

Spaces of domino tilings

*Nicolau C. Saldanha, Carlos Tomei,
Mario A. Casarin Jr., Domingos Romualdo*

Abstract: We consider the set of all tilings by dominoes (2×1 rectangles) of a surface, possibly with boundary, consisting of unit squares. Convert this set into a graph by joining two tilings by an edge if they differ by a *flip*, i.e., a 90° rotation of a pair of side-by-side dominoes. We give a criterion to decide if two tilings are in the same connected component, a simple formula for distances and a method to construct geodesics in this graph. For simply connected surfaces, the graph is connected. By naturally adjoining to this graph higher dimensional cells, we obtain a CW-complex whose connected components are homotopically equivalent to points or circles. As a consequence, for any region different from a torus or Klein bottle, all geodesics with common endpoints are equivalent in the following sense. Build a graph whose vertices are these geodesics, adjacent if they differ only by the order of two flips on disjoint squares: this graph is connected.

0. Introduction

In this paper we consider tilings of a region consisting of unit squares by *dominoes*, i.e., pairs of adjacent squares. Tilings of a rectangle of integral sides were counted by Kasteleyn ([5]). More recently, Lieb and Loss ([6]) showed how to count tilings of general regions by making use of determinants. Conway and Lagarias ([1]) studied the problem of tiling a subset of \mathbb{R}^2 with a given set of tiles, by group-theoretical techniques. Thurston ([9]) adapted these techniques to study domino tilings, producing a necessary and sufficient condition for a simply connected region of the plane to be tileable by dominoes.

We are interested in studying T , the set of all possible tilings of a fixed region. Given a tiling, we perform a *flip* by lifting two dominoes and placing them back in a different position: clearly, the two dominoes must form a square of side 2. Two tilings are *adjacent* in T if we move from one to the other by a flip. Turn T into a graph by joining adjacent tilings by edges and define connected components of T and distance between tilings in the usual way. We obtain a very operational criterion (the equivalent Theorems 1.1, 1.2 and 3.1) for two tilings to belong to the same connected component of T ; as a corollary, if the region is simply connected, T is connected. Our techniques provide us with a fair understanding of the combinatorial, topological and metric structure of T : thus, for example, each connected component of T is a lattice and we describe in Theorem 3.2 a simple formula for the distance between tilings and a characterization of shortest routes between points. In a sense to be detailed in Section 3, all such routes are equivalent: a topological version of this statement is that T induces naturally a CW-complex whose connected components are contractible (Theorem 3.4). More generally, we consider quadriculated surfaces (defined in Section 4) and obtain analogous results to Theorems 3.1, 3.2 (Theorem 4.1) and 3.4 (Theorem 4.3).

Acknowledgements: The first two authors receive support from SCT and CNPq, Brazil. The other two were supported by a grant for undergraduates by CNPq. We would like

to thank one referee for a very careful reading of the manuscript, often leading to a much clearer text.

1. Connected components of T

Let A be a finite subset of the lattice \mathbb{Z}^2 . We say that two points of A are *adjacent* if the distance between them is equal to 1. In this case we say they are connected by an *edge*, the line segment joining them. The set A thus receives a graph structure. Closely related to A is the set $\hat{A} \subseteq \mathbb{R}^2$, the interior of the union of closed squares of side 1 (in the usual position) with centers in A : we often identify A and \hat{A} and a point p of A with the unit square whose center is p . The graph A is called *connected* (or *simply connected*) if \hat{A} is. Without real loss, we always assume A to be connected. A *covering of A by edges* is a set of edges such that each point of A is the extremity of precisely one edge. Equivalently, we speak of *tilings of \hat{A} by dominoes*, each domino covering two unit squares connected by an edge.

A point of \mathbb{Z}^2 shall be called *white* (resp., *black*) if the sum of its coordinates is even (resp., odd); \hat{A} is therefore painted black and white like a chessboard. Edges of A connect points of different colours. Clearly, if A admits a tiling, the number of white squares equals that of black squares. In Figure 1.1, A is not tileable even though the colour condition is satisfied.




Figure 1.1

Our first aim is to state a necessary and sufficient condition for two tilings to be in the same connected component of T . In order to do this, we define combinatorial invariants for these components. We start with an explicit and easily computable description of such invariants, which is then rephrased in the vocabulary of homology theory.

A *cut* of \hat{A} is a simple oriented polygonal line in \hat{A} consisting of a sequence of edges of squares and joining two points in the boundary of \hat{A} . The *flow* of a tiling across a cut is defined to be the number of dominoes crossing the cut, where the domino is counted positively (resp., negatively) if its white square is to the left (resp., right) of the cut.

Consider a cut of \hat{A} disconnecting it into two sets \hat{A}_ℓ and \hat{A}_r to the left and right of the cut, respectively. The flow of any tiling of \hat{A} across this cut is clearly given by the number of white squares in \hat{A}_ℓ minus the number of black squares in \hat{A}_ℓ : this must be equal to the number of black squares in \hat{A}_r minus the number of white squares in \hat{A}_r . For a cut which does not disconnect \hat{A} , on the other hand, the flow may admit different values for different tilings, as in Figure 1.2. It is easy to see, however, that for a fixed cut, adjacent tilings in




Figure 1.2

$T(A)$ have the same flow: these are therefore invariants for the connected components of T . This is the easy part of Theorem 1.1 below.

Theorem 1.1: (combinatorial version) *Assume \hat{A} has genus n . Choose n disjoint cuts in \hat{A} which jointly do not disconnect \hat{A} . Two tilings t_1 and t_2 are in the same connected component of T if and only if their flows across each of the n chosen cuts are equal.*

In particular, if A is simply connected, T is connected. We will give two other equivalent versions of this theorem, and will prove the last one, after constructing the necessary tools.

Let us see how we can associate to two tilings t_1 and t_2 an element of $H_1(\hat{A}; \mathbb{Z})$, which we shall denote by $[t_1 - t_2]$. We first build two CW-complexes \mathbf{A} and \mathbf{A}^* with $\mathbf{A} \subseteq \hat{A} \subseteq \mathbf{A}^*$ and such that the inclusions are homotopy equivalences. For \mathbf{A} , the 0-cells are the points of A (the centers of the squares of \hat{A}), the 1-cells are the edges between points of A , the 2-cells are the open unit squares with all vertices in A . For \mathbf{A}^* , the 2-cells are the squares of \hat{A} , the 1-cells are their sides and the 0-cells their vertices, where the common side of two adjacent squares gives us only one 1-cell, as do the common vertices of adjacent squares but common vertices of non-adjacent squares are *not* identified unless the two squares are adjacent to a third one. In Figure 1.3, we show A (big dots), \mathbf{A} (big dots and thin lines) and \mathbf{A}^* (thick lines); notice that point p gives rise to two 0-cells in \mathbf{A}^* . For future use, call A^* the set of 0-cells of \mathbf{A}^* .




Figure 1.3

We shall be interested in a few related homology and cohomology moduli of the above spaces. Since the two CW-complexes are homotopy equivalent to \hat{A} , $H_1(\mathbf{A}^*; \mathbb{Z}) = H_1(\mathbf{A}; \mathbb{Z})$. By Poincaré-Lefschetz duality (See Sections 26 and 28 of [4]), on the other hand, $H_1(\hat{A}; \mathbb{Z}) = H^1(\mathbf{A}^*, \partial\mathbf{A}^*; \mathbb{Z})$. The equality is induced by the natural identification between

$C_1(\mathbf{A}; \mathbb{Z})$ and $C^1(\mathbf{A}^*, \partial\mathbf{A}^*; \mathbb{Z})$, where C_k and C^k are the usual spaces of k -complexes or cocomplexes.

Consider each edge (or domino) as a 1-cell in \mathbf{A} and always orient it from black to white; domino tilings correspond therefore to elements of $C_1(\mathbf{A}; \mathbb{Z})$ with boundary always equal to the sum of all white vertices minus the sum of all black vertices. The difference between two domino tilings t_1 and t_2 is therefore a closed element of $C_1(\mathbf{A}; \mathbb{Z})$: call the corresponding homology class $[t_1 - t_2]$. In Figure 1.4, we show how, given two tilings t_1 and t_2 ((a) and (b), resp.), we represent the class $[t_1 - t_2]$ in (c), consisting of a sum of cycles in $H_1(\hat{A}, \mathbb{Z})$.




Figure 1.4

This homology class turns out to have the following properties:

- (a) $[t_1 - t_1] = 0$,
- (b) $[t_1 - t_2] = -[t_2 - t_1]$,
- (c) $[t_1 - t_2] + [t_2 - t_3] = [t_1 - t_3]$,
- (d) if t_1 and t_2 are adjacent, $[t_1 - t_2] = 0$.

Properties (a), (b) and (c) are obvious. If t_1 and t_2 are adjacent, the cycle $[t_1 - t_2]$ is precisely the boundary of one of the 2-cells introduced in our construction, hence is exact, and (d) follows.

We defined $[t_1 - t_2]$ as an element of $H_1(\hat{A}; \mathbb{Z})$. From the duality stated above, we can also think of $[t_1 - t_2]$ as an element of $H^1(\mathbf{A}^*, \partial\mathbf{A}^*; \mathbb{Z})$; let us see a direct way of interpreting this cohomology class. In Figure 1.5.a we show the cocomplex corresponding to the tiling in Figure 1.4.a: notice that the cocomplex is represented in \mathbf{A}^* while the complex was represented in \mathbf{A} . The way to obtain the cocomplex from the tiling should be clear: take the 1-cells (edges) on boundaries of dominoes to zero and 1-cells crossing dominoes, when oriented so that the white square is at the left, to 1. The cocycle corresponding to the cycle in Figure 1.4.c is shown in Figure 1.5.b; notice that it is zero at the boundary and therefore corresponds to a class in $H^1(\mathbf{A}^*, \partial\mathbf{A}^*; \mathbb{Z})$, as claimed.

Consider the cocomplex in $C^1(\mathbf{A}^*; \mathbb{Z})$ taking any edge to 1 if oriented with white at the left: the difference $\langle t \rangle$ between this cocomplex and four times the cocomplex associated with a tiling t is a cocycle since it takes the boundary of any square to zero; notice that its value on $\partial\mathbf{A}^*$ does not depend on t . Since \mathbf{A}^* is a closed disk with holes, the map

induced by inclusion from $H^1(\mathbf{A}^*; \mathbb{Z})$ to $H^1(\partial\mathbf{A}^*; \mathbb{Z})$ is injective. The cohomology class in $H^1(\mathbf{A}^*; \mathbb{Z})$ corresponding to $\langle t \rangle$ does not therefore depend on the tiling t . In Figure 1.5.c we represent $\langle t \rangle$ for the tiling in Figure 1.4.a: in order to recover the tiling from the cocomplex, place central edges of dominoes over the triple arrows; in particular, different tilings correspond to different cocomplexes. We shall see the uses of the cocycle $\langle t \rangle$ in the next section.




Figure 1.5

To a cut Γ we associate an element $[\Gamma] \in H^1(\hat{A}, \mathbb{Z}) = H^1(\mathbf{A}, \mathbb{Z})$ as follows. Each element of C_1 (for \mathbf{A}) is mapped to an integer: a 1-cell which does not cross Γ is taken to 0; if the 1-cell crosses Γ , it is taken to ± 1 , according to orientation (if the 1-cell crosses Γ from right to left it is taken to 1). This map is a cocycle, i.e., the boundary of a 2-cell is taken to 0. This is the usual construction of a cohomology class from a curve such as a cut. Since the cohomology of a disk with n holes D is known to be generated by the classes of n such curves not disconnecting D , the n cuts mentioned in Theorem 1.1 form a basis of $H^1(\hat{A}, \mathbb{Z})$. Notice that the flow of a tiling t across a cut Γ is $[\Gamma](t)$; the difference between flows for two tilings t_1 and t_2 is the usual pairing (between cohomology and homology) $[\Gamma] \frown [t_1 - t_2]$. In particular, if $[t_1 - t_2] = 0$, the flows of t_1 and t_2 coincide on any cut.

Theorem 1.2: (homological version) *The tilings t_1 and t_2 are in the same connected component of T if and only if $[t_1 - t_2] = 0$.*

Both versions, Theorems 1.1 and 1.2, are equivalent. Indeed, from the remarks above, triviality of the class $[t_1 - t_2]$ is equivalent to the equality of the corresponding flows across the cuts mentioned in Theorem 1.1.

2. Height sections

We begin this section by discussing height functions, originally presented by Thurston ([9]). Height sections, which are appropriate extensions of the concept of height functions, are the main tools in our proofs of Theorems 3.1 and 3.2. The height function (or section) corresponding to t is in fact obtained by integrating $\langle t \rangle$; we nevertheless give an elementary and independent description of these objects.

Consider a (parametrized) polygonal line consisting of edges of unit squares with vertices in $(\mathbb{Z} + 1/2)^2$. We assign numerical values to the parametrized vertices by a sort of

integration process: in particular, it may happen that to a point on the line correspond two different values. Take an initial value (say 0) and assign it to the origin of the polygonal line. When walking along an edge with a white (resp., black) square to its left, add (resp., subtract) 1 to the value at the starting point of the edge in order to get the value at the endpoint. Notice that if the line joins P to Q and the integration process starting with a for P leads to b for Q then integration from Q to P along the same line starting with b yields the same value a at P .

If the endpoint coincides with the starting point of the line, how do the two values assigned to this point relate? It is not hard to see that we add (resp., subtract) 4 each time we surround a white (resp., black) square counter-clockwise, with reversed signs for opposite orientation. By the obvious additivity properties with respect to paths of integration, the value obtained when returning to the original point is the following. For each white (resp., black) square, take 4 (resp., -4) times the winding of the path around it and sum over all squares.

Thus, the value mod 4 at the endpoint does not depend on the integration path, and is given (up to a global additive constant) by the function $\phi : (\mathbb{Z} + 1/2)^2 \rightarrow \mathbb{Z}/(4)$ defined as $\phi(x, y) = 0$ if $\lfloor x \rfloor = x - 1/2$ and $\lfloor y \rfloor = y - 1/2$ are both even, $\phi(x, y) = 1$ if $\lfloor x \rfloor$ is odd and $\lfloor y \rfloor$ is even, $\phi(x, y) = 2$ if $\lfloor x \rfloor$ and $\lfloor y \rfloor$ are both odd, $\phi(x, y) = 3$ if $\lfloor x \rfloor$ is even $\lfloor y \rfloor$ is odd. However, integration along the boundary of a domino, or, more generally, of a simply connected tileable region assigns the same value to the starting point and endpoint. Notice that the situation above is very similar to two other more familiar constructions: the calculation of the area of a planar region by Green's Theorem and the computation of a complex integral by adding residues.

We now discuss height functions and their relation to tilings in the case when the closure of \hat{A} is a closed disk. Assume therefore that the closure of \hat{A} is a (topological) closed disk. Let $A^* \subseteq (\mathbb{Z} + 1/2)^2$ be, as above, the set of vertices of squares in \hat{A} . Choose a *basepoint* $p_0 = (x_0, y_0) \in A^*$, p_0 in the exterior boundary of \hat{A} , and a *base value* $v_0 \in \mathbb{Z}$ so that $v_0 \bmod 4 = \phi(x_0, y_0)$. Given a tiling t , we define a function θ from A^* to \mathbb{Z} at a typical point p by integrating along any path contained in boundaries of dominoes, starting from the basepoint p_0 with initial value $v_0 = \theta(p_0)$ and reaching p with value $\theta(p)$. This function does not depend on choices of paths. Indeed, as in the paragraph above, θ is locally well defined; our hypothesis on the global topology of \hat{A} guarantees that θ is also globally well defined. Given any path contained in boundaries of dominoes joining points p_1 and p_2 , integration along this path starting with $\theta(p_1)$ yields $\theta(p_2)$. Also, different choice of basepoint or base value produce the same height function up to an additive constant in $4\mathbb{Z}$. We call θ the *height function* of t ; in Figure 2.1 we show an example of a domino tiling and the corresponding height function.

Two points in A^* are called *adjacent* if the distance between them is 1 and the segment joining them is in the closure of \hat{A} . It is easy to see that a height function satisfies the following properties:

- (a) $\theta(x, y) \bmod 4 = \phi(x, y)$,
- (b) the values of θ at adjacent points never differ by more than 3,




Figure 2.1

(c) the values of θ at points which are adjacent along a segment contained in the boundary of \hat{A} always differ by exactly 1.

Conversely, given a function τ satisfying conditions (a), (b) and (c) as above, we obtain a tiling t as follows: join two adjacent points in A^* if the values of τ at such points differ by precisely 1, thus obtaining the contours of the dominoes of t . It remains to prove that the construction actually gives rise to a tiling by dominoes. Indeed, each square of \hat{A} is surrounded by four points of A^* and from conditions (a) and (b) exactly three of these sides are drawn in the above process: the fourth one (which cannot lie on the boundary, by (c)) indicates which way the domino covering our square goes. Furthermore, the height function θ corresponding to t is equal to τ , up to an additive constant in $4\mathbb{Z}$: these two constructions are the inverse of each other. We thus defined a bijection between the space of tilings T and the class of functions satisfying the three above conditions, i.e., height functions, modulo additive constants in $4\mathbb{Z}$.

Let us consider how to extend these concepts to the general case. First, there can be nasty points in $(\mathbb{Z} + 1/2)^2$ with all four edges arriving at it being part of the boundary of \hat{A} : as we have already seen in the construction of \mathbf{A}^* , such a point ought to be interpreted as *two* points in A^* with adjacency relations defined in the obvious way that assures the local good behaviour of A and A^* . Of course, height functions are free to assume different values at these two points.

A more serious problem comes from the consistency of θ along boundaries if \hat{A} is not simply connected. If inside one of the holes of \hat{A} the number of white squares is different from the number of black squares, no height function can exist because we get multivaluedness when following the boundary. Still, it is easy to construct such regions which admit tilings, as in Figure 1.2.

In cohomological terms, it is clear what is going on. The height function θ was obtained by integrating $\langle t \rangle$: this was possible because this cocycle is *exact*, i.e., corresponds to the cohomology class 0 in $H^1(\mathbf{A}^*; \mathbb{Z})$. In other words, $\langle t \rangle$ is the coboundary of θ . Now, the cohomology of a disk is trivial but if \hat{A} is not simply connected $H^1(\mathbf{A}^*; \mathbb{Z})$ is non-trivial and it may well happen (as in the example mentioned in the previous paragraph) that the cohomology class of $\langle t \rangle$ is non-zero.

What we need is not height functions but height sections of a certain fibre bundle with base space A^* and fibre \mathbb{Z} . In this bundle, a fibre is not an additive group: there is

no natural 0 nor addition on each fibre. We are allowed to add an integer to an element of a fibre (thus getting another element of the same fibre) or to subtract elements of the same fibre (thus getting an integer). We are also allowed to compare elements of the same or neighbouring fibres, but otherwise we are not allowed to compare elements of different fibres. The congruence class mod 4 of an element of a fibre is, however, well defined.

We begin the construction by choosing a basepoint p_0 in A^* . Consider next the set \mathcal{P} of all paths in A^* going from p_0 to some other point of A^* , i.e., functions ξ from sets of the form $\{0, 1, \dots, m\}$ to A^* such that $\xi(0)$ is the basepoint and $\xi(i)$ and $\xi(i+1)$ are always neighbours in A^* . Our bundle shall be obtained from $\mathcal{P} \times \mathbb{Z}$ by a quotient: the projection from $\mathcal{P} \times \mathbb{Z}$ to A^* just takes a pair (path, integer) to the endpoint $\xi(m)$ of the path. Two pairs (ξ_1, k_1) and (ξ_2, k_2) are identified if the following conditions hold. First, ξ_1 and ξ_2 must have the same endpoint. Second, consider ξ the path obtained by following ξ_1 and then following ξ_2 backwards; let ℓ be the sum of the windings of ξ around white squares *not in \hat{A}* minus the sum of the windings of ξ around black squares, again *not in \hat{A}* : identify the two pairs if $k_1 - k_2 = 4\ell$. This defines the desired *height bundle*, or \mathcal{H} . The allowed operations on this bundle have the obvious definitions in terms of representatives of the equivalence classes.

Another essentially equivalent interpretation for \mathcal{H} is as a (not necessarily connected) covering space for \hat{A} , or, equivalently, \mathbf{A}^* . Indeed, take the fibres as defined over A^* and extend them to edges of \mathbf{A}^* by the provided identification between neighbouring points. Finally, define fibres over the squares of \mathbf{A}^* in the essentially unique possible way: it is always possible to do it for each such square because the four identifications around it are compatible. The name ‘height section’ should generate no confusion: it is always to be understood as a section of \mathcal{H} restricted to A^* .

We construct the height bundle for the region shown in Figure 2.2(a), in a manner which is slightly different from the one described above. Start by drawing cuts as indicated, and consider the sub-CW-complex $\hat{\mathbf{B}}$ of \mathbf{A}^* , obtained by removing the 1- and 2-cells intersecting the cuts. Take now the cartesian product $\hat{\mathbf{B}} \times \mathbb{Z}$. This is necessarily isomorphic to the restriction of \mathcal{H} to $\hat{\mathbf{B}}$, since $\hat{\mathbf{B}}$ is contractible. In order to construct \mathcal{H} , it suffices to extend this bundle to the missing cells. This shall be done by choosing appropriate additive shifts between consecutive fibres, indicated again in Figure 2.2(a). How are those shifts obtained? Consider, for example, the two paths ξ_1 and ξ_2 in the picture; the equivalence relations defined in the construction of \mathcal{H} yield $(\xi_1, -4) = (\xi_2, 0)$. Of course, a different choice of cuts would give rise to isomorphic bundles. To construct an isomorphism, start by identifying (arbitrarily, but respecting orientation) a pair of fibres with the same base point and extend (in the only possible way) the identification to the entire bundle. It is only necessary to check that the above construction yields a well defined map: this follows from the definition of the bundles. Thus, from now on, we speak of *the* height bundle over A^* .

We define the *height section* corresponding to a tiling t by integration just as we did for height functions: the bundle \mathcal{H} is constructed in such a way that the definition of the section does not depend on choices of paths. Indeed, for two arbitrary paths along boundaries of dominoes with same starting point and endpoint, integration yields two




Figure 2.2

values at the endpoint which are identified in the construction of the bundle. As with height functions, basepoint and base value contribute only with a constant in $4\mathbb{Z}$. Figure 2.2(b) shows an example of a tiling and its height section; Figure 2.2(c) shows how we would write the same height section with a different set of cuts. The large difference between numbers on neighbouring points at opposite sides of cuts does *not* correspond to a jump of the height section: remember that such neighbouring fibres are attached with an additive shift. Again, a different choice of cuts would not have changed the height section itself, but only the notation employed.

A height section satisfies conditions (a), (b) and (c) with the appropriate invariant interpretation: in conditions (b) and (c), the difference between values of the section at neighbouring points is to be computed using the identification of neighbouring fibres intrinsic to the definition of the bundle, and *not* as a difference between the (cut-dependent) integers used in our examples. We again have a natural bijection between T and the class of sections of \mathcal{H} satisfying the three conditions, modulo constants in $4\mathbb{Z}$.

We now list some convenient properties of height sections. The difference of two height sections is a function with domain A^* and values in $4\mathbb{Z}$; this is well defined up to an additive constant. Consider a cut Γ connecting two points x_a and x_b in different boundary components of A^* . Let f_1 and f_2 be the flows across Γ of two tilings t_1 and t_2 . We claim that

$$\theta_2(x_b) - \theta_1(x_b) - \theta_2(x_a) + \theta_1(x_a) = -4(f_2 - f_1).$$

From the previous remark, the left hand side is a well defined integer. Indeed, the function $\theta_2(x) - \theta_1(x) - \theta_2(x_a) + \theta_1(x_a)$, from the cut Γ to the integers, can be computed, starting

at $x = x_a$ (where it is clearly equal to 0) and ending at $x = x_b$ by the integration processes, yielding the result. Also, two tilings are adjacent iff their corresponding height sections differ at a single (interior) point by ± 4 once the additive constants have been chosen so that they agree at a boundary point. Finally, the maximum or minimum of two or more height sections is again a height section, since properties (a), (b) and (c) are preserved.

It is clear from property (c) of height sections (in particular, functions), that if two height sections for the same region agree at one point of the boundary, they agree on the entire connected component of the boundary containing that point. We may thus assume without loss that height sections always agree on the exterior boundary. Our main interest, however, is on relating height sections for tilings t_1 and t_2 with $[t_1 - t_2] = 0$. In this case, the corresponding height sections agree on the *entire* boundary. Indeed, consider two sections θ_1 and θ_2 which agree on the exterior boundary. The identity above, which relates the values of the sections at one point of the boundary to the values at another point, immediately yields the equality of the sections at all boundary components, since flows for both tilings are equal, by the equivalence between Theorems 1.1 and 1.2.

There is then a lattice structure (and a partial order) on $T(A)$, induced by the corresponding order on height sections: remember that height sections are assumed to agree on the exterior boundary.

3. Distances in T

We state yet a different version of Theorems 1.1 and 1.2.

Theorem 3.1: (height section version) *The tilings t_1 and t_2 are in the same connected component of T if and only if their corresponding height sections θ_1 and θ_2 coincide on the whole boundary.*

As shown at the end of the previous section, $[t_1 - t_2] = 0$ iff the corresponding height sections agree on the whole boundary: the equivalence between Theorems 1.2 and 3.1 is now clear. A corollary of this theorem is that each connected component of T is a lattice, with a maximum and minimum height sections. The non-trivial part of Theorem 3.1 reduces therefore to the following: there always exists a path (in $T(A)$) joining two height sections coinciding on the boundary of \mathbf{A}^* , in which consecutive height sections differ at a single point.

Theorem 3.2: *Suppose the tilings t_1 and t_2 are such that their corresponding height sections θ_1 and θ_2 coincide on the whole boundary. Then t_1 and t_2 are in the same component of $T(A)$ and*

$$d(t_1, t_2) = \frac{1}{4} \sum_{p \in \mathbf{A}^*} |\theta_1(p) - \theta_2(p)|.$$

Also, the diameter of a connected component of $T(A)$ is the distance between the minimum and the maximum of all height sections in that component.

As we shall see, Theorem 3.1 and Theorem 3.2 follow easily from the lemma below.

Lemma 3.3: *Let $\theta_1 < \theta_2$ be two height sections coinciding on the boundary of \mathbf{A}^* . Then there exists θ_3 , θ_3 adjacent to θ_1 , with $\theta_1 < \theta_3 \leq \theta_2$.*

Proof:

To satisfy adjacency, the new section θ_3 must be constructed as follows: choose a point p_0 in A^* and define $\theta_3(p) = \theta_1(p)$ for $p \neq p_0$ and $\theta_3(p_0) = \theta_1(p_0) + 4$. The section θ_3 is a height section iff conditions (a), (b) and (c) hold. Condition (a) is trivially satisfied. Condition (c) is satisfied provided p_0 lies in the interior of A^* . Condition (b) holds if and only if p_0 is a local minimum of θ_1 . Indeed, local minimality and condition (a) guarantee that, if p is a neighbour of p_0 , $\theta_1(p) - \theta_1(p_0) = 1$ or 3 . Thus, $\theta_3(p) - \theta_3(p_0) = -3$ or -1 . Finally, to obtain $\theta_1 < \theta_3 \leq \theta_2$, we must choose p_0 with $\theta_1(p_0) < \theta_2(p_0)$. We only have to prove then that such a point p_0 exists. Consider that (non-empty) part B of the domain where $\theta_2 - \theta_1$ is maximum; we will see that there must be such a p_0 in B .

When height sections are just height functions, select p_0 so that $\theta_1(p_0)$ is minimum in B . We prove that p_0 is a local minimum in A^* . Let p be a neighbour of p_0 . If p is in B , we have $\theta_1(p) > \theta_1(p_0)$ by hypothesis. If p is not in B , let $x_i = \theta_i(p_0)$; conditions (a), (b) and (c) allow two possible values for each of $\theta_i(p)$: call these y_i and z_i with $y_i < x_i < z_i$. Clearly, $z_2 - z_1 = x_2 - x_1 = y_2 - y_1$ and, since $p \notin B$, $\theta_2(p) - \theta_1(p) < \theta_2(p_0) - \theta_1(p_0)$, whence $\theta_1(p) = z_1$ and $\theta_2(p) = y_2$, proving our claim in this second case.

The difficulty in the proof for sections lies in the fact that it doesn't make sense to look for global minima. From the previous arguments, however, a local minimum in B (which is necessarily a local minimum in A^*) is what we need. Suppose by contradiction that no such local minima exist: every point in B has a neighbour where θ_1 is smaller. Since B is finite, there exists a cycle $p_0, p_1, \dots, p_{N-1}, p_N = p_0$ of points of B with $\theta_1(p_i) > \theta_1(p_{i+1})$. Assume without loss that the cycle is simple (i.e., has no self-intersections), turns counterclockwise in the plane and encloses minimum area. It is clear that this minimum area is greater than 1.

We claim that $\theta_1(p_i) - \theta_1(p_{i+1}) = 1$. Indeed, if this is not the case, the difference equals 3. The edge joining p_i and p_{i+1} is the central edge of a domino which is common to both tilings t_1 and t_2 . The two points to the left of the oriented segment $p_i p_{i+1}$ also belong to B , and we may therefore insert these two points between p_i and p_{i+1} thus obtaining a new cycle with smaller enclosed area. If the new cycle is not simple, take a simple subcycle of it.

Also, the segments $p_i p_{i+1}$ and $p_{i+1} p_{i+2}$ form a right angle, since otherwise we would have a difference of 3 on one of the two edges. Finally, we cannot have p_i, p_{i+1}, p_{i+2} and p_{i+3} vertices of the same square traversed counterclockwise: otherwise, omit p_{i+1} and p_{i+2} to get a cycle with smaller area. It follows that the polygonal line joining midpoints between consecutive points of the cycle never turns left and this contradicts the fact that the cycle turns counterclockwise. ■

Proof of Theorems 3.1 and 3.2:

This lemma (plus induction) tells us that we can move from a smaller to a larger height section by flips; in particular, we can go from any height section to the maximum, thus proving the connectivity of classes of height sections with given boundary values (Theorem 3.1). The inequality $d(t_1, t_2) \geq \frac{1}{4} \sum_{p \in A^*} |\theta_1(p) - \theta_2(p)|$ is an obvious consequence of the fact at the end of the previous section relating height sections of adjacent tilings. As to

the non-trivial half of the distance formula, a shortest path is to move from one section to the maximum of the two and then to the other; we could equally well have first moved to the minimum and the distance would be the same. Our claim about the diameter follows from the distance formula; this, of course, finishes the proof of Theorem 3.2. ■

It is clear from the proof above that we know which flips to perform in order to get closer to a tiling t_2 starting from a tiling t_1 : simply compute both height sections and look for local minima of t_1 below t_2 , or local maxima of t_1 above t_2 . In this sense, there is a local characterization of the shortest paths in the graph $T(A)$.

Some of these paths should clearly be considered equivalent. For instance, let t_1 and t_2 be the tilings (a) and (d) in Figure 3.1: the two paths (abd) and (acd) are such an example. We render this notion precise by turning T into a CW-complex. The 0-cells are just the elements of T and the 1-cells connect adjacent tilings so that the notion of a connected component of T remains unaltered. The 2-cells are glued along squares whose edges are two independent flips (i.e., occurring on disjoint squares); in Figure 3.1, there is a 2-cell whose boundary is composed of the four 1-cells connecting the tilings in (a) to (b), (b) to (d), (d) to (c) and (c) to (a). Similarly, 3-cells correspond to three independent flips, and k -cells to k independent flips. The above mentioned equivalence of paths is of course homotopy and it turns out that all shortest routes between tilings are homotopic, as follows from the theorem below.




Figure 3.1

Theorem 3.4: *Each connected component of T is contractible.*

Proof:

For an arbitrary tiling t_0 (i.e., a 0-cell) we construct a homotopy from the identity to a constant function taking the entire connected component to t_0 . Start by those points which are furthest from t_0 : from each such point t there are a few (say k) possible flips. All these flips must necessarily approach the base tiling and must be independent; t is therefore the vertex of a k -cell corresponding to these k flips and it is easy to deform this cell, a k -dimensional cube, onto the walls of the cell that do not touch t without moving these walls. Repeat the process for all tilings different from t_0 , taking distances in decreasing order. ■

4. Quadriculated surfaces

In this section we generalize the constructions and results of the previous sections to the situation where \mathbf{A}^* is not a subset of the plane but a *quadriculated surface*. The idea of a quadriculated surface is very natural but its definition is somewhat technical: start with a finite collection of squares of unit side and glue certain pairs of sides (taking orientation of the sides into account) in such a way that the following two conditions hold. First, two sides of the same square are never identified. Two vertices of different squares are identified if they are the corresponding extremes of identified sides. Given an edge of a square and an incident vertex we can either replace the edge by the other edge on which the vertex lies or, if the edge is identified with an edge of some other square, pass to that edge and to the corresponding vertex. Performing these two operations in alternation, we see that a vertex in the surface (i.e., after identifications) corresponds to a sequence of vertices of squares; it is clear that such a sequence is either finite (if we reach the boundary, i.e., a non-identified side) or periodic. Our second condition is that periodic sequences must have length 4; intuitively, this says that the angles at vertices of squares are $\pi/2$ so that it is impossible to surround a point with less than or more than 4 squares. Of course, the surface may be non-orientable or not consistently colourable in black and white. As in the planar case, we consider only connected surfaces.

We can easily construct a quadriculated torus and a quadriculated Klein bottle by identifying opposite sides of a (quadriculated) rectangle in the usual way. More generally, any quadriculated torus can be constructed by taking the quotient of \mathbb{R}^2 by a sublattice of \mathbb{Z}^2 ; the construction of the general quadriculated Klein bottle is similar. It is easy to see that these are the only quadriculated surfaces with no boundary. Quadriculated cylinders and quadriculated Möbius bands are even easier to construct: start with any simply connected region in the plane and glue along congruent boundaries.

As for Euclidean manifolds, it is easy to define a developing map ([10]) from the universal cover of a quadriculated surface to the plane. Similarly, define the holonomy of a quadriculated surface: it is a homomorphism from the fundamental group of the surface to the group of isometries of \mathbb{Z}^2 . If the surface is a topological disk, it has trivial holonomy and may be thought of as some kind of Riemann surface over \mathbb{Z}^2 .

The notion of a domino tiling of a quadriculated surface is clear, as is the notion of a flip. As in the previous simpler situation, we want to characterize the connected components of T .

We now describe the correct generalization of $[t_1 - t_2]$, $\langle t \rangle$ and height sections. The cell complexes \mathbf{A} and \mathbf{A}^* are easily defined, as are A and A^* . Consider the original homological construction of $[t_1 - t_2]$: we draw an edge for each domino of either tiling, orienting those in t_1 from black to white and those in t_2 from white to black. The obvious difficulty in generalizing this construction is: there are no white or black squares now, and it may even be impossible to assign colour globally in a coherent way. This makes it clear that $H_1(\mathbf{A}^*; \mathbb{Z})$ is not the right place to try to define $[t_1 - t_2]$: we must instead use homology with local coefficients. Homology and cohomology with local coefficients are briefly described for the situation of interest in the Appendix to this Section. More precisely, let \mathcal{Z}_1 be a \mathbb{Z} -bundle over \mathbf{A}^* constructed as follows: put in \mathbb{Z} fibres over each square and glue fibres on neighbouring squares by identifying k on one fibre with $-k$ on the other. The gluing instructions provide us with fibres over edges and create no obstruction towards defining the fibre over a vertex because of our second condition on quadriculated surfaces; notice that on each square there is a privileged generator for the fibre, originally labeled 1, which we call *positive*. A more global characterization of \mathcal{Z}_1 is that its fibre twists along a given closed curve in \mathbf{A}^* iff this curve passes through an odd number of squares. If we try to colour squares alternatedly black and white we find that this is similar to constructing a section of \mathcal{Z}_1 : in particular, \mathbf{A}^* is bicolourable iff \mathcal{Z}_1 is trivial. It is now clear that our definition of $[t_1 - t_2]$ makes sense as an element of $H_1(\mathbf{A}; \mathcal{Z}_1)$: edges of any tiling t (i.e., edges connecting the centres of the two squares composing a domino) are oriented so that their boundaries come out as two points with positive coefficients.

The cohomological construction of $[t_1 - t_2]$ or $\langle t \rangle$ is of course similar but it has to be performed with different coefficients, as is to be expected from duality anyway. Let therefore \mathcal{Z}_2 be a \mathbb{Z} -bundle constructed over \mathbf{A}^* as follows: first put in fibres over each square as before, but now each generator of the fibre corresponds to a possible orientation for the square. Glue fibres on neighbouring squares so that orientations don't match (thus constructing the fibres over edges); again, our second condition on quadriculated surfaces states that the fibre is well defined on vertices. Equivalently, \mathcal{Z}_2 twists along a given closed curve iff the curve inverts either colour or orientation, but not both. The (very general) version of Poincaré duality for sheaves (as in [8]) guarantees that $H_1(\mathbf{A}^*; \mathcal{Z}_1) = H^1(\mathbf{A}^*, \partial\mathbf{A}^*; \mathcal{Z}_2)$; we provide a sketch of a direct proof of this isomorphism in the Appendix. Also, over any edge of a square, there is a natural correspondence between orientations for the edge and generators of the fibre of \mathcal{Z}_2 over the edge: choose an adjacent square, orient the square, and take the corresponding generator of the fibre of \mathcal{Z}_2 to correspond to the counterclockwise orientation for the edge. It is now easy to define $\langle t \rangle \in C^1(\mathbf{A}^*; \mathcal{Z}_2)$: for edges not crossing dominoes, take the corresponding generator; for edges crossing dominoes, take -3 times the same generator. Again, this gives us an element of $H^1(\mathbf{A}^*; \mathcal{Z}_2)$ whose restriction to the boundary does not depend on the tiling t but it is important to notice that since the map induced by the inclusion from $H^1(\mathbf{A}^*; \mathcal{Z}_2)$ to $H^1(\partial\mathbf{A}^*; \mathcal{Z}_2)$ is usually not injective, this does not mean that the cohomology class of $\langle t \rangle$ does not depend on t : in Figure 4.1, the two tilings of the torus produce $\langle t \rangle$'s which are not cohomologous in

$H^1(\mathbf{A}^*; \mathcal{Z}_2) = \mathbb{Z}^2$ (notice that \mathcal{Z}_2 is trivial). On the other hand, the hypothesis $[t_1 - t_2] = 0$ (in $H_1(\mathbf{A}^*; \mathcal{Z}_1) = H^1(\mathbf{A}^*, \partial\mathbf{A}^*; \mathcal{Z}_2)$) guarantees that $\langle t_1 \rangle$ and $\langle t_2 \rangle$ are cohomologous (in $H^1(\mathbf{A}^*; \mathcal{Z}_2)$).




Figure 4.1

As an additional example, consider the cylinders in Figures 4.2(a) and (b) and the Möbius bands in Figures 4.2(c) and (d). In (a), \mathcal{Z}_1 and \mathcal{Z}_2 are both non-trivial and $H_1(\mathbf{A}; \mathcal{Z}_1)$ (which by Poincaré duality equals $H^1(\mathbf{A}^*, \partial\mathbf{A}^*; \mathcal{Z}_2)$) is trivial (as discussed in the Appendix); all tilings are therefore homologous and the reader can easily check that T is connected. In (b), \mathcal{Z}_1 and \mathcal{Z}_2 are both trivial and $H_1(\mathbf{A}; \mathcal{Z}_1) = \mathbb{Z}$; there are 4 connected components in T classified by $[t_1 - t_2]$. In (c), \mathcal{Z}_1 is trivial, \mathcal{Z}_2 is not and $H_1(\mathbf{A}; \mathcal{Z}_1) = \mathbb{Z}$; T has 3 components in (c), again classified by $[t_1 - t_2]$. Finally, in (d), \mathcal{Z}_1 is non-trivial, \mathcal{Z}_2 is trivial, $H_1(\mathbf{A}; \mathcal{Z}_1) = 0$ and T is connected.




Figure 4.2

Our next step is to construct the height bundle \mathcal{H} and the height section θ in it; examples will be given in Figure 4.3. It is convenient to construct both simultaneously and, unlike the previous simpler situation of planar regions, the structure of \mathcal{H} depends to a certain extent on the tiling t . The fibres of \mathcal{H} are to be copies of \mathbb{Z} with no distinguished zero and not even a privileged orientation defining order on fibres; we are allowed to add to an element of a fibre of \mathcal{H} an element of the corresponding fibre of \mathcal{Z}_2 and we are allowed to compare for ‘equality’ elements of neighbouring fibres. In order to build \mathcal{H} and θ , take

\mathcal{Z}_2 on the vertices of \mathbf{A}^* and ‘forget’ the zero section and the exact way of identifying two neighbouring fibres: we shall take the old zero section to be θ (by definition) and glue neighbouring fibres with an additive shift. This shift is described, of course, by $\langle t \rangle$, so that $\langle t \rangle$ is the ‘derivative’ of θ by construction. As before, \mathcal{H} can be thought of as a covering space over \mathbf{A}^* .

Since \mathcal{H} is not the same for all t we have to explain how we can ever compare different height sections. The first observation is that the structure of \mathcal{H} depends on the cohomology class of $\langle t \rangle$ in $H^1(\mathbf{A}^*; \mathcal{Z}_2)$ only. Indeed, if $\langle t_1 \rangle$ and $\langle t_2 \rangle$ are cohomologous, their difference is by definition a coboundary and therefore a sum of coboundaries of ‘delta functions’, i.e., functions with support given by a single vertex. Construct a discrete path from t_1 to t_2 by adding one such ‘delta function’ at each step. The intermediate cocomplexes in this path usually do not correspond to tilings at all but they still allow for the construction of \mathcal{H} (and even θ) at intermediate steps. The isomorphism of consecutive height bundles (but not sections) is clear and our claim follows. If, furthermore, $[t_1 - t_2] = 0$ (as an element of $H^1(\mathbf{A}^*, \partial\mathbf{A}^*; \mathcal{Z}_2)$), the ‘delta functions’ are all in the interior of \mathbf{A}^* and, for the same procedure of taking intermediate bundles and sections, consecutive sections coincide on the boundary. Thus, in this case, θ_1 and θ_2 are sections coinciding on $\partial\mathbf{A}^*$ of the same bundle \mathcal{H} . We are thus ready to compare θ_1 and θ_2 in the relevant case $[t_1 - t_2] = 0$ if $\partial\mathbf{A}^*$ is non-empty, by the connectedness of \mathbf{A}^* . If $\partial\mathbf{A}^*$ is empty, however, we have to consider if the above construction of intermediate bundles and sections introduces any ambiguity in the identification of the two height bundles. For planar regions, height sections were well defined up to a constant. If \mathcal{Z}_2 is trivial, i.e., if it admits at least one non-zero section, the same thing happens. Otherwise, height sections are well defined given \mathcal{H} : the difference between two of them (obtained by integration from the same tiling) is a section of \mathcal{Z}_2 , hence 0. Thus, if $\partial\mathbf{A}^*$ is empty and \mathcal{Z}_2 is non-trivial, there is no ambiguity in comparing height sections, but if $\partial\mathbf{A}^*$ is empty and \mathcal{Z}_2 trivial we are free to add constants (i.e., sections of \mathcal{Z}_2) to any of the two sections. In any case, $[t_1 - t_2] = 0$ if and only if the height bundles for t_1 and t_2 are isomorphic and the height sections θ_1 and θ_2 coincide on boundaries.

We should be able to characterize height section by properties similar to (a), (b) and (c) above. Properties (b) and (c) do not change: just remember to interpret them as taking place inside \mathcal{H} (and not some cut-dependent system of coordinates you may want to use). Property (a), however, has to be rephrased a bit more carefully. Inside each fibre of \mathcal{H} there exists a class of elements with the ‘right’ congruence mod 4, i.e., those elements which differ from the height section used for the construction of \mathcal{H} by a multiple of 4. We call the union of such subsets \mathcal{H}_ϕ , a subset of \mathcal{H} ; \mathcal{H}_ϕ is not quite a fibre bundle however since it is defined only over A^* and can not be naturally extended to \mathbf{A}^* since the height section itself is only defined over A^* . Property (a) now says that a height section must assume values in \mathcal{H}_ϕ . It should be clear that again these properties characterize height sections.

In Figure 4.3 we show the height sections for the four tilings in Figure 4.2. Notice how simple it is to construct such height sections: work as if the region were planar and at the end the identifications will be automatically provided. In these examples a value x for the

height section on a point at the left cut corresponds to a value of $1 - x$, x , $-3 - x$ and x in Figures (a), (b), (c) and (d), respectively, for the corresponding point at the right cut.




Figure 4.3

The Möbius band in example (c) in Figure 4.3 illustrates an interesting point when compared to the tilings in Figure 4.4. Here, the cohomology group $H^1(\mathbf{A}^*; \mathcal{Z}_2)$ is isomorphic to $\mathbb{Z}/(2)$. As discussed, the structure of \mathcal{H} depends on the cohomology class of $\langle t \rangle$ in $H^1(\mathbf{A}^*; \mathcal{Z}_2)$ only. However, two different tiling t_1 and t_2 for this region induce cohomology classes $\langle t_1 \rangle$ and $\langle t_2 \rangle$ differing by a multiple of 4 in $H^1(\mathbf{A}^*; \mathcal{Z}_2)$ and being therefore equal. The height section in Figure 4.4(a) does not appear at first to be in the same bundle as Figure 4.3(c) but an appropriate renaming of the fibres as in 4.4(b) shows that, as predicted, the bundles are indeed isomorphic (they have the same gluing instructions) even though the subsets \mathcal{H}_ϕ are different in the two cases. Similarly, the height section in 4.4(c) can be renamed as in 4.4(d) to fit inside the bundle for 4.3(c) but the values of the section at the boundary are different.




Figure 4.4

This shows that if t_1 and t_2 are in the same connected component then $[t_1 - t_2] = 0$. It is disconcerting at this point to realize that the converse is *false*: in the simple example shown in Figure 4.4, \mathbf{A}^* is a cylinder, \mathcal{Z}_1 and \mathcal{Z}_2 are both trivial and the reader will have no trouble checking that $[t_1 - t_2] = 0$ (or in computing height sections). No flip, however, is possible.




Figure 4.5

If we try to follow the proof of Lemma 3.3 in this example, we see what the problem is. The height sections differ by 4 along the entire central zig-zag (which actually contains all points of A^* not on the boundary). No point is, however, a local minimum or maximum for any of the two height sections.

Let us consider this counter-example from a slightly different point of view. By a *ladder* we mean a sequence of parallel dominoes side by side such that: two neighbouring dominoes always touch along one edge of the longer side, each domino in the ladder has two neighbours in it and these two neighbours touch the domino at different squares. In Figure 4.4 the two tilings consist of two ladders each. The important thing about ladders is that they are totally immune to flips. So, if t_1 and t_2 are in the same connected component then $[t_1 - t_2] = 0$ and t_1 and t_2 have precisely the same ladders. It may surprise the reader that this rather ad-hoc condition is actually necessary and sufficient.

Theorem 4.1: *Two domino tilings t_1 and t_2 are in the same connected component of T if and only if $[t_1 - t_2] = 0$ and t_1 and t_2 have precisely the same ladders. Furthermore, if this is the case, the distance between them is given by*

$$d(t_1, t_2) = \frac{1}{4} \sum_{p \in A^*} |\theta_1(p) - \theta_2(p)|;$$

in the case where there is no boundary and \mathcal{Z}_2 is trivial, the additive constants in the height sections are to be chosen so that the right hand side is minimum.

The right hand side of the distance formula makes sense (and is an integer): $\theta_1(p)$ and $\theta_2(p)$ are in the same fibre of \mathcal{H} , $\theta_1(p) - \theta_2(p)$ is an element of the corresponding fibre of \mathcal{Z}_2 , whose absolute value is in \mathbb{Z} . As with Theorems 3.1 and 3.2, we isolate the inductive step in a lemma.

Lemma 4.2: *Let \mathbf{A}^* be a quadriculated surface and let t_1 and t_2 be two different tilings of it with $[t_1 - t_2] = 0$ and such that neither of them has ladders; let θ_1 and θ_2 be the*

corresponding height sections. Assume that θ_1 and θ_2 coincide on a non-empty set (possibly the boundary). Then there exists a tiling t_3 of the same region, obtained from t_1 by a flip and such that the corresponding height section θ_3 always lies between θ_1 and θ_2 .

Proof:

As in Lemma 3.3, let B be that part of the domain where $|\theta_1 - \theta_2|$ is maximum; by hypothesis, B is neither empty nor equal to \mathbf{A}^* . We claim there is a point of B where we can perform a flip on t_1 in order to obtain t_3 : we call such a point (with a certain abuse of notation) a *local minimum* of θ_1 . Again as in Lemma 3.3, therefore, our aim is to prove the existence of such a local minimum. Suppose by contradiction there is no such point: we show the existence of a ladder.

Let p and p' be two neighbouring points in B . We say that, when moving from p to p' , θ_1 changes as if trying to get further from θ_2 if $|x_1 - \theta_2(p')| < |\theta_1(p') - \theta_2(p')|$, where x_1 is the element of the fibre of \mathcal{H} over p' which belongs to \mathcal{H}_ϕ , is different from $\theta_1(p')$ and satisfies $|x_1 - \theta_1(p)| \leq 3$. A point p in B is a local minimum of θ_1 if and only if it has no neighbour p' in B such that, when moving from p to p' , θ_1 changes as if trying to get further from θ_2 . Thus, since B is finite, there exists a cycle $p_0, p_1, \dots, p_{N-1}, p_N = p_0$ of points of B such that, when going from p_i to p_{i+1} , θ_1 changes as if trying to get further from θ_2 . Call such cycles *monotonic*. We may interpret a cycle as a 1-complex; we call two monotonic cycles *adjacent* if their difference is the boundary of a square in \mathbf{A}^* . Two monotonic cycles are *homotopic* if they can be joined by a sequence of adjacent monotonic cycles; thus, monotonic cycles break into homotopy classes. If a cycle does not reverse orientations, we can consistently speak of left and right; since an orientation reversing cycle yields an orientation preserving one by running along it twice we assume from now on, without loss, that we are dealing with orientation preserving cycles. It makes sense therefore to speak of left and right of a cycle and, given two adjacent cycles, we can naturally order them by saying that one is to the left and the other one to the right.

Claim: Inside each homotopy class there are a leftmost and a rightmost monotonic cycles.

Supposing the opposite, it would always be possible to push a cycle to the left (say), obtaining a closed sequence $c_0, c_1, \dots, c_{M-1}, c_M = c_0$ of adjacent monotonic cycles such that c_{i+1} is to the left of c_i . The contradiction arises from proving that the existence of a closed sequence of cycles as above implies that \mathbf{A}^* is a torus or a Klein bottle and that the height sections θ_1 and θ_2 never coincide. By going to the universal cover and using the developing map as in [10], each cycle c_i becomes a periodic line \tilde{c}_i in \mathbb{Z}^2 , the period being an orientation preserving isometry of \mathbb{R}^2 preserving \mathbb{Z}^2 , thus either a translation or a rotation of period 2 or 4. If the period of \tilde{c}_0 is not a translation, the curve \tilde{c}_0 surrounds a certain signed area, which decreases in the process of passing from \tilde{c}_i to c_{i+1} , contradicting the fact that the isometric curves \tilde{c}_N and \tilde{c}_0 enclose equal areas; thus, the period is a translation. Also, any isometry taking the infinite curve \tilde{c}_0 to \tilde{c}_N is another translation, since rotations would move remote points by distances far greater than M ; also, the two translations are linearly independent since passing from \tilde{c}_i to c_{i+1} moves curves to the left. The cycle c_0 gives rise to a closed curve in \mathbf{A}^* by connecting neighbouring points; similarly, the points $c_i(0)$ are joined to produce a second closed curve, based on the same point $c_0(0)$. These curves can be interpreted as elements of $\pi_1(\mathbf{A}^*, c_0(0))$ and the above translations are their

representations under holonomy. By extending the discrete homotopy of cycles to a map from the rectangle $[0, N] \times [0, M]$ to \mathbf{A}^* we see that these two elements of $\pi_1(\mathbf{A}^*, c_0(0))$ commute, thus generating a copy of \mathbb{Z}^2 inside $\pi_1(\mathbf{A}^*, c_0(0))$. The only compact surfaces, however, for which the fundamental group contains a copy of \mathbb{Z}^2 are a torus or a Klein bottle, since any other surface is hyperbolic and there is no copy of \mathbb{Z}^2 inside the isometries of the hyperbolic plane (see [10]). Since this construction is performed in B , $B = A^*$ and the two height sections never meet. The proof of the claim is thus complete.

Consider these two extreme cycles: they behave very similarly to the least area cycle in the proof of Lemma 3.3. In fact, repeating the same steps, we see that the polygonal line joining midpoints between consecutive points of the leftmost (resp., rightmost) cycle never turns left (resp., right). Now, since these two cycles are homotopic these two polygonal lines turn by the same angle and it follows that neither turns at all: both cycles are zig-zag lines exactly like boundaries of ladders. Furthermore, to the left of the leftmost cycle or to the right of the rightmost cycle t_1 and t_2 must each have a ladder since we cannot have arrived at the boundary. This contradicts the hypothesis and ends the proof of the Lemma. ■

Proof of Theorem 4.1:

All we have to prove is that if $[t_1 - t_2] = 0$ and t_1 and t_2 have the same ladders then t_1 and t_2 are in the same component and the distance between them is smaller than or equal to the expression at the right hand side in the statement of the theorem. Let therefore t_1 and t_2 be tilings as above. Start by removing all ladders from \mathbf{A}^* : we have to prove that the tilings on each connected component of whatever remains are in the same connected of T . It is clear that on each such connected component height bundles for t_1 and t_2 are isomorphic and the height sections coincide on whatever remains of the old boundary and differ by a constant on boundaries of removed ladders. We claim that we can never have a connected component of the boundary consisting of the boundary of a ladder only: indeed, if this happened, the only way to tile the neighbourhood of this boundary component would be with a new ladder. It follows that θ_1 and θ_2 coincide on the entire boundary of each connected component of whatever remains after removing ladders. We can therefore assume without loss of generality that t_1 and t_2 have no ladders.

When \mathbf{A}^* has boundary, Lemma 4.2 (plus induction) finishes with the proof. If \mathbf{A}^* has no boundary, we must consider two cases. If \mathcal{Z}_2 is non-trivial, the two height sections must coincide at some point by topological reasons: if they did not, their difference would yield a global choice of generators for \mathcal{Z}_2 (the difference is to be a positive multiple of the chosen generator) hence a trivialization of \mathcal{Z}_2 (since the fibre is one-dimensional). If \mathcal{Z}_2 is trivial, add a constant to θ_2 in order to make the right hand side of the distance formula minimum: it is clear that now θ_1 and θ_2 coincide at some point. ■

As in the planar case, we know which flips to perform in order to get closer to a tiling t_2 starting from a tiling t_1 , assuming, of course, $[t_1 - t_2] = 0$. Start by computing the (isomorphic) height bundles and the height sections θ_1 and θ_2 which must coincide on the boundary. In case there is no boundary and \mathcal{Z}_2 is trivial, adjust the additive constant to make distance minimum (this may allow for one or two answers). Now flip at any local

extremum of θ_1 if that takes the section closer to θ_2 . Again, there is a local characterization of shortest paths in T . However, not all paths are homotopic anymore.

Theorem 4.3: *If \mathbf{A}^* has boundary or \mathcal{Z}_2 is non-trivial, all connected components of T are contractible. If \mathbf{A}^* is a torus or a Klein bottle and \mathcal{Z}_2 is trivial, there are two kinds of connected components of T : some consist of one single isolated point which corresponds to a tiling constructed entirely from ladders; others are homotopy equivalent to \mathbb{S}^1 .*

Proof:

When \mathbf{A}^* has boundary or \mathcal{Z}_2 is non-trivial, the proof is entirely analogous to the planar case. From now on, assume the other situation. Notice first that if a tiling contains a ladder, it must consist of ladders only: only a ladder fits beside a ladder. We now prove that a tiling with height section θ_1 which admits no flips must be of this type. Assume first that \mathbf{A}^* is a torus, the quotient of \mathbb{R}^2 (quadriculated by \mathbb{Z}^2) by a 2-dimensional sublattice L of \mathbb{Z}^2 . Raise the tiling to the universal cover in order to obtain an L -periodic tiling of the plane. Taking θ_2 to be $\theta_1 + 4$, as in Lemma 4.2, there must exist a monotonic cycle c which, raised to the universal cover \mathbb{R}^2 , must connect the origin to some other point of L ; without loss, this point is of the form (x, y) with $x \geq y \geq 0$. By the triviality of \mathcal{Z}_2 , x and y must be of the same parity. Raise θ_1 to a height function $\tilde{\theta}_1$ in the plane: we can assume without loss that $\tilde{\theta}_1(0, 0) = 0$. Also, the value of $\tilde{\theta}_1$ decreases along the lifted monotonic cycle \tilde{c} . The value of $\tilde{\theta}_1(x, y)$ must be precisely $-2x$: a smaller value is impossible for any height function by conditions (a), (b) and (c) and a larger value does not allow for a monotonic decreasing path from the origin to (x, y) . If $x = y$ (in which case $x > 0$) the monotonic path \tilde{c} must be a zig-zag going from the origin to (x, y) which can not cross dominoes and must therefore be a side of a ladder. Otherwise, the values of the height function at 0 and (x, y) are enough to dictate the values on a parallelogram with vertices at these points. Since, as in the proof of Lemma 4.2, there exist monotonic cycles through every point, the whole height function is well determined and the tiling must look like a garden variety brick wall, constructed from ladders going both ways. We take care of the Klein bottle by going to the orientable double cover, which is a torus.

For the other cases we claim that the universal cover of the corresponding connected component of T consists of all height sections *without* identifying sections which differ by a constant. It is clear that this is a covering map and what the CW-complex structure for this space must be. It is enough to prove that this space is connected and contractible since the quotient group will obviously be \mathbb{Z} , or, more precisely, $4\mathbb{Z}$. Since after identifications this space is known to be connected, it is enough to prove that we can move by flips from a section θ_1 to $\theta_1 + 4$. From the previous paragraph, we can perform *some* flip on θ_1 , without loss an increasing one, to obtain θ_2 ; but now θ_2 intersects $\theta_1 + 4$ at the flipped point and by Lemma 4.2 and Theorem 4.1 we can move θ_2 to $\theta_1 + 4$ by flips. The proof that the space of sections is contractible is similar to what we already saw in the previous cases, the fact that there are infinitely many cells being no source of trouble: a point contained in a cell such that its furthest vertex from the base section is at a distance d starts moving at time $1/2^d$.

This argument also shows that the generator of the fundamental group of a connected component of T is the path from θ to $\theta + 4$. Actually, such a closed path is a deformation retract of the connected component, but we give no details. In Figure 4.6 we show the four steps of such a cycle for the only non-trivial component of $T(\mathbf{A}^*)$, where $\mathbf{A}^* = \mathbb{R}^2/(2\mathbb{Z})^2$ (move from (a) to (b) to (c) to (d) to (e)). ■




Figure 4.6

Appendix: Homology and cohomology with local coefficients

This Appendix contains a brief review of the main facts about homology and cohomology with local coefficients which are necessary or convenient for us. More specifically, we apply the general constructions to our examples. Readers which know enough about the subject to compute homology and cohomology in simple examples and who are acquainted with Poincaré duality in this context are encouraged to skip the Appendix altogether. There are good expositions of the subject in [11] and [7]; for the more general theory of sheaves, the reader may consult [8].

We begin with a description of $H_1(\mathbf{A}, \mathcal{Z}_1)$; as in the usual homology, this is obtained from a chain complex of additive groups $C_2 \rightarrow C_1 \rightarrow C_0$ by taking the quotient Z_1/B_1 , where Z_1 is the kernel of the second boundary map and B_1 is the image of the first. Recall that the fibre of \mathcal{Z}_1 over a square of \mathbf{A}^* has a *positive* and a *negative* generator; thus, the fibre of \mathcal{Z}_1 over vertices of \mathbf{A} also has a positive and a negative generator. Thus, the generators of the fibre of \mathcal{Z}_1 over an edge of \mathbf{A} are positive on one of extrema and negative on the other. Over a square of \mathbf{A} , the generators are alternately positive and negative over the four vertices. The additive groups C_i are generated by (formal) products of an oriented i -cell in \mathbf{A} by a generator of the fibre over it. Thus, generators of C_0 are vertices with an orientation, i.e., a sign, as in Figure A.1(a). Similarly, generators of C_1 and C_2 are indicated in Figures A.1(b) and (c); the second equality in (b) is a notational convenience. The action of the boundary maps $C_{i+1} \rightarrow C_i$ over generators is indicated in Figures A.1(d) and (e); notice that the composition of both is zero.

We provide a similar description of the relative cohomology group $H^1(\mathbf{A}^*, \partial\mathbf{A}^*; \mathcal{Z}_2)$. Again, our first task is to construct C^2, C^1, C^0 and the coboundary maps. The fibre of \mathcal{Z}_2 over a square of \mathbf{A}^* is isomorphic to \mathbb{Z} . An orientation for the square and a sign (which again alternates between neighbouring squares) determine a generator of the fibre: changing one of these ingredients alters the generator. Thus, the additive group C^2 is generated by the map taking a given oriented square of \mathbf{A}^* to the generator of \mathcal{Z}_2 over this same square corresponding to the orientation of the square and the plus sign: we




Figure A.1

denote such generators as in Figure A.2(a). Generators of C^1 are maps taking a given non-boundary oriented edge of \mathbf{A}^* to the generator of \mathcal{Z}_2 defined as follows: choose any of the two adjacent squares to the edge, orient it so that the induced orientation on the boundary equals the original orientation of the edge, and take the generator of \mathcal{Z}_2 over it corresponding to its orientation and the plus sign. It is easy to check that this map does not depend on the choice of the adjacent square; we denote the generators of C^1 as in Figure A.2(b). Generators of C^0 are maps taking an interior vertex of \mathbf{A}^* to one of the generators of the fibre of \mathcal{Z}_2 over it; the choice of the generator is indicated by an orientation and signs for the neighbouring squares as in Figure A.2(c). Coboundary maps over generators are indicated in Figures A.2(d) and (e). In order to obtain the cohomology group $H^1(\mathbf{A}^*; \mathcal{Z}_2)$, drop the restrictions that edges or vertices must be interior.




Figure A.2

We recall the basic facts concerning Poincaré duality. Ordinary Poincaré duality ([4]) works by identifying $C_k(M)$ for a given triangulation with $C^{n-k}(M)$ for the dual triangulation, where M is an n -dimensional oriented closed manifold (the orientation is used in the identification procedure). This identification commutes (up to signs) with boundary and coboundary operations and thus induces, by taking quotients, the identifications between $H_k(M)$ and $H^{n-k}(M)$. In Lefschetz duality ([4]), we work with oriented n -manifolds with boundary and identify $C_k(M)$ with $C^k(M, \partial M)$ again by looking at dual triangulations. More generally, we can consider local coefficients \mathcal{Z} and, still for an oriented

manifold with boundary, essentially the same construction yields an identification between $C_k(M; \mathcal{Z})$ and $C^{n-k}(M, \partial M; \mathcal{Z})$. One way to get rid of the orientability hypothesis is to let the cohomology coefficients take care of the problem ([8]): if \mathcal{Z}_0 is the \mathbb{Z} -bundle over M with generators corresponding to (local) orientations, the appropriate generalization of Poincaré's construction provides the identification

$$C_k(M; \mathcal{Z}) = C^{n-k}(M, \partial M; \mathcal{Z} \otimes \mathcal{Z}_0).$$

By taking quotients, we obtain the duality we need:

$$H_k(M; \mathcal{Z}) = H^{n-k}(M, \partial M; \mathcal{Z} \otimes \mathcal{Z}_0). \quad (*)$$

Going back to our context, it is easy to see that $\mathcal{Z}_2 = \mathcal{Z}_1 \otimes \mathcal{Z}_0$, where \mathcal{Z}_0 is constructed as above. The identification

$$H_1(\mathbf{A}; \mathcal{Z}_1) = H^1(\mathbf{A}^*, \partial \mathbf{A}^*; \mathcal{Z}_2)$$

is a special case of (*). Notice, however, that our descriptions of the chain and cochain complexes yield an explicit construction of this bijection: just match corresponding letters in Figures A.1 and A.2.

We compute the homology groups which appear in Figure 4.2. Whenever the coefficient bundle is trivial, we are dealing with the usual homology group with coefficients in \mathbb{Z} ([11]), and this takes care of (b) and (c). Otherwise, by invariance of homology under deformation retracts, we are reduced to computing $H_1(S^1; \mathcal{Z})$, where \mathcal{Z} is the non-trivial \mathbb{Z} -bundle over S^1 . Consider the very simple CW-decomposition of the circle with a single edge having both extrema attached to the same 0-cell. The groups C_0 and C_1 are both cyclic and the boundary map takes a generator of C_1 to twice a generator of C_0 ; thus, $H_1(S^1; \mathcal{Z}) = 0$, as claimed, and $H_0(S^1; \mathcal{Z}) = \mathbb{Z}/(2)$. In the comments concerning Figure 4.4, we state that $H^1(\mathbf{A}^*; \mathcal{Z}_2) = \mathbb{Z}/(2)$, where \mathbf{A}^* is a Möbius band and \mathcal{Z}_2 is non-trivial: again, by invariance under deformation retracts, it suffices to compute $H^1(S^1; \mathcal{Z})$ with \mathcal{Z} as above. The groups C^0 and C^1 are both cyclic and C^2 is trivial; the coboundary takes a generator of C^0 to twice a generator of C^1 , so that $H^1(S^1; \mathcal{Z}) = \mathbb{Z}/(2)$ and $H^0(S^1; \mathcal{Z}) = 0$.

5. Final remarks

A. Calisson tilings

A *calisson* is the union of two equilateral triangles with a common side. Calisson tilings of simply connected regions in the plane admit height functions ([9]) with a strong visual interpretation: by looking at a calisson tiling, you can see it as a figure of a pile of (3-dimensional) cubes, the calissons being their faces ([2]). In close analogy with what we did in this paper, we can define height sections for calisson tilings of other regions. In this context, we perform a *flip* by lifting three calissons forming a hexagon and placing them back in the only possible different configuration. Clearly, two tilings are adjacent by a flip if their height sections differ at a single point. Under the pile-of-cubes interpretation, a

flip corresponds to adding or removing a cube. Thus, for simply connected regions, the space of tilings is connected and the distance between two tilings is given by the number of non-common cubes. Height sections might be useful for a more careful study of the space of calisson tilings of more complicated regions.

B. The adjacency matrix of A

The adjacency matrix of A is of the form

$$\begin{pmatrix} 0 & M \\ M^T & 0 \end{pmatrix},$$

provided white vertices are listed before black vertices. The sign of $\det M$ is not natural: it depends on the order in which the vertices are listed. Tilings of \mathbf{A}^\star correspond to monomials in the expansion of the determinant of M . Indeed, such a monomial (up to sign) corresponds to a set of 1's in M with exactly one element in each row or column: each 1 gives rise to an edge of A and it is clear that the associated set is a covering by edges. Tilings are thus naturally divided into two classes according to the sign of the corresponding monomial and we shall say that two tilings have the same or opposite parities if the corresponding monomials have the same or opposite signs (there is, however, no natural definition of an 'even' and an 'odd' tiling). It is easy to see that adjacent tilings always have opposite parities.




Figure 5.1

The absolute value of the determinant of M is the difference between the number of tilings of each parity: in [3], it is shown that, when \mathbf{A}^\star is a simply connected surface, this difference is 0 or 1. On the other hand, if A is not simply connected, this difference can have any value (see [3] or consider instead a $4 \times (2n - 1)$ rectangle with $n - 1$ vertical isolated dominoes removed from its interior). Since A being simply connected implies the connectivity of $T(A)$, it might be thought that the correct generalization to non-simply connected regions would be that, on each connected component of $T(A)$, this difference still is 0 or 1. In the examples shown in Figure 5.1, however, there are always three connected

components with differences of 1, -2 and 1 (in the natural order). Indeed, in both cases it is easy to see that $\det(M) = 0$ by considering the element of the kernel indicated in the Figure; the fact that there are three connected components follows from Theorem 1.1, by merely constructing tilings with different flows, and it is just as easy to see that two of these components have a single element each, always with the same parity: our claim follows.

C. Higher dimensions

The obvious generalization of Theorem 1.1 to higher dimensions is false even if \mathbf{A}^* is a topological closed ball contained in \mathbb{Z}^n (although the definition of $[t_1 - t_2]$ still works, and properties (a), (b), (c) and (d) as above still hold). In dimension 3, let

$$A = \{(0, 0, 0), (0, 0, 1), (0, 0, 2), (0, 1, 0), (0, 1, 1), (0, 1, 2), (1, 0, 1), (1, 0, 2), (1, 1, 0), (1, 1, 1)\}.$$

The tiling

$$([0, 1], 1, 0), ([0, 1], 0, 2), (1, [0, 1], 1), (0, 0, [0, 1]), (0, 1, [1, 2])$$

has no adjacent tilings (since there is no square to flip) but is not the only one: consider

$$(0, [0, 1], 0), (0, [0, 1], 1), (0, [0, 1], 2), (1, 0, [1, 2]), (1, 1, [0, 1]).$$

As another example, now in dimension 4, let $A = \{0, 1\}^4$ be the cube of side 2. The tiling

$$\begin{aligned} &([0, 1], 0, 0, 0), ([0, 1], 1, 1, 1), (0, 0, 1, [0, 1]), (0, 1, [0, 1], 0), \\ &(0, [0, 1], 0, 1), (1, 1, 0, [0, 1]), (1, 0, [0, 1], 1), (1, [0, 1], 1, 0) \end{aligned}$$

again has no neighbours but is not the only one. By the way, we know of no satisfactory extension of the idea of height sections to higher dimensions: the definition of $\langle t \rangle$ as a $(n - 1)$ -cocycle still works but, even if this is exact, its integral is not close to being unique.

References

- [1] J. H. Conway and J. C. Lagarias, Tilings with polyominoes and combinatorial group theory, *J. Comb. Theor.* **A53**, 183-208 (1990).
- [2] G. David and C. Tomei, The problem of the calissons, *Amer. Math. Monthly*, **96**, 429-431 (1989).
- [3] P. A. Deift and C. Tomei, On the determinant of the adjacency matrix for a planar sublattice, *J. Comb. Theor.* **B35**, 278-289 (1983).
- [4] M. J. Greenberger and J. R. Harper, *Algebraic Topology, A First Course (Revised)*, Benjamin-Cummings, Menlo Park, California, 1981.
- [5] P. W. Kasteleyn, The statistics of dimers on a lattice I. The number of dimer arrangements on a quadratic lattice, *Physica* **27**, 1209-1225 (1961).
- [6] E. H. Lieb and M. Loss, Fluxes, Laplacians and Kasteleyn's theorem, *Duke Math. Jour.*, **71**, 337-363 (1993).

- [7] E. H. Spanier, Algebraic Topology, New York: McGraw-Hill, 1966.
- [8] R. G. Swan, The Theory of Sheaves, The University of Chicago Press, Chicago, 1964.
- [9] W. P. Thurston, Conway's tiling groups, Amer. Math. Monthly, **97**, 8, 757-773 (1990).
- [10] W. P. Thurston, Geometry and Topology of Three-manifolds, Department of Mathematics, Princeton University, 1979.
- [11] G. W. Whitehead, Elements of Homotopy Theory, GTM 61, Springer-Verlag, New York, 1978.

Nicolau C. Saldanha, IMPA and PUC-Rio

e-mail: nicolau@impa.br, <http://www.impa.br/~nicolau/>

Carlos Tomei, IMPA and PUC-Rio

e-mail: tomei@impa.br

Mario A. Casarin Jr., NYU

e-mail: casarin@math1.cims.edu

Domingos Romualdo, Harvard University

e-mail: romualdo@husc.harvard.edu

Instituto de Matematica Pura e Aplicada, Estrada Dona Castorina, 110

Jardim Botânico, Rio de Janeiro, RJ 22460-320, BRAZIL

Departamento de Matemática, PUC-Rio, Rua Marquês de São Vicente, 225

Gávea, Rio de Janeiro, RJ 22453-900, BRAZIL

Courant Institute of Mathematical Sciences, 251, Mercer St.

New York, NY 10012, USA

Department of Economics, Harvard University, One Oxford St., Science Center 325

Cambridge, MA 02138, USA