

In a tiling of P by this set of triangles, no edge of a triangle interior to its corresponding t_i lies in the boundary of P since its length is different than the length of edges of P . Thus this tiling of P by triangles gives a tiling of P by tiles t_i . \square

COROLLARY 12. *Given a set of trapezoids (quadrilaterals with two edges parallel) of area ≥ 1 and a polygon P with n edges and area $O(n^k)$, the problem of tiling P with the trapezoids is NP-complete.*

PROOF. A triangle can be tiled by nine trapezoids so that new vertices appear only in the interior and at the center of the edges of the triangle (see Figure 9). Given a finite set of triangles, we can subdivide each one into different trapezoids




Fig. 9. Nine trapezoids can tile a triangle with vertices in the center of the edges of the triangle.

so that the trapezoids are forced to fit together to recreate the triangles, as we did before with triangles.

The same argument shows that a tiling by trapezoids gives a tiling by the triangles. □

THEOREM 13. *Given a general polygonal region (not simply connected) with n edges, the problem of finding a parallelogram tiling is NP-complete.*

Note that this problem is different from the previous two (Corollaries 11 and 12) in that we do not fix the tiles beforehand. Tiling a polygonal region with a *previously defined* set of parallelograms is polynomial in the size of the polygonal region.

PROOF OF THEOREM 13. The problem is certainly in NP: there are at most $n^2/4$ parallelograms in any tiling.

We will show that solution of this problem allows us to solve the **subset sum** problem: given a finite set of positive integers $A = \{a_i\}$ and a positive integer K , find a subset $A' \subset A$ such that

$$\sum_{a_i \in A'} a_i = K.$$

That the subset sum problem is NP-complete is due to Karp [Ka]. We can assume in this problem that $K < \frac{1}{2} \sum_A a_i$.

Suppose we are given an instance of the subset sum problem, with a set of positive integers $A = \{a_1, a_2, \dots, a_k\}$ in increasing order. Let $B = \{b_1, b_2, \dots, b_k\}$ be a set of k positive integers (in *decreasing* order) such that there are no disjoint subsets $B_1, B_2 \subset B$ which have the same sum. For example, we can take $b_i = 2^{k-i}$.

Let $E = \{(\pm a_i, \pm b_i)\}_{i=1}^k \subset \mathbb{Z}^2$. We construct a polygonal region P with edges in $E' = E \cup \{(\pm 1, 0), (0, \pm 1)\}$. The boundary of P will consist of two polygons, the outer one P_1 and the inner one P_2 , so that the region P is an annulus contained between P_1 and P_2 . The edges of the outer boundary P_1 in cclw order from $(0, 0)$ are

$$P_1 = (a_1, -b_1)(a_2, -b_2) \cdots (a_k, -b_k)(1, 0)(a_k, b_k)(a_{k-1}, b_{k-1}) \cdots (a_1, b_1), (0, 1), \\ (-a_1, b_1)(-a_2, b_2) \cdots (-1, 0)(-a_k, -b_k) \cdots (-a_1, -b_1)(0, -1).$$




Fig. 10. The annular region P .

Note that P_1 is convex. P_2 is the square which, starting from $(2K, 0)$, has edges $(1, 0)(0, 1)(-1, 0)(0, -1)$ in cclw order. See Figure 10.

Suppose we are given a tiling of P by parallelograms; it has edges from the set E' . The higher edge $(-1, 0)$ of P_1 must be matched with the top edge of P_2 , and similarly for the lower edge of P_1 and the lower edge $(1, 0)$ of P_2 . There are no other edges $(\pm 1, 0)$ in the boundary P_1 .

The left edge of P_2 must be matched with the vertical edge at the origin. If we consider the parallelogram chain γ from the origin to the left edge of P_2 , the x -direction lengths of the parallelograms in γ sum to $2K$; and these lengths are all in A .

No edge in γ can occur more than once, since each edge type occurs only once (in each orientation) on the boundary P_1 . The total change in the y direction of γ is 0; but by the definition of B this means that for every edge (a, b) in this path the corresponding edge $(a, -b)$ occurs.

Let A' be the set of $a \in A$ whose corresponding edge occurs in γ . Then we have $2 \sum A' = 2K$ or $\sum A' = K$. Thus by construction of a tiling we have found a solution to the subset sum problem. \square

It is interesting to note that a solution to the subset sum problem gives us quite easily a tiling: Suppose we have a solution to the subset sum problem $A' = \{a_{i_1}, a_{i_2}, \dots, a_{i_j}\} \subset A$. We take a path γ inside P from the origin to P_2 to be the sequence of edges:

$$\gamma = (a_{i_1}, b_{i_1})(a_{i_1}, -b_{i_1})(a_{i_2}, b_{i_2})(a_{i_2}, -b_{i_2}) \cdots (a_{i_j}, b_{i_j})(a_{i_j}, -b_{i_j}).$$

Note that γ ends at $(2K, 0)$.

We can tile P so that this path occurs in the tiling, as follows. We can think of P with the path γ adjoined as a simple polygon P' . The boundary of P' starting from the origin traverses P_1, γ, P_2 clockwise, and then γ in the reverse direction.

Each edge of absolute type $[(a_i, b_i)]$ occurs exactly four times in the boundary of P' : once in the upper half of P_1 , once in the lower half, and twice on γ . The matching of these edges is trivial: each edge on P_1 is matched to the corresponding one on γ .

The edges $(0, \pm 1)$ are matched straight across; those edges $(\pm 1, 0)$ are matched to P_2 as explained above. Any other absolute edge type of P' occurs exactly twice (twice on the boundary of P_1), once in each orientation. Thus finding the good matching in P' is trivial. It is also easy to see that it is peripherally monotonic, therefore P' is tileable.

6. Some Problems

1. How many parallelogram tilings of a regular $2n$ -gon are there?
2. Can we use an algorithm similar to that in Section 2 to tile parallelohedra by parallelopipeds? A parallelohedron is a three-dimensional polyhedron with faces which are parallelograms. What if the faces are simply polygonal?
3. Is there a general criterion for sets of tiles for which the tiling problem is not NP-complete?

Acknowledgments. I would like to thank Claire Kenyon for suggesting the problem and discussing the results, and the referees for helpful comments in clearing up my definitions.

References

- [B] R. Berger. *The Undecidability of the Domino Problem*. Memoirs of the American Mathematical Society, No. 66. AMS, Providence, RI, 1966.
- [C] B. Chazelle. Triangulating a simple polygon in linear time. *Proceedings of the 31st Annual Symposium on the Foundations of Computer Science*, 1990, pp. 220–230.
- [GJP] M. S. Garey, D. S. Johnson, and C. H. Papadimitriou, unpublished results, 1977.
- [KS] S. Kannan and D. Soroker. Tiling polygons with parallelograms. Preprint, IBM.
- [Ka] R. M. Karp. Reducibility among combinatorial problems, in *Complexity of Computer Computations* (R. E. Miller and J. W. Thatcher, eds.). Plenum, New York, 1972, pp. 85–103.
- [Ke] R. Kenyon. Self-Similar Tilings. Thesis, Princeton University, 1990.
- [R] J. Robinson. Undecidability and non-periodicity of tilings in the plane. *Invent. Math.*, **12** (1971), 177–209.
- [T] W. P. Thurston. Conway's tiling groups. *Amer. Math. Monthly*, October 1990, pp. 757–773.