

Conway's Tiling Groups

WILLIAM P. THURSTON; *Princeton University*

DR. THURSTON's numerous distinctions and honors include the Oswald Veblen Prize in Geometry, the Alan T. Waterman Award of the NSF, and the Fields Medal. His Ph.D. is from Berkeley, in 1972.



1. Introduction

John Conway discovered a technique using infinite, finitely presented groups that in a number of interesting cases resolves the question of whether a region in the plane can be tessellated by given tiles. The idea is that the tiles can be interpreted as describing relators in a group, in such a way that the plane region can be tiled, only if the group element which describes the boundary of the region is the trivial element 1.

A convenient way to describe the construction is by means of the *Cayley graph* or *graph* of a group. If G is a group, then its graph $\Gamma(G)$ with respect to generators g_1, g_2, \dots, g_n is a directed graph whose vertices are the elements of the group. For each vertex $v \in \Gamma(G)$, there will be n outgoing edges, labeled by the generators, and n incoming edges: the edge labeled g_i connects v to vg_i .

It is convenient to make a slight modification of this picture when a generator g_i has order 2. In that case, instead of drawing an arrow from v to vg_i and another arrow from vg_i back to v , we draw a single undirected edge labeled g_i . Thus, in a drawing of the graph of a group, if there are any undirected edges, it is understood that the corresponding generator has order 2.

The graph of a group is automatically homogeneous: for every element $g \in G$, the transformation $v \rightarrow gv$ is an automorphism of the graph. Every automorphism of the labeled graph has this form. This property characterizes graphs of groups: a graph whose edges are labeled by a finite set F such that there is exactly one incoming and one outgoing edge with each label at each vertex is the graph of a group if and only if it admits an automorphism taking any vertex to any other.

Whenever R is a relator for the group, that is, a word in the generators which represents 1, then if you start from $v \in \Gamma$ and trace out R , you get back to v again. If G has presentation

$$G = \langle g_1, g_2, \dots, g_n \mid R_1 = 1, R_2 = 1, \dots, R_k = 1 \rangle,$$

the graph $\Gamma(G)$ extends to a 2-complex $\Gamma^2(G)$: sew k disks at each vertex of $v \in \Gamma(G)$, one for each relator R_i , so that its boundary traces out the word R_i . An exception is made here for relations of the form $g_i^2 = 1$, since this relation is already incorporated by drawing g_i as an undirected edge. The 2-complex $\Gamma^2(G)$ is simply-connected: that is, every loop in $\Gamma^2(G)$ can be contracted to a point. In fact, if the loop is an edge path, the sequence of edges it follows describes a word in the

generators. The fact that the path returns to its starting point means that the word represents the identity. A proof that this word represents the identity by making substitutions using the relations R_i can be translated geometrically into a homotopy of the path in $\Gamma^2(G)$.

As a very simple example, the symmetric group S_3 is generated by the transpositions $a = (12)$ and $b = (23)$. They satisfy the relation $(ab)^3 = 1$. The graph is a hexagon, with undirected edges, alternately labeled a and b .

A slightly more complicated example is S_4 . It is generated by three elements $a = (12)$, $b = (23)$, and $c = (34)$. A presentation is

$$S_4 = \langle a, b, c \mid a^2 = b^2 = c^2 = 1, (ab)^3 = (bc)^3 = (ac)^2 = 1 \rangle.$$

To construct its graph, first make some copies of the ab hexagon for the S_3 subgroup generated by a and b , and similarly make some copies of bc hexagons. The subgroup generated by a and c is $Z_2 \times Z_2$, and its graph is a square with edges labeled alternately a and c . Make copies also of ac -squares. Take one copy of each polygon, and fit them together around a vertex, gluing an a edge to an a edge, etc. Around the perimeter of this figure, keep gluing on a copy of the polygon that fits. If you do this systematically, layer by layer, you will have constructed a polyhedron—it is a truncated octahedron. All the edges from the underlying octahedron are labeled b , while the squares produced by truncating the vertices are labeled $acac$.

The reader may enjoy working out the graph of the alternating group A_5 , using generators $a = (12)(34)$, and $b = (12345)$. Note that they satisfy the relations $b^5 = 1$ and $ab^3 = (135)^3 = 1$. Try kicking around the construction, with white $ababab$ hexagons and black $bbbbbb$ pentagons.




FIG. 1.1. Soccerball. A soccerball is constructed from 12 pentagons, obtained by rotating and shrinking the faces of a regular dodecahedron, together with 20 hexagons centered at the vertices of the dodecahedron.

Of course, graphs of groups don't always work out so nicely or so easily, but often, for simple presentations, they can be worked out, and they tend to have a nice geometric flavor.

2. Lozenges

We will begin with a relatively easy tiling problem. Suppose we have a plane ruled into equilateral triangles, and a certain region R bounded by a polygon π

whose edges are edges of the equilateral triangle network. When can R be tiled by figures, let us call them lozenges, formed from two adjacent equilateral triangles?




FIG. 2.1. A region tiled by lozenges. A portion of an equilateral triangle subdivision of the plane, tiled by lozenges.

To analyze this problem, we first establish a labeling convention. We arrange the triangulation of the plane so that one set of edges is parallel to the x -axis, or at 0° . Label these directed edges a , label b the directed edges pointing at 120° , and c the edges pointing at 240° . This labeling is homogeneous, so it is the graph of a group A . We can read off relators for A by tracing out the boundary curves of triangles: A satisfies $abc = 1$ and $cba = 1$. If desired, the first relation could be used to eliminate c ; the second relation then says that $ba = ab$. The group A is $\mathbb{Z} + \mathbb{Z}$, as we could have seen anyway by its action on the plane.

The shape of the polygon π is determined by the sequence of edges it traces out; this is a word in the generators a, b, c of A . Rather than thinking of it as a word, we prefer to think of it as an element $\alpha(\pi)$ in the free group F with generators a, b, c . The fact that π closes up is equivalent to the condition that the homomorphism $F \rightarrow A$ send $\alpha(\pi)$ to the identity.

If a lozenge is placed in the triangular network, its boundary can be traced by one of three elements, depending on its orientation: that element is either $L_1 = aba^{-1}b^{-1}$, $L_2 = bcb^{-1}c^{-1}$, or $L_3 = cac^{-1}a^{-1}$. The precise word depends on the starting point on the boundary of the lozenge, but starting from a different vertex only changes the word by a circular permutation; the two choices give conjugate elements of F . The lozenge group L is defined by these relators, that is

$$L = \langle a, b, c \mid L_1 = L_2 = L_3 = 1 \rangle.$$

Actually, the three relations say that the three generators commute with each other, so that $L = \mathbb{Z}^3$.

We claim that if the region R can be tiled by lozenges, then the image $I(\pi)$ of $\alpha(\pi)$ in L must be trivial. In fact, suppose that we have such a tiling. If R consists of a single tile, the claim is immediate. Otherwise, find a simple arc in R which cuts R into two tiled subregions R_1 and R_2 . By induction, we may assume that $I(\pi_1)$ and $I(\pi_2)$ are both trivial, where π_i is a polygonal curve tracing around ∂R_i . But $I(\pi) = I(\pi_1) * I(\pi_2)$, so $I(\pi)$ is also trivial.




FIG. 2.2. Three-dimensional interpretation of lozenge tiling. If a region R can be tiled by lozenges, then the lozenge pattern lifts to the 2-skeleton of a cubical tiling of R^3 , oriented diagonally to the plane of the lozenges.

There is a very direct geometric interpretation: think of the graph $\Gamma(L)$ as the 1-skeleton of a cubical tessellation of space, oriented so that cubes are on their corners: more precisely, so that the two endpoints of any path labeled abc are on the same vertical line. The 2-complex $\Gamma^2(L)$ is the union of the faces of the cubes. A lozenge in the plane is the orthogonal projection of a square face of a cube. Given a path π in the plane, arrange it (for notational purposes only) so the base point $*$ lies below the base point 1 of $\Gamma(L)$. Lift it edge by edge to a path in $\Gamma(L)$. When you make a complete circuit around π , you may or may not come back to the starting point in $\Gamma(L)$. The invariant $I(\pi) \in L$ is the ending vertex. This invariant of necessity lies in the kernel of the map $L \rightarrow A$, which is isomorphic to \mathbb{Z} : it can be described simply as the net rise in height.

If R can be tiled by lozenges, the tiling itself can be lifted, tile by tile, into $\Gamma^2(L)$, that is, into the 2-skeleton of the cubical tessellation. This gives another proof that the invariant $I(\pi)$ must be 1 if R can be tiled. In fact, if you look at a tiling by lozenges, you can imagine it so that it springs out at you in a three-dimensional picture.




FIG. 2.3. Nontilable region. The region in the plane enclosed by the polygonal curve cannot be tiled by lozenges, since when it is lifted to the cubic network, it fails to close.

Algebraically, given the word representing π , the net rise in height is simply the sum of the exponents. The condition is that π heads at a bearing of 0° , 120° , or 240° the same length of time it heads at a bearing of 60° , 180° , or 300° .

This condition can be seen in an alternative way using a coloring argument. The triangles in the plane have an alternating coloring, with abc triangles colored white and cba triangles colored black. Each lozenge covers one triangle of each color—therefore, if R can be tiled, the number of white triangles must equal the number of black triangles. The difference in fact can be shown to be the net rise in height of α , as measured in main diagonals of cubes. The coloring consideration really gives a more elementary derivation that $I(\pi)$ must vanish for a tiling to be possible. However, this and related coloring arguments in general cannot give as much information as $I(\pi)$. One way to think of it is that coloring arguments are the abelian part of the group theory. If the group is abelian as in the present case, or more generally if the subgroup consisting of invariants $I(\pi)$ for *closed* paths is abelian, then that information is sufficient.

The algebraic condition that $I(\pi) = 1$, is not sufficient to guarantee a tiling by lozenges. There are curves π which go around nearly a full circle, with the lift in $\Gamma(L)$ rising considerably, and then instead of closing, they circle around another loop which brings them down to the starting height. If R could be tiled by lozenges, it could be divided into two regions by a fairly short path along edges of lozenges: but the rise in height for one side would be forced to be still positive, which would be a contradiction. We will return later to give a necessary and sufficient condition for a tiling by lozenges, along with a formula for a tiling if such exists.




FIG. 2.4. Potentially tileable region. The boundary curve of this region lifts to a closed curve, so it meets the group-theoretic tiling condition. An actual tiling will be shown in 4.1, High lozenge tiling.

3. Tribone Tilings

Here is another example, for which other methods seem inadequate. I first heard this problem in an electronic mail inquiry from Carl W. Lee (ms.uky.edu!lee) in Kentucky.

Last semester, a number of us here became interested in a combinatorial problem that was making the rounds. I'm sure you already have heard of it, and we heard a rumor that John Conway had solved it. It concerned a triangular array of dots. The problem was to pack in as many segments as possible, where each segment covered three adjacent dots in one of the three directions, and no two segments were allowed to touch. Is there any size configuration that admits a packing such that each dot is covered? Do you know anything about the status of this problem? Thanks in advance.




FIG. 3.1. Triangle of hexagons. A triangular array of hexagons, eight on a side. Can this be tiled by tribones?

I hadn't heard of it, but I asked Conway about it. We sat down together, and he worked it out.

This question can be alternately formulated in terms of a triangular array of hexagons. The problem is to show that one cannot tessellate the region using tiles made of three hexagons hooked linearly together. More generally, one can ask for the minimum number of holes left in an attempt to tile the region by these tiles.

If the region has side length n , then the number of hexagons is $n(n+1)/2$. A first, necessary condition is that n or $n+1$ is divisible by 3, that is, n is congruent to 0 or 2 mod 3. Note that if it is ever possible to solve the problem when n is congruent to 2 mod 3, one can extend the solution by adding a row of tiles along one side, to derive a solution for $n+1$.

Label each side in the hexagonal grid with an a , b , or c , according to the direction of the edge: a if it is parallel to the x axis, b if the angle from the x -axis to the edge (measured counterclockwise) is 60° , and c if this angle is 120° . Thus, the sides of every hexagon are labeled $abcabc$.

This labeling gives the 1-skeleton of the grid the structure of a group graph, where the group is

$$A = \langle a, b, c | a^2 = b^2 = c^2 = (abc)^2 = 1 \rangle.$$

The group is a group of isometries of the plane, generated by 180° revolutions about the centers of the edges; it also contains the 180° revolutions about the centers of the hexagons. The group A is sometimes called the $(2, 2, 2, 2)$ -group.

A path π in the 1-skeleton of the hexagonal grid now is determined by a word in the generators of A . We prefer to think of this in a slightly different way: π determines an element $\alpha(\pi)$ in the free product $F = \mathbb{Z}_2 * \mathbb{Z}_2 * \mathbb{Z}_2$. We are

particularly interested in closed paths, that is, elements of the kernel of $F \rightarrow A$. Unfortunately, this kernel is infinitely generated: it is a free group whose generators are given by arbitrary paths p_1 , followed by a circuit around one of the three hexagons at the endpoint of p_1 , followed by the p_1^{-1} .




FIG. 3.2. Tribones in three orientations. There are three possible orientations for a tribone, in an array of hexagons. With our labeling convention, they are labeled in three different ways.

The standard tile, let us call it a *tribone*, can be laid in the plane in three different orientations. Circuits around the tribones in these three orientations trace out the elements

$$T_1 = (ab)^3 c (ab)^3 c$$

$$T_2 = (bc)^3 a (bc)^3 a$$

$$T_3 = (ca)^3 b (ca)^3 b.$$

If π is a simple closed circuit in the plane such that the region R bounded by π can be tiled by these tribones, then the image $I(\pi)$ of $\alpha(\pi)$ in the tribone group

$$T = \langle a, b, c | a^2 = b^2 = c^2 = T_1 = T_2 = T_3 = 1 \rangle$$

must be trivial.

The relation T_1 says that c conjugates $(ab)^3$ to its inverse. Observe that a and b also conjugate $(ab)^3$ to its inverse—in fact, this is already true in F . In other words, $(ab)^3$ generates a normal subgroup, and it commutes with every word of even length. Similarly, $(bc)^3$ and $(ca)^3$ generate normal subgroups. Together, the three elements generate a normal abelian subgroup J of T .

To form a picture of T , let us first look at the quotient group $T_0 = T/J = \langle a, b, c | a^2 = b^2 = c^2 = (ab)^3 = (bc)^3 = (ca)^3 = 1 \rangle$. The graph of T_0 can readily be constructed: take an infinite collection of three types of hexagons, with their edges labeled by the relations C_1, C_2 , and C_3 . These glue together to form a hexagonal pattern in the plane, where each vertex has one a edge, one b edge, and




FIG. 3.3. Second hexagonal group. The group T_0 also has a graph isomorphic to the edges of a hexagonal tiling of the plane.

one c edge incident to it. The group T_0 acts faithfully as a group of isometries of the plane, generated by reflections in the edges of this hexagonal tiling: it is a triangle group. It is curious that even though the groups A and T_0 and the labeled graphs $\Gamma(A)$ and $\Gamma(T_0)$ are different, when the labels are stripped they become isomorphic.

If the region R can be tiled by tribones, then $\alpha(\pi)$ must map to the trivial element of T , so it maps to the trivial element of T_0 . In our case, the region is a triangular array of hexagons, and its boundary can be taken as $\alpha(\pi) = (ab)^n(ca)^n(bc)^n$.

Obviously, if n is a multiple of 3, the image $I(\pi)/J$ in T_0 is trivial. In the other case, that n is 2 more than a multiple of 3, it is also trivial. This is easily seen by tracing out the curve in our array of hexagons, or by noticing that one can add additional tribones along one edge to form a triangular region with side length $n+1$, which is a multiple of 3. Since we have pushed π only across tribones, $I(\pi)$ is the same for the two cases.

Since T_0 was not sufficient to detect the nontriviality of $I(\pi)$, we need to finish our job, and build a picture of T . First, look at the path in the graph of T_0 determined by the element T_1 . Start at a vertex $*$ where the circuit $C_1 = ababab$ goes counterclockwise around a hexagon. Then T_1 goes counterclockwise around this hexagon, then along the c edge, clockwise around the C_1 hexagon through that vertex, and back along the c edge to close. In particular, the signed total of C_1 -hexagons enclosed (counted according to degree of winding with counterclockwise circuits counted positively), is 0.

It is not hard to describe now the full group T , which is an extension of the form $J = \mathbb{Z}^3 \rightarrow T \rightarrow T_0$. We can interpret an element of T to be a vertex v in the graph of T_0 , together with a path p from $*$ to v , subject to the equivalence relation that




FIG. 3.4. Alternate image of tribone. By construction, the tribone relations are satisfied in the groups T and hence $T_0 = T/J$. This is the image of one of the tribone relators in the graph of the group T_0 . Note how it encircles two AB -hexagons, once clockwise and once counterclockwise.

if q is another path from $*$ to v , then $p \sim q$ if the signed totals of C_1 , C_2 , and C_3 hexagons are all 0. (Of course, if we pick one path such as p from $*$ to v , then other paths from $*$ to v are determined by three arbitrary integers, which specify these signed totals.) With this definition, the relations T_i are obviously satisfied, hence the group so constructed is at least a quotient group of T . But we have already seen that the kernel J of the map $T \rightarrow T_0$ is abelian, and generated by C_i . In the construction, this kernel is the free abelian group on the C_i , so it must in fact give T .

Once we know T , we can read $I(\pi)$ by inspection. As we saw, it suffices to consider the case $n = 3k$; the invariant is $C_1^k C_2^k C_3^k$, which is obviously not 1, so the tiling is impossible.

One can ask whether this method gives a lower bound on the number of holes one is forced to leave, in a partial tiling of R by tribones. To study this question, we should examine the subgroup K of T generated by elements of the form $I(\gamma)$, where γ is a path in the graph of A going from $*$ to some point v , circumnavigating a hexagon, and returning. In other words, K is the kernel of the map $T \rightarrow A$. Note that $\alpha(\gamma)$ has the form $gabcabcg^{-1}$, where g is arbitrary. In the group T_0 , $abcabc$ acts as a translation. The conjugates of $abcabc$ in T_0 are translations in three different directions spaced at 120° angles, and the subgroup they generate is isomorphic to \mathbb{Z}^2 . In K , there are actually an infinite number of different conjugates of $abcabc$: if g acts as a translation in T_0 , then the commutator $gabcabcg^{-1}cbacba$ is trivial in T_0 , but it might not be trivial in T : this path may enclose an arbitrary number m of hexagons of type C_1 , and an equal number of type C_2 and C_3 .

The subgroup K is therefore a nilpotent group, generated by $s = abcabc$, $t = bcabca$, and $u = C_1 C_2 C_3$, with presentation

$$K = \langle s, t, u \mid [s, u] = [t, u] = 1, [s, t] = u^3 \rangle.$$




FIG. 3.5. Alternate image of a triangle. The triangle word $(ab)^n(ca)^n(bc)^n$ of size $n = 3m$ or $n = 3m + 2$ maps to the trivial element in T_0 . In the diagram above, if $n = 3m$, trace the word starting at the center. If $n = 3m + 2$, start b from the center.

It is easy to check that every element of K is realized as $I(\pi)$, for some simple closed curve π in the plane.

Even though the invariants associated with triangular regions take larger and larger values in I , this does not give any information limiting the number of holes: for instance, three holes $g_1 abcabcg_1^{-1}$ can yield u^k , for arbitrarily high k . In fact, it is possible to tessellate the triangular region of size n with tribones except for 1 hole, if $n \equiv 1(3)$, by placing the hole exactly in the middle, and then arranging concentric triangular layers of tribones around this hole. From these examples, tribone tilings with 3 holes are easily constructed when $n \equiv 0(3)$ or $2(3)$. It does give some information, however: in the case that $n \equiv 2(3)$ or $n \equiv 0(3)$, the conjugacy class changes ("increases") with n , which implies that the length of the minimum closed loop enclosing all the holes has to go to infinity with n . In the case $n \equiv 1(3)$, the conjugacy class of $I(\pi)$ is constant—since the region can always be tiled with a single hexagon missing, $I(\pi)$ is conjugate to $abcabc$. However, the actual word changes with n , which implies that the missing hole cannot be too close to the boundary. Perhaps a careful analysis would show that if there is a single hole, it must be exactly in the center of the triangle.

4. Dominoes and Lozenges Revisited

Conway's tiling groups are quite versatile, provided you can work out the group determined by the tiles. Even when (or perhaps especially when) the invariant $I(\pi)$ gives no information which could not have been easily obtained by other means, the geometric picture of the graph of the group can sometimes be exploited to give not just an algebraic criterion, but a precise geometric criterion for the existence of a tiling.

When G is a tiling group (with presentation given by a set of tiles), we define a measure of area in $\Gamma^2(G)$ to be the area defined by projection to the plane: the area of a 2-cell is the area of a corresponding tile. When the algebraic invariant $I(\pi)$ is 1, the curve π bounding R lifts to a closed $\tilde{\pi}$ in $\Gamma(G)$. We can ask, what is

the minimum area of a surface S in $\Gamma^2(G)$ with boundary $\bar{\pi}$? This area is necessarily at least as great as the area of R . If it is equal, then the images of the 2-cells of S must be disjoint, so that they form a tiling of R . There are several approaches which are sometimes successful for calculating this minimal area, but there is one particular situation when there is a really definitive solution: when $\Gamma^2(G)$ can be enlarged, by adding 3-cells, to make a contractable 3-manifold. In this situation, there is a "max flow min cut" principle which guarantees an efficient algorithm for finding a minimal surface.

Rather than going on with the general theory, we will illustrate this with two examples. First we revisit the lozenge question.

If R is a union of triangles in the plane, and if v and w are vertices in R , possibly on the boundary, define $d(v, w)$ to be the minimum length of a positively directed edge-path in R (possibly going on the boundary) joining v to w . This "distance" function d is not symmetric, since we cannot simply reverse an edge path. Any closed positively directed edge path has length a multiple of 3, so the $d(v, w)$ is defined modulo 3 independent of path. The three vertices of a triangle take the three distinct values modulo 3. If R is connected, it is always possible to find at least one positively directed path from v to w , so $d(v, w)$ is well-defined.

Consider the lifting of any tiling of R by lozenges to the cubical network, $\Gamma^2(L)$. This is determined by a height function $h(v)$ for the vertices v . We can choose the vertical scale so that h is integer-valued, and each edge of a lifted lozenge increases in height by 1; the edge of the triangular network covered by the lozenge lifts to a diagonal of a square, and decreases in height by 2. It follows that $h(w) - h(v) \geq d(v, w)$.

The boundary path π determines a unique height function h on its vertices, up to constants. This gives a necessary condition that R can be tiled: for any two vertices v and w on π , $h(w) - h(v) \geq d(v, w)$.

If π satisfies this necessary condition, then there is a unique maximally high lozenge tiling: define

$$h(x) = \min_{v \in \pi} \{d(v, x)\}.$$

To produce the actual tiling, place a lozenge so as to cover an edge where the height changes by 2. Since the three vertices of a triangle take distinct values modulo 3, and since h increases by at most 1 along any edge, each triangle has exactly one edge where h changes by 2; therefore, the collection of lozenges is a tiling.




FIG. 4.1. High lozenge tiling. The "highest" lozenge tiling compatible with the boundary curve.

There is a simple algorithm for quickly computing h , and the tiling: rather than spell it out, we will describe the analogous algorithm for dominoes.

A closed path π in a square grid can be described by an element $\alpha(\pi)$ of the free group $F(x, y)$, which maps to the trivial element of the $A = \mathbb{Z}^2$. If the region R bounded by π can be filled with dominoes, then the image $I(\pi)$ of $\alpha(\pi)$ in the domino group $G = \langle x, y | xy^2 = y^2x, yx^2 = x^2y \rangle$ must be trivial.

What does the graph of G look like? We can construct a picture in \mathbb{R}^3 , as follows. Fill the xy -plane with a black and white checkerboard pattern. Above the black square $[0, 1] \times [0, 1]$, construct a right-handed helix, joining $(0, 0, 0)$ by a line segment to $(0, 1, 1)$, to $(1, 1, 2)$, $(0, 1, 3)$, $(0, 0, 4)$, and so on: the x and y coordinates here marching forever around the boundary of the square, while the z coordinate increases by 1 each move. Similarly, $(0, 0, 0)$ is connected to $(0, 1, -1)$, etc. Construct a similar helix above each black square. Label each edge x or y , according to its image in the plane. Note that this creates left-handed helices above the white squares. The boundary of any domino in the plane lifts to a closed path in this graph we have constructed. Since the graph has a simply-transitive group of isometries, it is the graph of a group. Since it satisfies the domino relations, it is at least a quotient group of the domino group G . It is not hard (and strictly speaking, it is not logically necessary) to verify that this graph is indeed the graph of G .

The curve π lifts to a curve $\tilde{\pi}$ in the graph of G . A convenient way to denote this, in the plane, is to record the height of the lift next to each vertex of π in the




FIG. 4.2. The domino group. The graph of the domino group is a union of square helices over the squares of a checkerboard, alternating in handedness. A domino anywhere in the plane lifts to this graph, starting at any point. This illustration shows two coils of four neighboring helices.




FIG. 4.3. Domino tiling. A tiling by 9 dominoes, lifted to the graph of the domino group.

plane. The rule is simple: one can start with 0 at some arbitrary vertex. Along any edge of π which has a black square to its left, the height increases by 1. Along any edge with a white square to its left, the height decreases by 1. A necessary condition that R can be filled with dominoes is that the height after traversing once around the curve is 0.




FIG. 4.4. Domino roof. This is the tiling which the algorithm yields, when applied to a 16×16 square grid. This is the tiling which has the highest lifting to the graph of the domino group of any tiling by dominoes.

There is a criterion and construction for a domino tiling, analogous to the construction for lozenges. Here is how the formula can be worked out, on a sheet of grid paper. Begin, as above, by labeling the height of each vertex of π . The heights consist of the integers in some interval, $[n, m]$. We will construct a height function on all vertices of R , beginning with $n + 1$, and working up. Suppose, inductively, that we have finished with all vertices of height less than or equal to k . For each vertex v of height k , and for each edge e leading from v which has a black square on its left, consider the second endpoint w of e . If the height of w has been previously defined, and if it is not greater than $k + 1$ leave it as is. If the height is defined and greater than $k + 1$, then a domino tiling is impossible: give up. Otherwise, define the height of w to be $k + 1$.

If this procedure reaches a successful conclusion, each edge of R has a difference of heights of its two endpoints of either 1 or 3. (Note that the height modulo 4 is determined by the point in the plane.) Erase all the edges whose endpoints have a difference of height of 3. What is left is a picture of a tiling by dominoes.

5. Triangles

Here is a related sequence of tiling problems which are resistant to direct attempts at general solution, but translate nicely into the realm of group theory.

Consider, again, a triangular array of dots, with N dots on each side. Is it possible to subdivide this array into disjoint triangular arrays of dots with M on each side? We suggest the reader indulge in experimentation with a few cases,




FIG. 4.5. Domino bubble. This illustration shows both the highest and the lowest tiling by dominoes of a standard checkerboard. They are isomorphic, differing only by a 90° rotation of the checkerboard (interchanging colors). The upper tiling is shown in the upper plane as well as the upper surface of the bubble, the lower tiling in the lower plane and the lower surface of the bubble. The bubble they form encloses the lift of any tiling by dominoes. Possible tilings are "like" Lipschitz functions in the square with Lipschitz constant 1, as measured in the Manhattan metric. The limits of domino tilings, lifted to the graph of the group, as the grid size goes to zero, are exactly such Lipschitz functions.

before reading further. For example, the cases $M = 2$ with N ranging from 2 to 12 are interesting.

...

As in the case of the tribones, this translates into a tiling problem: given a triangular array of hexagons with N hexagons per side, can one tile it by tiles T_M which are triangular arrays of hexagons M per side? We can express this with notation as in the case of tribones: label the edges of the underlying hexagonal tiling by a 's, b 's and c 's. Given a path π in the plane, it is described by an element $\alpha(\pi)$ of $F = \langle a, b, c | a^2 = b^2 = c^2 = 1 \rangle$. If the region R bounded by π can be tiled by the copies of T_M , then the image $I(\pi)$ of $\alpha(\pi)$ is trivial in the group

$$G_M = \langle a, b, c | a^2 = b^2 = c^2 = 1, t_M = 1 \rangle,$$

where t_M represents the boundary curve of the tile T_M ,

$$t_M = (ab)^M (ca)^M (bc)^M.$$

A parallelogram of hexagons with M hexagons on one side and $M + 1$ on the other can be tiled by two copies of T_M . This implies that $(ab)^M$ commutes with $(bc)^{M+1}$ and with $(ca)^{M+1}$, and so forth.

These relations imply that $(ab)^M$ commutes with $(bc)^{M(M+1)}$, and they also imply that $(ab)^{M+1}$ commutes with $(bc)^{M(M+1)}$. Combining these two facts, it follows that (ab) commutes with $(bc)^{M(M+1)}$. Geometrically, one can tile an $M \times M(M+1)$ parallelogram and an $M+1 \times M(M+1)$ parallelogram. Their difference is a $1 \times M(M+1)$ parallelogram: this can be tiled in a certain algebraic sense as the difference of the two.

It will simplify the picture at this point if we pass to the subgroups F^e and G_M^e generated by words of even length. Since all relations have even length, the wordlength modulo 2 describes a homomorphism of F and G_M to \mathbb{Z}_2 , and these subgroups have index 2. The group F^e is the free group on 2 generators, but a more symmetric description is

$$F^e = \langle x, y, z | xyz = 1 \rangle,$$

where $x = ab$, $y = bc$, and $z = ca$. A presentation for the group G_M^e is obtained by adjoining relations coming from t_M to F^e : it requires two relations, one obtained by transcribing t_M directly, and the other transcribing the conjugate of t_M by an element of odd length. Using $t_M = 1$ and $bt_Mb = 1$, we obtain

$$G_M^e = \langle x, y, z \mid xyz = 1, x^My^Mz^M = 1, x^{-(M+1)}y^{-(M+1)}z^{-(M+1)} = 1 \rangle.$$

G_M^e has an interesting alternate generating set: $X = x^M$, $X' = x^{-(M+1)}$, together with Y, Y', Z and Z' defined similarly, clearly generate. We have already seen that X, Y , and Z commute with X', Y' , and Z' .

The elements $s = X^{M+1}$, $t = Y^{M+1}$, and $u = Z^{M+1}$ commute with everything in G_M^e , so they generate a central subgroup J which is Z^3 or a quotient. Let $G_M^0 = G_M^e/J$. We will analyze the structure of G_M^0 , and from that construct G_M^e .

In G_M^0 , X, Y , and Z satisfy relations

$$XYZ = 1, \quad X^{M+1} = Y^{M+1} = Z^{M+1} = 1.$$

These relations describe the orientation-preserving $(M+1, M+1, M+1)$ triangle group, which acts as a discrete group of isometries on the Euclidean plane if $M=2$ and on the hyperbolic plane if $M>2$. We have not checked that these generate *all* the relations on X, Y , and Z , but we immediately deduce that the subgroup H of G_M^0 generated by X, Y , and Z is a quotient of this triangle group. But there is a homomorphism f of the original group G_M to the full triangle group (including reflections), defined by sending a, b , and c to reflections in the sides of a $\pi/(M+1), \pi/(M+1), \pi/(M+1)$ triangle. The relation $t_M = 1$ is satisfied, since in this group $(ab)^M = ba$ so that $(ab)^M(ca)^M(bc)^M = (ba)(ac)(cb) = 1$. Note that f sends X to ba , Y to ac and Z to cb , that is, to the standard generators of the $(M+1, M+1, M+1)$ triangle group, and it sends s, t , and u to 1. Therefore, H is isomorphic to the orientable $(M+1, M+1, M+1)$ triangle group.

A similar analysis shows that the subgroup H' generated by X', Y' and Z' is the orientable (M, M, M) triangle group. This group acts on the sphere, the Euclidean plane, or the hyperbolic plane when $M=2, M=3$, or $M \geq 4$. The analogous homomorphism f' maps G_M to the full (M, M, M) triangle group, mapping a, b , and c to the standard generators.

The two subgroups H and H' intersect (as seen from the effects of f and f') they generate G_M^0 , and they commute with each other. Therefore, G_M^0 is the product $H \times H'$ of the two triangle groups.

Now we need to determine the kernel J of the quotient $G_M^e \rightarrow G_M^0$, and the structure of the central extension. As in the tribone case, we can do this geometrically, in terms of areas enclosed by curves. The graph Γ of the full $(M+1, M+1, M+1)$ triangle group is formed from copies of three kinds of $2(M+1)$ -gons, with perimeters labeled $(ab)^M$, $(ca)^M$ and $(bc)^M$, with one of each kind meeting at each vertex. Arrange the orientation so that 1 is an "even" vertex that is, the a, b , and c edges emanating from 1 are in counterclockwise order. Then the relation t_M based at v encloses positively one copy of each type of polygon, while the conjugate bt_Mb encloses negatively one copy of each type of polygon.

Similarly, the graph Γ' of the full (M, M, M) triangle group is made from three kinds of $2M$ -gons. Starting at the 1, which we suppose is an even vertex, the relation t_M encloses positively one copy of each type of polygon, while bt_Mb encloses negatively one copy of each. However, in the case $M=2$, the entire graph is finite: it is the 1-skeleton of a cube, and the number of polygons enclosed by a curve is well-defined only modulo 2.

First let's deal with the case $M > 2$. We can define an extension K of G_M^0 as an equivalence relation on elements of F^e , as follows. An element g of F^e determines paths $p(g)$ in Γ and $p'(g)$ in Γ' . We define g to be equivalent to h if $p(g)$ ends at the same point as $p(h)$, $p'(g)$ ends at the same point as $p'(h)$, and if the closed loop $p(g)p^{-1}(h)$ encloses the same numbers of ab -polygons, bc -polygons, and ca -polygons as $p'(g)p'^{-1}(h)$.

In particular, an element of the kernel of the map of K to $H \times H'$ maps to closed loops in both pictures, and is determined by the triple of differences of the number of polygons enclosed. The elements s , t and u map to $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$. It follows that $K = G_M^e$, and $J = \mathbb{Z}^3$ (provided $M > 2$).

The boundary of the size N triangle T_N can be described by the element $t_N = (ab)^N(ca)^N(bc)^N$. The path $p(t_N)$ in Γ closes only when N is 0 or $-1 \pmod{M+1}$, while the path $p'(t_N)$ closes only when N is 0 or $-1 \pmod{M}$. Since M and $M+1$ are relatively prime, there are four solutions modulo $M(M+1)$: $0, M, M^2-1, -1$. For values of N satisfying one of these congruence conditions, the invariant in G_M^0 is 0, so the invariant is in J ; it is a positive multiple of $(1, 1, 1)$ in all but the trivial case $N = M$.

THEOREM (Conway). *When $N > M > 2$, the triangular array T_N of hexagons cannot be tiled by T_M 's.*

This analysis has an interesting variation case $M = 2$. Given two elements g and h of F^e , we can define them to be equivalent if $p(g)$ and $p(h)$ have the same endpoints, $p'(g)$ and $p'(h)$ have the same endpoints, and if the numbers of polygons of the three types enclosed by the path $p(g)p(h)^{-1}$ is a multiple k of $(1, 1, 1)$ which has the same parity as the number of polygons enclosed by $p'(g)p'(h)^{-1}$. This defines a central extension of $H \times H'$ by \mathbb{Z}^3 modulo the subgroup generated by $s^2t^2u^2 = 1$. To justify that this group is in fact G_2^e , we must prove that $s^2t^2u^2 = (ab)^{12}(ca)^{12}(bc)^{12} = 1$ in this group, or even better, that it is possible to tile T_{12} . Such a tiling can be found fairly easily—see FIGURE 5.1, the 12-stack by 2-stacks.




FIG. 5.1. The 12-stack by 2-stacks. The triangle T_{12} can be tiled by T_2 's.

The computation of the mod 2 invariant for tilings by T_2 's can be rather annoying when done directly. However, there is a neat trick, which enables one to see this invariant geometrically: most regions which have a multiple of 3 hexagons can be tiled easily by T_2 's along with tribones. The boundary $abababcabababc$ of a tribone maps to closed paths in both Γ and Γ' . In Γ , it encloses a net of 0 of each type of hexagon, as we saw before. In Γ' , this curve winds counterclockwise 1.5 revolutions about an ab -face of the cube, goes down a c -edge to the opposite face, winds 1.5 revolutions counterclockwise (with respect to the orientation of the square), and goes up again to close. It is therefore equivalent, in terms of which kinds of squares it encloses, to $abcabc$, which is an odd multiple of $(1, 1, 1)$.

Therefore, if a region can be tiled with a collection of T_2 's together with an odd number of tribones, it cannot be tiled with T_2 's. For $0 < N < 12$, only for the values 2, 3, 5, 6, 8, 9, 11 is the number of tiles a multiple of 3. One quickly finds that in the cases T_3 , T_5 , T_6 and T_8 there is a tiling by one tribone and the rest T_2 's, while T_2 , T_9 , and T_{11} can be tiled.

Given any tiling or partial tiling of T_k , with $k > 1$, it can be extended to a tiling or partial tiling of T_{k+12} by adding a $12 \times k$ parallelogram, together with a T_{12} . The $12 \times k$ parallelogram can be tiled by subdividing into 2×6 and 3×6 parallelograms.

THEOREM (Conway). *A triangular array T_k of hexagons can be tiled by T_2 's if and only if k is congruent to 0, 2, 9, or 11 modulo 12.*

REFERENCE

[CoLa] J. H. Conway and J. C. Lagarias, *Tiling with Polyominoes and Combinatorial Group Theory*.