Tiling a Polygon with Rectangles

Claire Kenyon[*]

LIP
ENS-Lyon
46, allée d'Italie
69364 Lyon Cedex 07
France

Richard Kenyon[†]

UFR de Mathématiques
Institut Fourier
BP 74
38402 Saint Martin d'Heres
France

Abstract

We study the problem of tiling a simple polygon of surface n with rectangles of given types (tiles). We present a linear time algorithm for deciding if a polygon can be tiled with $1 \times m$ and $k \times 1$ tiles (and giving a tiling when it exists), and a quadratic algorithm for the same problem when the tile types are $m \times k$ and $k \times m$.

1 Introduction

We present algorithms for tiling a simple region of Z^2 (i.e. polygonal region without holes) with rectangular tiles.

In the first part, we present a linear time algorithm (in the surface area n of the region) when there are two tile types, the $1 \times m$ and $l \times 1$ rectangles. Previously the only known algorithms other than exhaustive search were for the case of dominoes ($m = 2, l = 2$), where several combinatorial methods were known (using matchings, or max-flow-min-cut algorithms). Our algorithm generalizes a domino tiling algorithm of W. P. Thurston [1] based on ideas of J. H. Conway and Lagarias which rely on geometric group theory. After the fact, one can view this problem as a flow problem as in the case of dominoes.

In the second part, we give a quadratic algorithm for tiling with $k \times l$ and $l \times k$ rectangles (i.e. when there is only one tile type, up to rotation). This case appears more complicated, but is interesting because there seems to be no natural interpretation as a flow problem.

In the last section, we define the distance between two tilings of a region in terms of the number of elementary transformations needed to go from one to the other, and study the maximum distance.

It is known from Mike Robson [2] that tiling a region with holes with $1 \times m$ and $n \times 1$ tiles is NP-complete as soon as $m \geq 2$ and $l \geq 3$. On the other hand, efficient tiling algorithms and criteria have been obtained for restricted classes of polygons [3][4]. As far as rectangles are concerned, from [5] we see that our tiling problem is NP-complete for polygonal regions with holes as soon as $k \geq 2$ and $l \geq 3$ (optimal tile salvage problem).

[*]Supported by CNRS URA 1398.

[†]supported by CNRS.

2 Tiling a polygon with $1 \times m$ and $l \times 1$ tiles

We will only present the algorithm in the case $m = l = 3$, since generalizing the algorithm is straightforward from that case. By a *polygon* we will mean a simple polygon in $\mathbf{R}^2$, having vertices in $\mathbf{Z}^2$, and vertical or horizontal edges. A polygon P is described by its perimeter ∂P, which is a sequence of directed horizontal and vertical edges of length 1.

2.1 The method of Conway and Lagarias

Let a be a symbol associated with each horizontal step (of length one) to the right and b be the symbol associated with each step upward. Let a^{-1} denote a horizontal step to the left, and b^{-1} a vertical step down; then the perimeter of a polygon, read in a counterclockwise sense from an arbitrary starting point, is a word in a, b, a^{-1}, b^{-1}.

Let $F = \langle a, b | \rangle$ denote the free group generated by a and b and in general let $\langle A|R \rangle$ denote the group generated by the elements of A with relations $r = e$ for $r \in R$. Following Thurston [1], we define the tiling group of a set of tiles to be the quotient of the free group F by the relations describing the tiles. Let $G = \langle a, b | [a^3, b], [a, b^3] \rangle$ (where the notation $[x, y]$ is shorthand for $xyx^{-1}y^{-1}$).

Let P be the polygon in Z^2 which we want to tile. Choose a point on the boundary of P arbitrarily and label it by e. Going around P from that point in counterclockwise order, we get a word w in F. The word w can also be thought as a word in G in a natural way.

Fact 1 *[1] If P can be tiled, then $w = e$ in G.*

Proof idea: First, the polygon given by a single tile gives $w = e$: a horizontal 3×1 tile corresponds to the relation $a^3ba^{-3}b^{-1} = [a^3, b] = e$, and a vertical 1×3 tile corresponds to $ab^3a^{-1}b^{-3} = [a, b^3] = e$. Starting from any region P, removing a tile from P corresponds to changing w by applying one of the relators $[a^3, b]$ or $[a, b^3]$. □

The condition $w = e$ in G is not always sufficient. For example, the polygon figure 1 admits no tiling (with 1×3 and 3×1 tiles) even though w is trivial in G:

$$w = a^2b^3aba^{-2}b^{-3}a^{-1}b^{-1} = b^3a^3bb^{-3}a^{-2}a^{-1}b^{-1} = e.$$

If $w = e$ in G, then ∂P lifts to a closed path in the Cayley graph $\Gamma(G)$ of G. From that closed path, ∂P can be recovered by projection onto $Z^2 = G/[a,b]$ (defined to be the group obtained from G by adding the relator $[a,b]$, i.e. by quotienting by the normal closure of the relator $[a,b]$).

Define a 2-complex $\Gamma^2(G)$ by glueing disks onto $\Gamma(G)$, one disk for each occurence of the relators $[a,b^3]$ and $[a^3,b]$ in $\Gamma(G)$. A tiling of P then lifts to a surface in $\Gamma^2(G)$ spanning the loop w in $\Gamma(G) \subset \Gamma^2(G)$. Define a surface area in $\Gamma^2(G)$ by giving area 3 to each disk (i.e. the area of its projection on Z^2). In general, a surface spanning w in $\Gamma^2(G)$ does not necessarily project to a tiling of P, since the projection may be self-overlapping. But we have the following key observation:

Fact 2 *[1] If P is tileable, then any minimum area surface spanning w in $\Gamma^2(G)$ projects to a tiling of P.*

The problem is now reduced to finding a surface of minimum area.

In the case of domino tilings (when the tiles are 1×2 and 2×1, $\Gamma(G)$ is naturally embeddable in $\mathbb{R}^3$, and Thurston solves the problem by first defining a "height" function in $\Gamma(G)$, secondly observing that there is a lowest tiling T spanning w, whose highest point must be on the border w,and thirdly taking the highest point of w and showing that the tile of T at that point can be determined easily from the local conditions.

To generalize Thurston's approach, the main problem is finding the correct definition of a height function.

2.2 Defining a height function

In our case, the Cayley graph $\Gamma(G)$ is a bit difficult to visualize, since it cannot be naturally embedded in $\mathbb{R}^3$. However we only need to use a quotient graph of G. Let $H = G/a^3, b^3$. Since $a^3 = b^3 = e$ implies that $ab^3 = b^3a$ and that $a^3b = ba^3$, we have $H = \langle a, b | a^3, b^3 \rangle = \mathbb{Z}_3 * \mathbb{Z}_3$, the free product of $\mathbb{Z}_3$ by $\mathbb{Z}_3$. The Cayley graph $\Gamma(H)$ of H is simply a tree of triangles (see figure 2). given $x \in G$, let x' denote its projection in H and x'' denote its projection in $\mathbb{Z}^2 = G/[a,b]$. G is a semi-direct product of H by $\mathbb{Z}^2$, which means that an element of G is determined by its projections onto H and $\mathbb{Z}^2$:

Fact 3 *The map $G \to H \times \mathbb{Z}^2$ given by $x \mapsto (x', x'')$ is bijective.*

Now, the closed path w in $\Gamma(G)$ associated to the boundary of P maps to a closed path w' in $\Gamma(H)$.

Root $\Gamma(H)$ at some arbitrary vertex r, far from w', and define the height of $x \in H$ to be the distance from x to r in the graph $\Gamma(H)$. (For example, if n is the surface area of P, we can define r as the point $(ab)^{10n}$ in H).

Lemma 1 *If P is tileable, then there is a tiling of P whose highest point is on the boundary ∂P.*

Proof: Let T be a tiling of the polygon. Every point inside P is on the boundary of a tile, and so has an associated element of G (and so a height in H). Assume that a highest point x is in the interior of P. We claim that x cannot be at the corner of a tile. For, if x is, for example, the lower left corner of some tile, then x has neighbors labelled xa and xb. In $\Gamma(H)$, it is easy to see that no matter how x is with respect to r, either xa or xb is farther from r than x, hence higher: contradiction.

Without loss of generality, we can assume then that x is on the left side of a vertical tile. Then the tile on the left side of x must also be vertical (because x is not a corner and not on the boundary). Assume for example that x is as in figure 3. Then xb^2, which labels a neighbor of x, is also highest. Since it is not a corner either, this determines the position of the tile left of x: the two tiles surrounding x must form a rectangle.

Thus $y = xbab^{-1}$ and $z = xba^{-1}b^{-1}$ are two labels of the tiling. We observe that no matter where r is, one of the two points y and z must be at the same distance as x from r, hence also of maximum height. For example, z is also a maximum in figure 4. If z is on the border of P, we are done. Otherwise, by the same argument as above, the tile left of z must be vertical and form a rectangle with the previous two tiles. We can now replace the three vertical tiles by three horizontal tiles, forming a new, lower tiling (in the sense that every point of P is lower; see figure 5). Therefore a lowest tiling must have a highest point on its boundary. $\square$

Given P, we now consider the highest point x on the boundary ∂P of P. From the above lemma, we know that x is also the highest point of some tiling of P (if P is tileable). It is easy to see that x cannot be at a corner of ∂P. Assume that x is on a vertical side of ∂P, as in figure 6. The neighbors of x on ∂P are labelled xb and xb^{-1}. One of them must also have maximum height, for example xb. Since neither x nor xb can be at a corner of a tile, the only possibility is a vertical tile covering both x and xb.

Removing that tile, we get a region which can be tiled iff P can be tiled: iterating the process, we will obtain an algorithm for tiling P or deciding that P cannot be tiled by 1×3 and 3×1 tiles.

2.3 The algorithm

Input: a polygonal region P of $\mathbb{Z}^2$ of area n.
Initialization: Label some point of P by e. Let $r = (ab)^{10n} \in H = \langle a, b | a^3, b^3 \rangle$. Going around the boundary, label each point of ∂P by a word $w \in H$ and compute its height.
Repeat: Let x be a highest point of ∂P.
One of x's neighbors, say y, is also a highest point of

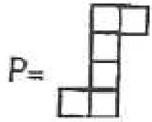

Figure 1: A necessary but not sufficient condition for tiling P

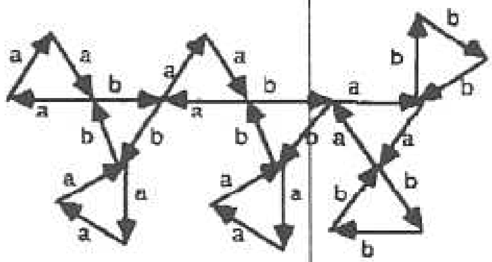

Figure 2: The Cayley graph of $H = \langle a, b | a^3 = b^3 = e \rangle$

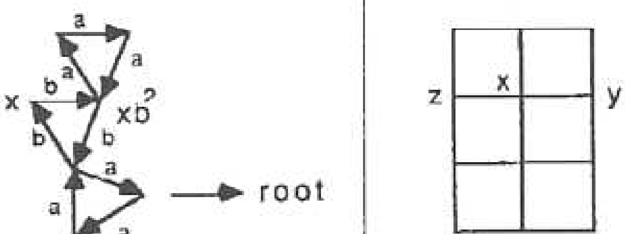

Figure 3: Proof of lemma 1: positioning the tiles surrounding the highest point

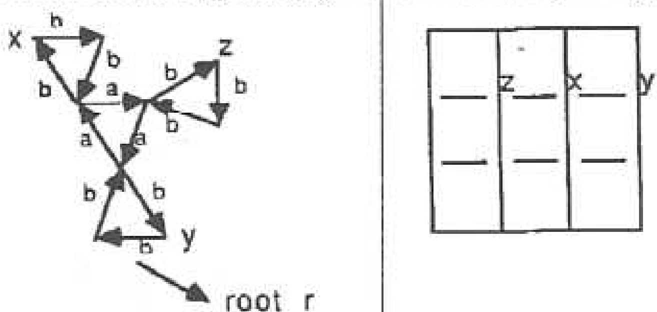

Figure 4: Proof of lemma 1: building a square with highest tiles

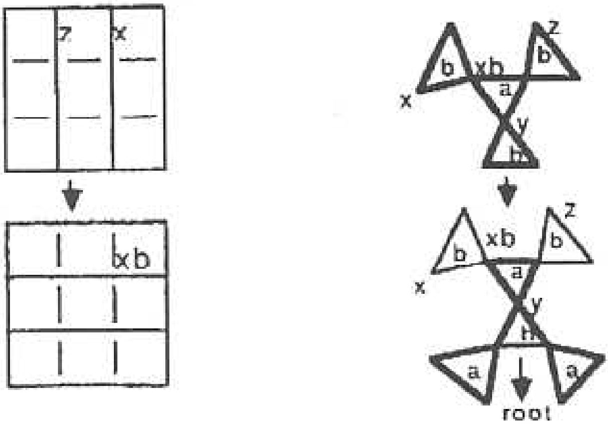

Figure 5: Proof of lemma 1: Removing interior maxima by changing the tiling of a square

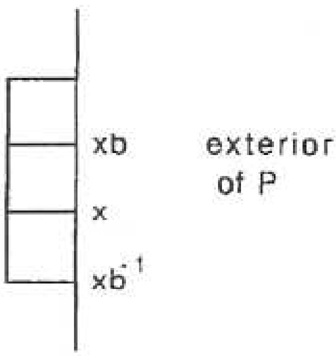

Figure 6: Tiling P: putting down the first tile

∂P. Take the tile T in P covering x and y (vertical if xy is a vertical segment, horizontal otherwise). Remove T from P.
Update ∂P, the labels and heights.
Until: P is tiled or some point of P is given two different labels (then P cannot be tiled).

It is easy to see that, since all the updates are local operations, the algorithm can be implemented so as to have linear running time.

If the tiles are $k \times 1$ and $1 \times l$ instead of 1×3 and 3×1, the algorithm can easily be generalized. The group to consider is $H = \langle a, b | a^k, b^l \rangle$, whose Cayley graph is a tree of alternating k-cycles and l-cycles. The height is similarly defined, as the distance to some arbitrary root (sufficiently far away). Given a highest boundary point, the lowest tiling must contain the tile which has x in the middle or co-middle of its long side.

Theorem 1 *A polygon of area n can be tiled by $1 \times l$ and $k \times 1$ tiles, or proved not to be tileable, in time linear in n.*

3 Tiling a polygon with $k \times l$ and $l \times k$ rectangles

We will only present the algorithm in the case $k = 2$ and $l = 3$, since generalizing the algorithm is easy from that special case.

3.1 The tiling group and its quotients

The two rectangular tiles define the relations $a^2b^3 = b^3a^2$ and $a^3b^2 = b^2a^3$, hence the tiling group is $G = \langle a, b | [a^2, b^3], [a^3, b^2] \rangle$. Let H be the quotient group $H = G/a^3, b^3 = \langle a, b | a^3, b^3 \rangle$, and K be the group $K = G/a^2, b^2 = \langle a, b | a^2, b^2 \rangle$. We have: $\mathbf{Z}^2 = G/[a, b]$. Given $w \in G$, let w', w'', w''' be the projections of w on H, K and $\mathbf{Z}^2$.

Fact 4 *The map $w \mapsto (w', w'', w''')$ is bijective.*

Defining a height function is now a more complex operation, since it apparently involves both quotient groups H and K. However, it turns out that using the labels in H already gives us "most" of the information that we need to construct a tiling.

3.2 Defining a height function

We label some arbitrary point of ∂P by e. Going around the boundary, we label each point of ∂P by an element of H. Rooting H at some point r sufficiently far away, we can define a height function as being the distance from r in the graph $\Gamma(H)$. A tiling of P extends the labellings and heights to all the points of P which are on the border of a tile (each tile has two interior points whose height is undefined).

Lemma 2 *If P is tileable, then there is a tiling of P whose maximum height h occurs either on the boundary or within distance 1 of the boundary (in which case there is a boundary point of height $h-1$).*

Proof sketch: Let x be a highest interior point. Using the same elementary arguments as in section 2, it is easy to see that x cannot be at the corner or in the middle of the short side of a rectangle: x must be on the long side of a rectangle. Doing case-by-case analysis, we find that the tiling must contain a "block" formed by three adjacent rectangles, as in figure 7 (where x, x', y, y' are all highest points of the tiling) for vertical tiles, or a rotated version for horizontal tiles.

Definition 1 *Given a tiling of P, a* block *is a set of 9 adjacent tiles, forming a 3×6 or 6×3 rectangle, such that the points inside the block which are labelled are all highest in the tiling.*

What is there below that block? An elementary case-by-case analysis shows that there are only three cases, illustrated figure 8: either there is another block, lined up with the previous block, forming a 6×6 rectangle, or there is another block, shifted by 3, or there are two horizontal tiles, forming a 6×5 rectangle.

In the first case, we replace the 6 vertical rectangles by 6 horizontal rectangles, eliminating the highest points. In the second case, we look at what is below the second block. In the third case, we put the two horizontal tiles above the block, moving the block down by 2, which either eliminates the highest points or moves them closer to the boundary of P below the block. Iterating the process, we finally get a block which touches the boundary, and the highest point has been pushed to within distance 1 of ∂P. □.

3.3 Putting down the first tile

Let x be the highest point on the boundary ∂P. Assume that x is on a vertical part of ∂P (since x is maximum, it cannot be at a corner), and that the interior of P is to the right of x. One of its neighbors on ∂P is also highest, say y.

If the tile at x is a vertical rectangle, then it is easy to see that it has to be a rectangle covering x and y (see figure 9 – otherwise some interior point would have height 2 more than the height of x).

If it is horizontal, then the points along one of the long sides of the rectangle have height equal to 1 more than the height of x: they must therefore be highest points of P. As in the previous section, by doing case analysis to study the position of the tiles adjacent to that first tile, we see that there has to be a block, in two possible positions.

Lemma 3 *Let B be a vertical block of a tiling of P. The area to the right of B must*
1) contain some points in the exterior of P or
2) be covered by another block at the same level (figure 10a) or
3) be covered by two vertical rectangles (figure 10b) or
4) be covered by one vertical rectangle and one block shifted by 3 (figure 10c,c') or
5) be covered by two blocks shifted by 3 (figure 10d).

The proof is by inspection.

Up to rotating the 6×6 figure to change the tiling, we can assume that case a never occurs in the tiling. We will now apply the lemma repeatedly. Let B_0 be the block at x. We say that a block B is *reachable* from B_0 if there is a sequence of blocks $B_1, B_2, ..., B_m = B$ such that B_i is in the right neighborhood of B_{i-1} (as

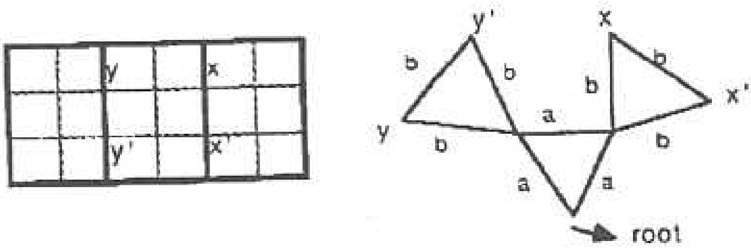

Figure 7: Interior highest points of a tiling

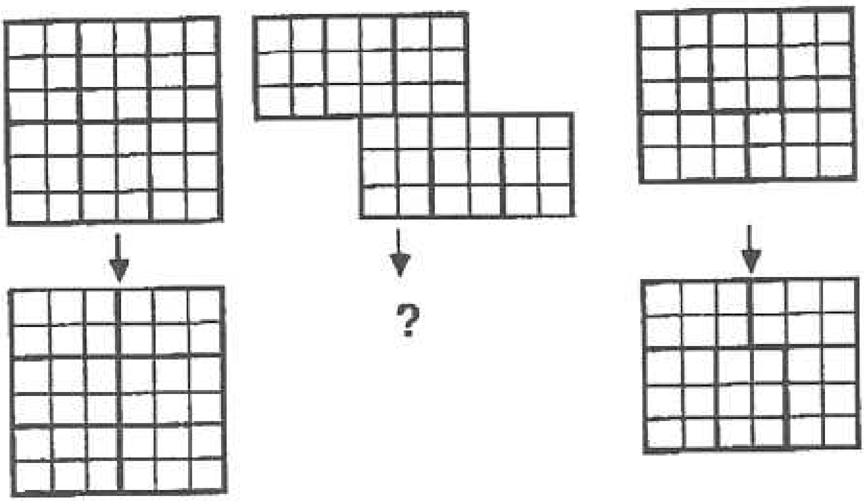

Figure 8: Pushing the highest point towards the boundary

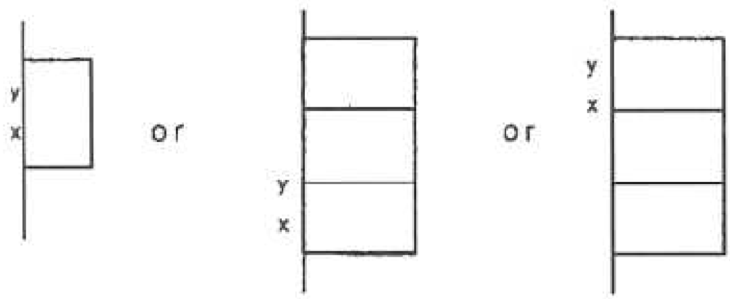

Figure 9: Putting down the first tile

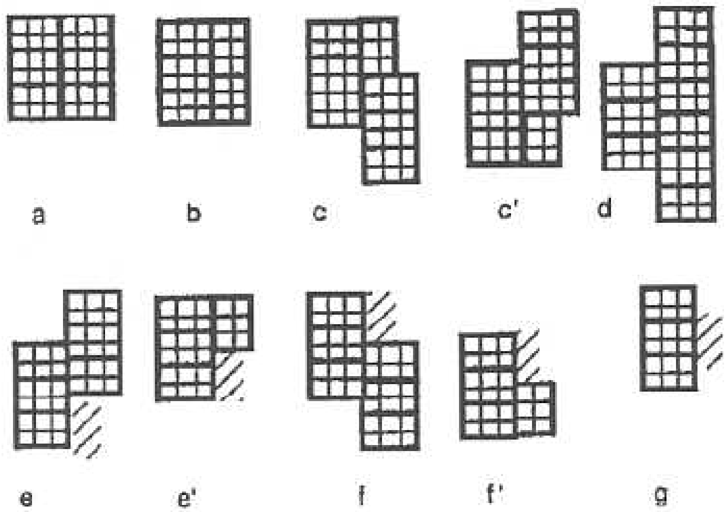

Figure 10: The right neighborhood of a block

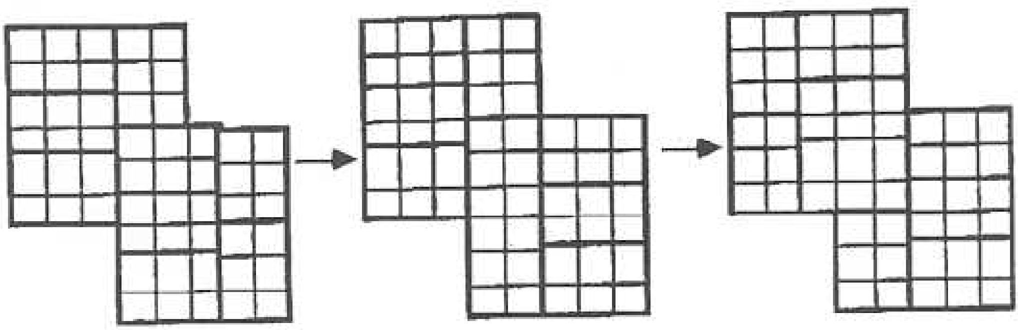

Figure 11: Pushing vertical tiles to the left

in cases c, c', e, e' and d of figure 10). We consider the set S of all blocks reachable from B_0 in the tiling.

If those blocks never touch the border ∂P outside B_0, the only cases which occur are b, c, c' and d. All the blocks farthest from B_0 have right neighborhoods as in figure 10b. Then we can "push" the vertical blocks towards the left (as in figure 11) and construct a new tiling of P such that the tile at x is vertical.

Otherwise, there must be a part S of the boundary ∂P reachable from x by a sequence of blocks, as in figure 12. We observe that in that tiling, there is a point x' of S with the same label as x in H.

Definition 2 *Given a tiling of P, a pair (x, x') of points of ∂P is* **tight** *if x and x' are both highest on ∂P, with the same label, and are linked by a sequence of blocks in the tiling.*

To have an algorithm similar to the one in the previous section, we would like to be able to say that we can always place a vertical tile at x, which would enable us to put down the first tile using local conditions only. Unfortunately tight pairs of points are a problem in this approach. However we will show that such pairs are actually independent of the tiling. In order to characterize them, we will use a different notion of height. Let $K = (a, b | a^2, b^3)$ be rooted at some arbitrary far away point r'. As we go around ∂P from c, points can also be labelled with a word $w \in K$. We define their K-height to be the distance from w to r' in the graph $\Gamma(K)$ (which in this case is just a line graph).

Lemma 4 *Given a tileable polygon P, a pair of points $x, x' \in \partial P$ is tight for some tiling of P iff the following 3 conditions are satisfied:*
1) x and x' are highest points of ∂P.
2) It is possible to construct a sequence of 3×6 rectangles (respectively, sequence of 6×3 rectangles) linking x to x', where every rectangle is to the right of (resp. below) the previous one and shifted by three, and all rectangles are inside P.
3)

$$|K\text{-}height(x) - K\text{-}height(x')| =$$

$$\begin{cases} 2|x - x'|/3 & \text{in case of figure 12} \\ 2|x - x'|/3 - 1 & \text{in case of figure 13} \end{cases}$$

where x and x' are the x-coordinates (resp. y-coordinates) of x and x' in $\mathbb{Z}^2$.

Proof idea: If (x, x') is tight for some tiling, then the three conditions of the lemma are clearly satisfied.

Conversely, given (x, x') which satisfy the three conditions of the lemma, and given a tiling of P, consider a shortest path from x to x' in P. One can prove that the path must only traverse horizontal rectangles of the tiling (otherwise the K-height would not have time to change by $2|x - x'|/3 - 1$). It is possible to show that those horizontal rectangles must be part of blocks, which form a sequence. Therefore (x, x') is tight □.

Consequence 1 *If (x, x') is tight for some tiling of P, then (x, x') is tight for every tiling of P.*

The three conditions of the lemma being independent of the tiling, can be checked even when a tiling of P is not known.

Finally, this gives us a criterion for putting down the first tile of the tiling: if the maximum x of ∂P is on a vertical side, then the first tile is always a vertical tile at x, except if there is x' such that (x, x') is tight: then the first tile at x must be horizontal (and its exact position is determined by whether the K-height increases or decreases when going from x to x').

3.4 A polynomial-time algorithm

The analysis of the previous subsection give a polynomial time algorithm for deciding whether a polygon can be tiled by 2×3 and 3×2 tiles, and giving a tiling if it exists. A high-level description follows.

Input: A polygon P with vertical and horizontal edges of integral lengths.
Initialization: Label the vertices of ∂P with words of H and of K, calculate their height and K-height.
Repeat: Take $x \in \partial P$ of maximum height. If there exists $y \in \partial P$ such that (x, y) is tight, then place a horizontal tile at x accordingly; else place a vertical tile at x appropriately.
Update P, ∂P, H-labels and heights, K-labels and K-heights.
Until: P is tiled or there is a conflict in the labels (in which case P is not tileable).

The exact running time of the algorithm depends on the implementation. With a little bit of work, one can see that there is an implementation for which the algorithm takes quadratic time.

For general $k \times l$ rectangles (with the two tile types $k \times l$ and $l \times k$), the algorithm can be generalized to yield a quadratic-time algorithm.

Theorem 2 *Given a polygon of area n, there is an algorithm for tiling the polygon with $k \times l$ and $l \times k$ tiles or proving that it is not tileable, with running time $O(n^2)$.*

We conjecture that the running time can be improved.

4 Properties of rectangle tilings

As a corollary of the above constructions we find:

Theorem 3 *Any two tilings of P by $m \times 1$ and $1 \times n$ tiles are obtained from one another by "rotations" of the form illustrated figure 14. That is, m vertical $1 \times n$ tiles can be replaced by n horizontal $m \times 1$ tiles and conversely.*

Any two tilings of P by 2×3 and 3×2 tiles are obtainable from one another by rotations of the form illustrated figure 15.

It is not hard to show that the distance in terms of the number of rotations between two tilings of P, in either of the above cases, is at most $O(n^{3/2})$ for area n. Furthermore, figure 16 shows two tilings of a $2k \times 3k$ square by 2×1 and 1×3 tiles having distance

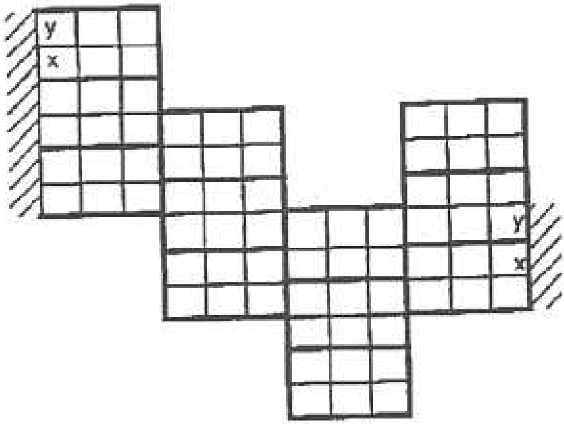

Figure 12: Reaching the boundary with a sequence of blocks

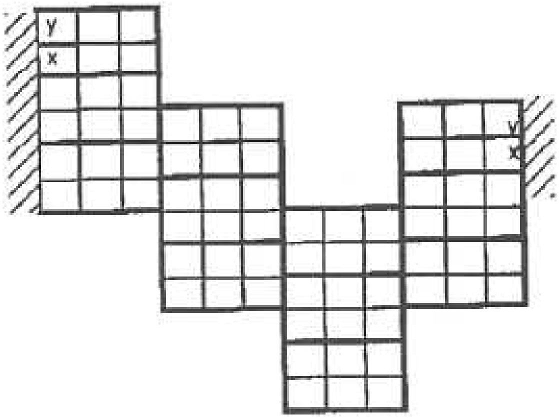

Figure 13: Chain of blocks

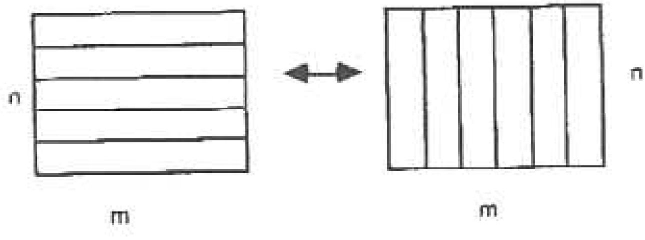

Figure 14: rotations for transforming $m \times 1$, $1 \times n$ tilings

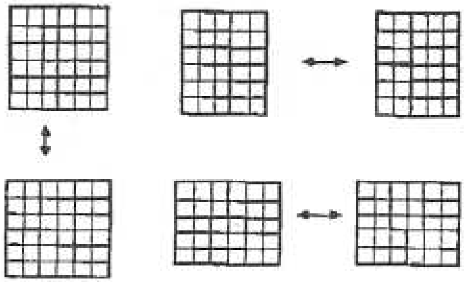

Figure 15: rotations for tranforming 3×2 tilings

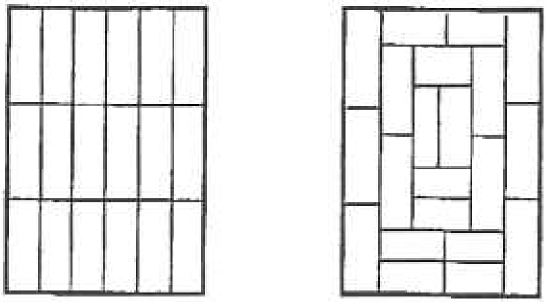

Figure 16: Two tiling which are "far apart"

$\Theta(k^3)$. This distance can be calculated by integrating the difference in height function over the tilings. Each rotation decreases the integral by a constant amount only.

References

[1] William P. Thurston, "Conway's Tiling Groups", *Am. Math. Monthly*, Oct. 1990, pp. 757-773.

[2] J.M.Robson, "Sur le pavage de figures du plan par des barres", *Actes des Journées Polyominos et Pavages*, June 1991, pp. 95-103.

[3] D.Beauquier, M.Nivat, "Pavages de pièces par h_m et v_n", *Actes des Journées Polyominos et Pavages*, June 1991, pp. 45-60.

[4] Eric Remila, "Trois algorithmes de pavages par h_m et v_n", *Actes des Journées Polyominos et Pavages*, June 1991, pp. 83-94.

[5] Fran Berman, David Johnson, Tom Leighton, Peter W. Shor, Larry Snyder, "Generalized Planar Matching", *Journal of Algorithms* 11, no 2, 1990, pp. 153-184.