On the connectivity of three-dimensional tilings

Juliana Freire Caroline Klivans Pedro H. Milet
Nicolau C. Saldanha

February 1, 2017

Abstract

We consider domino tilings of three-dimensional cubiculated manifolds with or without boundary, including subsets of Euclidean space and three-dimensional tori. In particular, we are interested in the connected components of the space of tilings of such regions under local moves. Building on the work of the third and fourth authors [19], we allow two possible local moves, the *flip* and *trit*. These moves are considered with respect to two topological invariants, the *twist* and *flux*.

Our main result proves that, up to refinement,

• Two tilings are connected by flips and trits if and only if they have the same flux.

• Two tilings are connected by flips alone if and only if they have the same flux and twist.

1 Introduction

Tiling problems have received much attention in the second half of the twentieth century: two-dimensional domino and lozenge tilings in particular, due to their connection to the dimer model and to matchings in a graph. A large number of techniques have been developed for solving various problems in two dimensions. For instance, Kasteleyn [13], Conway and Lagarias [5], Thurston [25], Cohn, Elkies, Jockush, Kuperberg, Larsen, Propp and Shor [12, 4, 7], and Kenyon, Okounkov and Sheffield [15, 14] have used very interesting techniques, ranging from abstract algebra to probability.

A number of generalizations of these techniques have been made to the three-dimensional case. Randall and Yngve [23] considered tilings of "Aztec" octahedral

2010 *Mathematics Subject Classification.* Primary 05B45; Secondary 52C20, 52C22, 05C70. *Keywords and phrases* Three-dimensional tilings, dominoes, dimers, flip accessibility, connectivity by local moves

and tetrahedral regions with triangular prisms, which generalize domino tilings to three dimensions. Linde, Moore and Nordahl [18] considered families of tilings that generalize rhombus (or lozenge) tilings to arbitrary dimensions. Bodini [2] considered tiling problems of pyramidal polycubes. And, while the problem of counting domino tilings is known to be computationally hard (see [22]), some asymptotic results, including for higher dimensions, date as far back as 1966 (see [11, 3, 10]).

Most relevant to the discussion in this paper are the problems of connectivity of the space of tilings under local moves. A *flip* is the simplest local move: remove two adjacent parallel dominoes and place them back in the only possible different position. In two dimensions, any two tilings of a simply connected region are flip connected; see e.g., [25, 24]. This is no longer the case when one considers tilings in three dimensions.

Even for simple three-dimensional regions, the space of tilings is no longer connected by flips. This is perhaps not surprising as the flip is inherently a two-dimensional move. The *trit* is a three-dimensional local move, which lifts three dominoes sitting in a $2 \times 2 \times 2$ cube, no two of which are parallel, and places them back in the only other possible configuration (see Figure 3). It is natural to ask if these two moves, the flip and trit are enough to connect all tilings of three-dimensional spaces. In general, the answer is again no. Here we consider connectivity of tilings taking into account two topological invariants. In doing so, we are able to characterize (up to refinement) when two tilings are connected by flips or flips and trits.

The first invariant is the *Flux* of a tiling. The flux of a tiling of a region R takes values in the first homology group $H_1(R; \mathbb{Z})$. The second invariant is the *twist* of a tiling. The twist assumes values either in $\mathbb{Z}$ or in $\mathbb{Z}/m\mathbb{Z}$ where m is a positive integer depending on the value of the Flux. For contractible regions (such as boxes), the twist assumes values in $\mathbb{Z}$. If R is a torus of the form $\mathbb{Z}^3/\mathcal{L}$, where $\mathcal{L}$ is spanned by $(a, 0, 0), (0, b, 0), (0, 0, c)$ (with a, b, c even positive integers) then the twist assumes values in $\mathbb{Z}$ if Flux is 0 and in some $\mathbb{Z}/m\mathbb{Z}$ otherwise. The twist was first introduced by Milet and Saldanha [19, 20] for particularly nice regions. In that context, the twist has a simple combinatorial definition. Unfortunately, it does not extend to the more general tiling domains considered here.

Our new definition of twist, and our introduction of Flux relies on the construction of auxiliary surfaces. The difficulty is that the required surfaces may not always exist. The difficulty is addressed by using the concept of *refinement*. A region R is refined by decomposing each cube of R into $5 \times 5 \times 5$ smaller cubes. Refinement guarantees the existence of auxiliary surfaces which, borrowing from knot theory, we call *Seifert surfaces*.

Informally, Flux measures how a tiling flows across a surface boundary. If two tilings are flip and trit connected, they must have equal Flux. The twist measures how "twisted" a tiling is by trits: under a trit move the twist changes by exactly one. A key property is that if two tilings are in the same flip connected component, then they must have equal twist. The converse is false in general.

The twist can also be interpreted as a discrete analogue of *helicity* arising in fluid mechanics and topological hydrodynamics, see e.g. [21, 1, 16]*. The helicity of a vector field on a domain in $\mathbb{R}^3$ is a measure of the self linkage of field lines. An important recent result shows that helicity is the only integral invariant of volume-preserving transformations [8].

Our main result is a characterization of the connectedness of three-dimensional tilings by flips and trits with respect to flux and twist.

Theorem 1. *Consider a cubiculated region R and two tilings t_0 and t_1 of R.*

(a) There exists a sequence of flips and trits taking a refinement of t_0 to a refinement of t_1 if and only if $\operatorname{Flux}(t_0) = \operatorname{Flux}(t_1)$.

(b) There exists a sequence of flips taking a refinement of t_0 to a refinement of t_1 if and only if $\operatorname{Flux}(t_0) = \operatorname{Flux}(t_1)$ *and* $\operatorname{Tw}(t_0) = \operatorname{Tw}(t_1)$.

In general, the refinement condition is necessary in the statement of the theorem. However, it is not known if the refinement condition may be dropped in certain special cases. For nice regions (such as boxes) there is empirical evidence, see [19, 9], that refinement is almost never necessary; for item (a) it may never be necessary.

Section 2 contains preliminaries for the regions we will consider. The two local moves, the flip and trit, are introduced in Section 3. Section 4 introduces the flux. In Sections 5 and 6 we work heavily with discrete surfaces leading to the definition of the twist in Section 7. Section 8 extends the concept of height functions to our setting where they are better described as height forms. Theorem 1 is proved in Section 9. We end with a discussion of further questions and conjectures concerning three-dimensional tilings.

The authors are thankful for the generous support of CNPq, CAPES, FAPERJ and a grant from the Brown-Brazil initiative.

2 Preliminaries

2.1 Cubiculated Regions

In this paper, we will consider tilings of certain three-dimensional regions. By a *cubiculated region R*, we will mean a cubical complex embedded as a finite

*The authors thank Yuliy Baryshnikov for bringing this concept to our attention.

polyhedron in $\mathbb{R}^N$, for some N, which is also a connected oriented topological manifold of dimension three with (possibly empty) boundary ∂R. We assume that: (i) interior edges of R are surrounded by precisely four cubes; (ii) cubes are painted *black* or *white* such that two adjacent cubes have opposite colors; and (iii) the number of black cubes equals the number of white cubes. It follows from this definition that ∂R is also a polyhedron and an oriented topological manifold of dimension two.

Example 2.1. A *box* is the cubiculated region $[0, L] \times [0, M] \times [0, N]$, where L, M and N are positive integers, at least one of them even. The *torus* $T = \mathbb{R}^3/\mathcal{L}$, where $\mathcal{L} \subset \mathbb{Z}^3$ is a three-dimensional lattice such that $(x, y, z) \in \mathcal{L}$ implies $x+y+z$ is even, is a cubiculated region without boundary.

We also consider the dual cubical complex $R^* \subset R$. Vertices of R^* are centers of cubes in R and edges of R^* join centers of adjacent cubes in R. There is a cube in R^* around each interior vertex of R: its eight vertices are the centers of the cubes in R adjacent to the interior vertex. The dual cubical complex R^* may or may not be a manifold, for instance, there may exist edges of R^* not adjacent to any cube of R^*.

We will also work with the graph $\mathcal{G}(R)$ and its dual $\mathcal{G}(R^*)$. The vertices and edges of $\mathcal{G}(R)$ are just the vertices and edges of R; in other words, $\mathcal{G}(R)$ is the 1-skeleton of R. Similarly, $\mathcal{G}(R^*)$ is the 1-skeleton of R^* which is a bipartite graph. Tiling regions are often simply regarded as subgraphs of $\mathcal{G}(R^*)$. It is important to note that the regions we are working with here must be topological manifolds, therefore it does not suffice to consider arbitrary subgraphs of the $\mathbb{Z}^3$ lattice.

A *domino* (or domino brick) is the union of two adjacent cubes in R and a *(domino) tiling* of R is a collection of dominoes with disjoint interior whose union is R. Thus, a tiling of R is equivalent to a matching of the graph $\mathcal{G}(R^*)$. When seen as an edge in $\mathcal{G}(R^*)$, a domino is called a *dimer*.

2.2 Cycles

An important concept to have in mind throughout this paper is the interpretation of the difference of two tilings as a union of disjoint cycles.

An *embedded cycle* is an injective continuous map $\gamma : \mathbb{S}^1 \to R^* \subset R$ whose image is a union of vertices and edges of R^*. Thus, an embedded cycle is a cycle in the graph theoretical sense for $\mathcal{G}(R^*)$.

We will also consider cycles homologically as elements of $Z_1(R^*; \mathbb{Z})$, the kernel of the boundary map from one to zero dimensional cells of R^*. Similarly, since a dimer connects a pair of vertices of opposite color, we may also think of dimers as oriented edges pointing from the center of a white cube to the center of a black cube (by convention), i.e., as generators of $C_1(R^*;\mathbb{Z})$, the one dimensional chain group.

With this point of view in mind, given two tilings t_0 and t_1, we define $t_1 - t_0$ to be the union of the dimers in both tilings, with the dimers in t_0 having their orientations reversed. Hence, $t_1 - t_0$ is the union of disjoint cycles; cycles of length 2 in the graph theoretical sense are called *trivial cycles* and are usually ignored. Figure 1 shows a simple example.

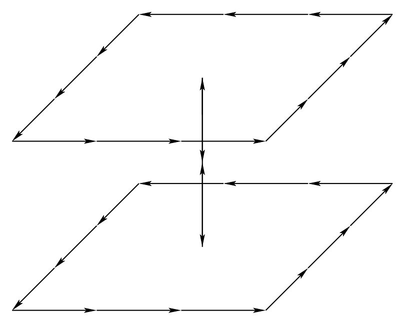

Figure 1: The two tilings in Figure 2 plotted together as a union of oriented dimers yielding three curves, one of which is trivial.

2.3 Refinements

A region R is *refined* by decomposing each cube of R into $5 \times 5 \times 5$ smaller cubes; the corners and the center are painted the same color as the original cube. This defines a new cubiculated region R'. We sometimes need to refine a region R not once but k times: we then call the resulting region $R^{(k)}$. As topological spaces, R and R' are equal. A tiling t is refined by decomposing each domino of t into $5 \times 5 \times 5$ smaller dominoes, each one parallel to the original domino. Again, this defines a new tiling t'; if we refine k times we obtain $t^{(k)}$.

Some comments on the choice of $5 \times 5 \times 5$ are in order. Dividing each cube into $2 \times 2 \times 2$ smaller cubes would erase the distinction between black and white and make the entire discussion trivial. Dividing into $3 \times 3 \times 3$ smaller cubes works for our purposes but the fact that the central cube in this small block has the opposite color as the corner is a source of unnecessary confusion. What we need, therefore, is a positive integer congruent to 1 mod 4; 5 being the smallest.

3 Local Moves: Flips and Trits

A *flip* is a move that takes a tiling t_0 into another tiling t_1 by removing two parallel dimers that form a $2 \times 2 \times 1$ “slab” and placing them back in the only other position possible. The *flip connected component* of a tiling t is the set of all tilings that can be reached from t after a sequence of flips. For instance, the $3 \times 3 \times 2$ box has 229 total tilings, but only three flip connected components. Two components contain just one tiling each; i.e. there are no possible flip moves from these tilings, see Figure 2.

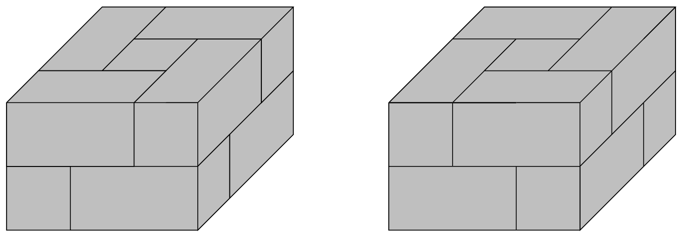

Figure 2: The two tilings of a $3 \times 3 \times 2$ box that have no possible flips. Note that the tilings differ only in that the “top” and “bottom” levels have been swapped.

A *trit* is a move involving three dominoes which sit inside a $2 \times 2 \times 2$ cube where each domino is parallel to a distinct axis. We thus necessarily have some rotation of Figure 3. The trit that takes the drawing at the left of Figure 3 to the drawing at the right is a *positive trit*. The reverse move is a *negative trit*. Notice that the $2 \times 2 \times 2$ cube need not be entirely contained in R. On the other hand, not just the six cubes directly involved in the trit but also at least one of the other two must be contained in R otherwise R is not a manifold.

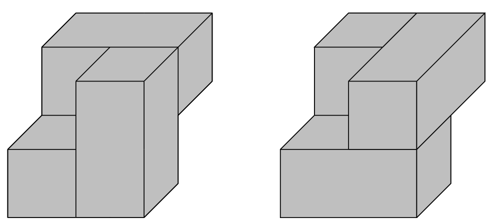

Figure 3: The anatomy of a positive trit (from left to right). The trit that takes the right drawing to the left one is a negative trit. The empty corners may represent either partial dimers that are not contained in the $2 \times 2 \times 2$ cube or cubes that are not contained in the region (for instance, if the region happens not to be a box).

Flips and trits behave well with respect to refinement. If t_0 and t_1 differ by a flip, then their refinements differ by a sequence of 125 flips. If t_0 and t_1 differ by a trit, then their refinements differ by a a trit and a sequence of flips. Therefore we have:

Proposition 3.1. *If t_0 and t_1 are connected by flips (resp. flips and trits) then their refinements are also connected by flips (resp. flips and trits).*

The converse does not hold however. Already in the $4 \times 4 \times 4$ box, there exist tilings t_0 and t_1 which are not mutually accessible by flips but such that $t_0^{(1)}$ and $t_1^{(1)}$ are.

4 Flux

In this section we define and develop the concept of *Flux* via homology theory, a related notion of *flux through a surface* will be give in Section 6.

Recall that if t_0 and t_1 are two tilings, then $t_1 - t_0 \in Z_1(R^*; \mathbb{Z})$, and the equivalence class $[t_1 - t_0]$ is an element of the homology group $H_1(R^*; \mathbb{Z})$ (which is naturally identified with $H_1(R; \mathbb{Z})$ since the inclusion $R^* \subset R$ is a homotopy equivalence).

Definition 4.1. *Let $t_\oplus$ be a fixed base tiling. The* Flux *of a tiling t is defined as:*

$$\mathrm{Flux}(t) = [t - t_\oplus] \in H_1(R^*; \mathbb{Z}).$$

Notice that if t_0 and t_1 differ by a flip then $t_1 - t_0$ is the boundary of a square. Similarly, if t_0 and t_1 differ by a trit then $t_1 - t_0$ is the boundary of a sum of three squares. In either case we have:

Proposition 4.2. *If t_0 and t_1 differ by flips and trits, then* $\mathrm{Flux}(t_0) = \mathrm{Flux}(t_1)$.

Again, the converse is not true. For $\mathcal{L} = 8\mathbb{Z}^3$, let R be the torus $\mathbb{R}^3/\mathcal{L}$. It is possible to construct two tilings of R for which the Flux is 0 but from at least one tiling no flip or trit is possible, see [9].

While the converse is not true, our main Theorem provides the correct converse to this statement – there exist refinements of the two tilings that are connected by flips and trits if and only if the Flux are equal. First, we show that the Flux is preserved under refinement.

Lemma 4.3. *If t' is the refinement of a tiling t, then* $\mathrm{Flux}(t') = \mathrm{Flux}(t)$.

Proof. An embedded cycle $\gamma : \mathbb{S}^1 \to R^*$ is refined by merely interpreting it as $\gamma' = \gamma : \mathbb{S}^1 \to (R')^*$. A tiling t of R is *tangent* to an embedded cycle if every vertex of R^* in the image $\gamma[\mathbb{S}^1]$ is one of the endpoints of a dimer d contained in the tiling t and in $\gamma[\mathbb{S}^1]$. Unfortunately, if t is tangent to γ it does not follow that t' is tangent to γ'. In this situation, we therefore define a *modified refinement* $t^{(1;\gamma)}$ which is tangent to γ. At each $5 \times 5 \times 5$ cube around a vertex of R^* belonging to the image of γ, perform a flip in t' if needed (as in Figure 4) to obtain the required tiling $t^{(1;\gamma)}$ tangent to γ'. We can iterate this procedure to define $t^{(k;\gamma)}$ which is tangent to $\gamma^{(k)}$ and connected to $t^{(k)}$ by flips in the neighborhood of γ. Similarly,

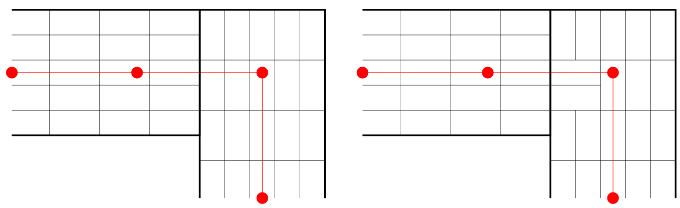

Figure 4: Performing flips after a refinement in order to obtain a tiling tangent to the boundary cycle.

consider two tilings t_0 and t_1 of R and let γ be the system of cycles $t_1 - t_0$. Then $t_0^{(k;\gamma)}$ and $t_1^{(k;\gamma)}$ are both tangent to $\gamma^{(k)}$. Moreover, a sequence of flips in the neighborhood of γ yields new tilings $t_0^{[k;\gamma]}$ and $t_1^{[k;\gamma]}$ such that $t_1^{[k;\gamma]} - t_0^{[k;\gamma]} = \gamma^{(k)}$. If $t_\oplus$ is the base tiling of R, take $t'_\oplus$ to be the base tiling of R'. It then follows from the construction of $t_1^{[k;\gamma]}$ that $\mathrm{Flux}(t') = \mathrm{Flux}(t)$. $\square$

5 Surfaces

In order to better understand the Flux, and to define the Twist, we will work heavily with discrete surfaces. Consider a cubiculated region R and its dual R^*. An *embedded discrete surface* in R^* is a pair (S, ψ) where:

- S is an oriented topological surface with (possibly empty) boundary ∂S;
- $\psi : S \to R^* \subset R$ is an injective continuous map whose image is a union of vertices, edges and squares of R^*.

We will sometimes abuse notation using S to refer to the domain surface S, the image $\psi[S]$ and the element of $C_2(R^*; \mathbb{Z})$ obtained by adding the squares in $\psi[S]$ (with the orientation given by S and ψ).

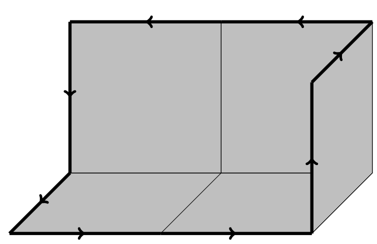

Figure 5: An embedded discrete surface which consists of five squares. The thicker line represents the oriented boundary of the surface.

Figure 5 shows a simple embedded discrete surface. Since both S and R are oriented, this defines an orientation transversal to S. In other words, at each square of S there is a well defined normal vector.

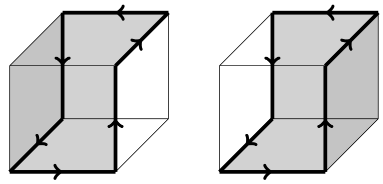

Figure 6: Two possible Seifert surfaces.

Consider a tiling t_0 of a region R and an embedded discrete surface S. The tiling and the surface are *tangent at the boundary* if t_0 and the system of cycles $\psi|_{\partial S}$ are tangent. In particular, if $\partial S = \emptyset$ then any tiling is tangent at the boundary to S. The tiling and the surface are *tangent* if they are tangent at the boundary and, furthermore, every vertex of R^* in S is an endpoint of a dimer in t_0 and S.

Refinements are important throughout the paper and some additional remarks on the subject are in order. Consider a tiling t_0 of R tangent to S. The refinement t'_0 of t_0 is not tangent to S but, as in Figure4, a few flips are sufficient to go from t'_0 to a tiling t_1 which is tangent to S. We will pay little attention to the distinction between t'_0 and t_1 and speak of t_1 as the refinement of t_0.

We shall also have to consider tilings t_0 of R which are tangent to ∂S but which cross S (that is, fail to be tangent to S). If t_0 crosses S at m vertices, then the refinement t'_0 crosses S at many more points (between $9m$ and $25m$). Again, a few flips take t'_0 to t_1 which crosses S at the original m points only; we often simply call t_1 the refinement of t_0, slightly abusing notation.

Definition 5.1. *A* (discrete) Seifert surface *for a pair of tilings* (t_0, t_1) *of a region R is a connected embedded discrete surface S where the restriction $\psi|_{\partial S}$ is the collection of nontrivial cycles of* $t_1 - t_0$*.*

By definition, both t_0 and t_1 are tangent at the boundary to S. The tiling t_0 is tangent to S if and only if t_1 is. As above, if S is a Seifert surface for (t_0, t_1) then its refinement S' is a Seifert surface for the pair of refinements (t'_0, t'_1).

Example 5.2. If t_0 and t_1 differ by a flip, a unit square is a valid Seifert surface for the pair (the simplest surface, but not the only one). If t_0 is obtained from t_1 after a single positive trit, we may assume that the situation is, perhaps after some rotation, as portrayed in Figure 6. Note that in order to build the surfaces portrayed in Figure 6, we need that the interior point of the surface in either case is a center of a cube in R. This condition is satisfied in at least one of the cases since ∂R is a manifold.

It follows from homology theory that in order for a discrete Seifert surface to exist, we must have $\mathrm{Flux}(t_0) = \mathrm{Flux}(t_1)$. The converse is not true: in the example of Figure 1 there exists a disconnected surface (two disks) and a connected surface which does not respect orientation (a cylinder) but no discrete Seifert surface for the pair. This is one of several occasions when taking refinements solves our difficulties.

Lemma 5.3. *Consider a cubiculated region R and two tilings t_0 and t_1 of R. If* $\mathrm{Flux}(t_0) = \mathrm{Flux}(t_1)$ *then for sufficiently large $k \in \mathbb{N}$ there exists a discrete Seifert surface in $(R^{(k)})^*$ for the pair (t_0, t_1).*

First, we need the following purely topological lemma. The statement is well-known for sufficiently nice regions, see e.g. [17]. As we were unable to find a proof at our level of generality, we include one here for completeness.

Lemma 5.4. *Consider a cubiculated region R and two tilings t_0 and t_1 of R. If* $\mathrm{Flux}(t_0) = \mathrm{Flux}(t_1)$ *then there exists a smooth Seifert surface in $(R^{(1)})^*$ for the pair (t_0, t_1).*

Proof. To simplify notation, write L for the difference $t_1 - t_0$. The refinement guarantees that L is contained in the interior of R. The hypothesis $\mathrm{Flux}(t_0) = \mathrm{Flux}(t_1)$ guarantees that L is a boundary, i.e., that there exists s in $C_2(R^{(1)})^*$ with $\partial(s) = L$.

For each vertex v of R, construct a small open ball b_v around v. For each edge e of R, construct a thin open cylinder c_e around e. Let $R_0 = R \setminus \{b_v \cup c_e\}$, R minus the union of all b_v and c_e. Let $R_1 = R \setminus \{b_v\}$, R minus the union of all b_v. Thus $R_0 \subset R_1 \subset R$. We construct a smooth Seifert surface in three stages: first in R_0, then extend it to R_1 and finally to R.

For R_0, we consider each square a of R and its coefficient s_a in s. Orient the square a so that $s_a \geq 0$. Construct S_0 in R_0 by taking s_a translated copies of a, with boundary falling outside R_0.

Next consider each edge e. There are two possibilities: e may or may not belong to the support of L. First assume it does not. Examine the boundary of c_e: we see a number of line segments (the intersection of the squares in S_0 with the boundary of c_e). This can be described by a family of $2k$ points in a circle, k positive and k negative. It is possible to match positive and negative points and draw curves joining them so that the curves do not cross. Indeed, by induction, take two adjacent points, one positive and one negative, and join them by a curve near the circle. For the other points construct segments taking to a smaller circle. Now use induction. Take the Cartesian product of these curves by e to construct a surface in R_1.

Now consider the case where e belongs to the support of L. Similar to above, we have k positive points on a circle, $(k + 1)$ negative points on the circle and a positive origin. The same inductive proof constructs smooth disjoint curves joining the points (the last negative point on the circle will be connected to the center). Again, take the Cartesian product of these curves with e to construct a surface in R_1. This completes the construction of S_1 in R_1.

We now extend the surface to R, that is, we extend it to each ball b_v. Again we consider two cases: v in the support of L and v not in support of L. First consider v not in the support of L. Examine the boundary of b_v, a sphere S_v. Our previous construction obtains a family of disjoint oriented simple closed curves in S_v. Attach disjoint disks contained in b_v with these curves as boundary: Take a point in S_v not in the cycles and call it infinity, so that S_v is identified with the plane. Cycles are now nested: start with an innermost cycle and close it; proceed along cycles to the outermost.

Next, consider the case where v is in the support of L. Again examine the sphere S_v. We have a family of disjoint oriented simple curves: one of them is a segment (with two endpoints), the others are closed (cycles). Take the "point at infinity" in S_v very near the segment so that in the plane the segment is outside the cycles. Close cycles from inner to outermost.

Finally, consider the segment. Notice it can be long, perhaps going several times around a face of the octahedron formed on the sphere by the adjacent cubes. Take a smooth 1-parameter family of diffeomorphisms keeping the endpoints of the segment fixed and taking the segment to a geodesic on the sphere (this may involve rotations around endpoints). Apply this family of diffeomorphisms on spheres of decreasing radii and complete the surface with a plane near the vertex. This completes the construction of S. □

Proof of Lemma 5.3. Construct a smooth Seifert surface $\psi_\infty : S \to R$ (as in Lemma 5.4). A sufficiently large value of k allows for an approximation ψ of ψ_∞ such that ψ is a discrete Seifert surface as follows:

Consider a smooth Seifert surface S as in Lemma 5.4. Let K be the maximum sectional curvatures of S, up to and including the boundary. Take n such that $K < 5^{(n-1)}$ and refine n times so that the radii of curvature at any point is always more than twice the diagonal of any cube. Now classify cubes near the surface (meaning with center at a distance $< 5/2$ from the surface, thus including cubes crossing the surface) as *above* or *below* S according to the orientation of S and the measure of the cube on each side of S. The tiled surface S lies between cubes which are *above* and cubes *below* (and therefore closely approximates S). The curvature estimate implies the good behavior of S. □

6 Flux through surfaces

We now introduce *flux through surfaces*. We relate this notion with the Flux(t) at the end of the section. The interpretation here provides motivation for our choice of terminology; we think of tilings as flowing through regions and across surfaces.

Let t be a tiling and S an embedded discrete surface such that t is tangent to S at the boundary. For each vertex v of R^* in the interior of S consider the only dimer d of t adjacent to v. Draw a small vector along d starting from v: its endpoint may be on S (if d is contained in S), above, or below S as determined by the normal vectors to squares of S.

Definition 6.1. *For a surface S and a tiling t, define the* flux of t through S, *$\phi(t;S) \in \mathbb{Z}$, as follows. Set*

$$\varphi(v;t;S) = \operatorname{color}(v) \cdot \begin{cases} +1, & \textit{endpoint above } \psi[S]; \\ 0, & \textit{endpoint on } \psi[S]; \\ -1, & \textit{endpoint below } \psi[S]; \end{cases} \qquad \phi(t;S) = \sum_v \varphi(v;t;S),$$

where the color of v is $+1$ if v corresponds to a black tile and -1 if v corresponds to a white tile.

Example 6.2. In Figure 7, we see four dimers intersecting the interior of a surface. The horizontal dimer completely contained in the surface will contribute 0 to the flux.

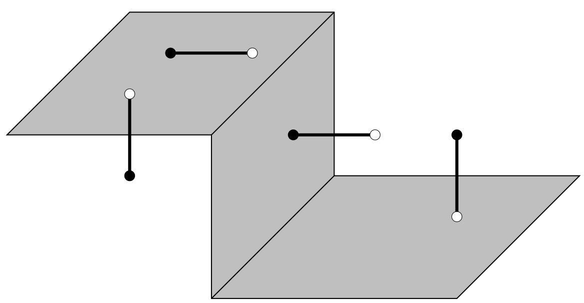

Figure 7: An example of flow through a surface. Only a few representative dimers have been shown; one sitting strictly in the interior of the surface, one above and one below.

The next Theorem shows that with the definition above, the flux of a tiling through the boundary of a manifold is always zero.

Theorem 2. *Let R be a cubiculated region and t a tiling of R. Let $\psi : S \to R^*$ be an embedded discrete surface with $\partial S = \emptyset$. Assume there exists a topological manifold $M_1 \subset R^*$ with $\partial M_1 = S$. Then, $\phi(t;S) = 0$.*

Proof. We start by showing the following enumerative result. Let b_{int} and w_{int} be the number of black and white vertices of R^* in the interior of M_1; let b_∂ and w_∂ be the number of black and white vertices of R^* on S. Then $2b_{\text{int}} + b_\partial = 2w_{\text{int}} + w_\partial$.

To prove this claim, let M_2 be a copy of M_1 with reversed orientation. Glue M_1 and M_2 along S to define a topological 3-manifold M. The manifold M is oriented, has empty boundary, and inherits from M_1 and M_2 a cell decomposition, with vertices painted black and white. Let $f_i(M)$ be the number of faces of dimension i in this cell decomposition. Since the Euler characteristic of M equals 0,

$$f_0(M) - f_1(M) + f_2(M) - f_3(M) = 0.$$

Also, $f_2(M) = 3f_3(M)$: since M has no boundary, each face must be shared by exactly two cubes and each cube has six faces.

Our enumerative claim is equivalent to saying that the number of black vertices of M equals the number of white vertices. In order to see this, we first build another 3-complex T for the topological manifold M in the following way: the vertices of T are the white vertices of M; the edges of T are the diagonals that connect white vertices in each (two-dimensional) face of M; the two-dimensional faces of T are the four triangles in each cube that form a regular tetrahedron with its four white vertices; finally, the three-dimensional faces of T come in two flavors: the regular tetrahedrons inside each cube, and cells around black vertices of M (these are regular octahedra if the vertex is in the interior of either M_i but may have some other shape if the vertex belongs to S).

We have $f_0(T) = 2w_{\text{int}} + w_\partial$, $f_1(T) = f_2(M)$ (each face of M contains exactly one edge of T), $f_2(T) = 4f_3(M)$ (each cube of M contains exactly four faces of T) and $f_3(T) = f_3(M) + 2b_{\text{int}} + b_\partial$ (one tetrahedron inside each cube, one other cell around each black vertex). Thus

$$\begin{aligned}0 &= f_0(T) - f_1(T) + f_2(T) - f_3(T)\\ &= (2w_{\text{int}} + w_\partial - 2b_{\text{int}} - b_\partial) + (3f_3(M) - f_2(M)) = 2w_{\text{int}} + w_\partial - 2b_{\text{int}} - b_\partial.\end{aligned}$$

Now, each black vertex on S must match either a white vertex in the interior of M_1, on $\partial M_1 = S$, or in the exterior of M_1: let b_1, b_2 and b_3 be the number of vertices of each kind so that $b_1 + b_2 + b_3 = b_\partial$. Define w_1, w_2 and w_3 similarly, so that $w_1+w_2+w_3 = w_\partial$. Clearly, $b_2 = w_2$. By definition, $\phi(t; S) = b_3-b_1-w_3+w_1$. Counting vertices in the interior of M_1 gives $b_1 - w_1 = w_{\text{int}} - b_{\text{int}}$ so that $\phi(t; S) = (b_\partial - w_\partial) - 2(b_1 - w_1) = 0$. □

The remainder of the section develops that the flux is not dependent on a precise surface, but can be defined in terms of homology classes. If S_0 and S_1 are oriented surfaces with $\partial S_0 = \partial S_1$ then let $S_1 - S_0$ denote the surface obtained by gluing S_1 and S_0 along the boundary and reverting the orientation of S_0. Furthermore, if $\psi_i : S_i \to R$ are continuous maps, let $\psi_1 - \psi_0 : S_1 - S_0 \to R$ denote the map defined by $(\psi_1 - \psi_0)(p) = \psi_i(p)$ if $p \in S_i$.

The image, $(\psi_1-\psi_0)[S_1-S_0]$, can be seen as an element of $H_2(R)$; the maps ψ_0 and ψ_1 are *homological* if $(\psi_1 - \psi_0)[S_1 - S_0] = 0 \in H_2(R)$. Note that if S_0 and S_1

are smooth or topological oriented surfaces with $\partial S_0 = \partial S_1$ and $\psi_i : S_i \to R$ are smooth or topological embeddings such that there exists a topological manifold $M_1 \subset R$ for which $\psi_1 - \psi_0 : S_1 - S_0 \to \partial M_1 \subset R$ is an orientation preserving homeomorphism then ψ_0 and ψ_1 are homological. On the other hand, if S_0 and S_1 are smooth or topological oriented surfaces with $\partial S_0 = \partial S_1$ then the map $\psi_1 - \psi_0 : S_1 - S_0 \to R$ is usually not an embedded surface.

Lemma 6.3. *Let R be a cubiculated region and t be a tiling of R. Let S_0 and S_1 be oriented surfaces with $\partial S_0 = \partial S_1$. Let $\psi_i : S_i \to R^*$ be embedded discrete surfaces with $(\psi_0)|_{\partial S_0} = (\psi_1)|_{\partial S_1}$. If t is tangent to $(\psi_i)|_{\partial S_i}$ and ψ_0 and ψ_1 are homological then $\phi(t, S_0) = \phi(t, S_1)$.*

We first construct a function ω, the winding number, taking integer values on vertices of R (and therefore also cubes of R^*). Consider v_0, v_1 vertices of R: we first show how to compute $\omega(v_1) - \omega(v_0)$. Consider a simple path γ along edges of R going from v_0 to v_1 (such a path exists since R is assumed to be connected). Count intersections of γ with (the image of) $\psi_1 - \psi_0$. Notice that γ intersects $\psi_1 - \psi_0$ at the centers of oriented squares: each intersection counts as $+1$ (resp. -1) if the tangent vector to γ coincides (resp. or not) with the normal vector to the square in $\psi_1 - \psi_0$. This total is $\omega(v_1) - \omega(v_0)$.

Notice that the value of $\omega(v_1) - \omega(v_0)$ does not depend on the choice of the path γ. Indeed, take two such paths γ_0 and γ_1 and concatenate them to obtain a closed path $\gamma_1 - \gamma_0$. Counting intersections with $\gamma_1 - \gamma_0$ as described in the previous paragraph defined a linear map from $C_2(R^*)$ to $\mathbb{Z}$ and therefore an element of $C^2(R^*)$ which is easily seen to be in $Z^2(R^*)$. Since $\psi_1 - \psi_0$ is assumed to belong to $B_2(R^*)$ their product must equal 0, yielding independence from path. A similar argument shows that if v_0 and v_1 both belong to the boundary ∂R then counting intersections with γ also defines an element of $Z^2(R^*)$ and therefore $\omega(v_1) - \omega(v_0) = 0$. We may therefore define w so that if $v \in \partial R$ then $\omega(v) = 0$; if $\partial R = \emptyset$ we have a degree of freeedom here and we choose w so that it assumes the value 0 somewhere.

Proof of Lemma 6.3. Our proof works by induction of $c = |\max \omega| + |\min \omega|$. If $c = 0$ then the surfaces ψ_0 and ψ_1 coincide and we are done. Let us consider the case $c = 1$; without loss of generality, ω assumes the values 0 and 1. Let $M_1 \subset R^*$ be the union of closed cubes centered at vertices v (of R) with $\omega(v) = 1$. We would like M_1 to be a 3D manifold with boundary. Unfortunately, that is not guaranteed. But this can easily be fixed. Start by refining R (and the surfaces), so that now M_1 is a union of $5 \times 5 \times 5$ blocks of cubes. If M_1 is *not* a manifold, this means thare are bad edges (two alternate blocks present, two absent) or bad vertices (more than one undesirable pattern). First fix the vertices by adding extra cubes; then the edges. This corresponds to constructing a chain of auxiliary surfaces

$$\tilde{\psi}_0 = \psi_0, \tilde{\psi}_1, \tilde{\psi}_2, \ldots, \tilde{\psi}_{N-1}, \tilde{\psi}_N = \psi_1$$

such that for any k the pair $\tilde{\psi}_{k-1}$, $\tilde{\psi}_k$ satisfies $c = 1$ and M_1 a manifold.

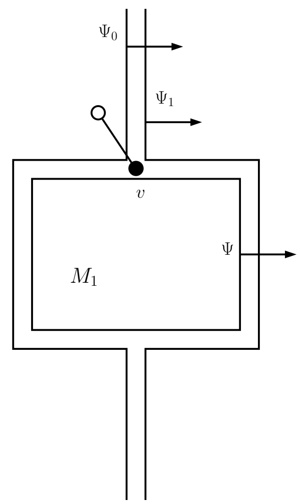

Figure 8: Auxiliary surfaces as in the proof of Lemma 6.3.

We are therefore left with the case where $\psi_1 - \psi_0$ is the boundary of a manifold M_1. It now appears that this case follows from Theorem 2. This is true but not as trivial as it may seem at first. As in Figure 8, let $S = \partial M_1$ and ψ be a parametrization of S. What Theorem 2 tells us is that $\phi(t; S) = 0$ but what we need is that $\phi(t; S_1) - \phi(t; S_0) = 0$. In other words, we need

$$\sum_v D(v) = 0; \qquad D(v) := \varphi(v; t; S_1) - \varphi(v; t; S_0) - \varphi(v; t; S). \tag{1}$$

If v belongs to at most two surfaces then $D(v) = 0$. On the other hand, for v as in Figure 8, $D(v) = -1$. Let Γ be the curve along which the three surfaces S_0, S_1 and S meet. A case by case analysis shows that for $v \in \Gamma$ we have $D(v) = 0$ if v's partner also belongs to Γ and $D(v) = -\operatorname{color}(v)$ otherwise. Since Γ is balanced this proves equation 1 and completes the proof of this case.

The general inductive step is now similar. Otherwise assume without loss of generality that $\max \omega = l > 0$. Let M_l be the union of cubes of R^* with center v with $\omega(v) = l$. As above, we may assume M_l to be a 3D manifold. Thus, M_l is a 3D manifold and its boundary ∂M_l consists of subsets of ψ_0 and ψ_1 meeting at curves (in the simplest example, M_l is a ball, the two subsets are disks meeting at a circle; this may get significantly more complicated but does not affect our argument). Modify ψ_0 to define ψ_2 by discarding the subset of ψ_0 in ∂M_l and attaching instead the subset of ψ_1 also in ∂M_l. The two surfaces ψ_0 and ψ_2

are homological by construction and the difference $\psi_2 - \psi_0$ is the boundary of M_1. The case $c = 1$ above shows that $\phi(t, S_0) = \phi(t, S_2)$. For the new winding number $\tilde{\omega}$ (defined for the pair ψ_1 and ψ_2) we have $\tilde{c} = c - 1$: by induction we are done. □

Lemma 6.4. *Consider a cubiculated region R, a tiling t of R, and an element $a \in H_2(R; \mathbb{Z})$.*

1. *There exist a nonnegative integer k and an embedded discrete surface $\psi : S \to (R^{(k)})^*$ such that $\partial S = \emptyset$ and $\psi[S] = a$.*

2. *Let k_0 and k_1 be nonnegative integers. Let $\psi_0 : S_0 \to (R^{(k_0)})^*$ and $\psi_1 : S_1 \to (R^{(k_1)})^*$ be embedded discrete surfaces such that $\partial S_0 = \partial S_1 = \emptyset$. Assume that $\psi_0[S] = \psi_1[S] = a$. Then $\phi(t^{(k_0)}; S_0) = \phi(t^{(k_1)}; S_1)$.*

Proof. As in the proof of Lemma 5.3, for any $a \in H_2(R; \mathbb{Z})$, there exists a smooth embedded surface $\psi_\infty : S \to R$ with $\psi_\infty[S] = a$ (where S is a smooth closed surface with $\partial S = \emptyset$). For sufficiently large k there is an embedded discrete surface $\psi : S \to R$ approximating ψ_∞ so that $\psi[S] = a$. The equality in item (2) follows from Lemma 6.3. □

Using this lemma, for $a \in H_2(R; \mathbb{Z})$, define $\phi(t; a)$ to be equal to $\phi(t^{(k)}; S)$ for any embedded discrete surface $\psi : S \to (R^{(k)})^*$ such that $\psi[S] = a$ (as an element of $H_2(R; \mathbb{Z})$).

Note that if t_0 and t_1 differ by a flip or trit then $\phi(t_0; a) = \phi(t_1; a)$.

Lemma 6.5. *Consider a cubiculated region R and tilings t_0, t_1 of R. If* $\mathrm{Flux}(t_0) = \mathrm{Flux}(t_1)$ *then $\phi(t_0; a) = \phi(t_1; a)$ for all $a \in H_2(R; \mathbb{Z})$.*

Proof. Let S be a discrete embedded surface in some refinement $R^{(k)}$ with $\partial S = \emptyset$. Adding $\phi(t_0; S) + \phi(t_1; S)$ gives a linear map from $C_1(R^*; \mathbb{Z})$ to $\mathbb{Z}$. Boundaries of squares are taken to 0 and therefore $B_1(R^*; \mathbb{Z})$ is contained in the kernel of this map. Hence we have a map from $H_1(R^*; \mathbb{Z})$ to $\mathbb{Z}$. By hypothesis $[t_1 - t_0] = 0 \in H_1$ and therefore t_1 and t_0 are taken to the same number. In other words, $\phi(t_0; S) = \phi(t_1; S)$. Since this holds for all S, $\phi(t_0; a) = \phi(t_1; a)$. □

The converse does not hold. For example, let $\mathcal{L}$ be spanned by the vector $(0, 0, 4)$; let $R \subset (\mathbb{R}^3/\mathcal{L})$ be the set of points (x, y, z) for which $0 \le x, y \le 4$. Then, one can construct tilings of R with different values of the Flux but $H_2(R) = 0$ and therefore $\phi(t; a)$ is always trivial [9].

Most importantly, it follows from Lemma 6.5 that the next definition is well-defined and can be seen as a function of the $\mathrm{Flux}(t)$ rather than t.

Definition 6.6. *Define the* modulus *of a tiling as*

$$m = \mu(\text{Flux}(t)) = \gcd_{a \in H_2} \phi(t; a);$$

so that for all $a \in H_2$ we have $\phi(t; a) \equiv 0 \pmod{m}$.

7 Twist

In this section, we define our second topological invariant of tilings, the twist. The twist was first introduced in [19]. There it has a very nice combinatorial definition but the construction involved is not well defined at the level of generality of this paper. Here we give an alternate definition of twist involving embedded surfaces. If the Flux is zero, the twist assumes values in $\mathbb{Z}$. Otherwise, the twist assumes values in $\mathbb{Z}/m\mathbb{Z}$ where m is the modulus of the tiling, as defined at the end of the previous section. Intuitively, the twist records how "twisted" a tiling is by trits; the value of the twist changes by exactly 1 after a trit move.

We start by defining the *flux around a curve*. Consider a cubiculated region R, a tiling t of R, and $m = \mu(\text{Flux}(t))$. If $\gamma : \mathbb{S}^1 \to R^*$ is an embedded cycle such that t is tangent to γ, then Lemmas 5.3 and 6.5 imply that there exists a nonnegative integer k and an embedded surface $\psi : S \to (R^{(k)})^*$ such that $\psi|_{\partial S} = \gamma^{(k)}$. Furthermore, if k_0 and k_1 are nonnegative integers and $\psi_0 : S_0 \to (R^{(k_0)})^*$ and $\psi_1 : S_1 \to (R^{(k_1)})^*$ are embedded surfaces such that $(\psi_i)|_{\partial S_i} = \gamma^{(k_i)}$, then $\phi(t^{(k_0;\gamma)}; S_0) = \phi(t^{(k_1;\gamma)}; S_1)$ (as elements of $\mathbb{Z}/m\mathbb{Z}$). These observations allow us to give the following definition.

Definition 7.1. *For a tiling t and a curve γ, define $\phi(t;\gamma) \in \mathbb{Z}/m\mathbb{Z}$, the* flux of t around γ, *to be*

$$\phi(t;\gamma) := \phi(t^{(k;\gamma)}; S) \in \mathbb{Z}/m\mathbb{Z}$$

for any surface $\psi : S \to (R^{(k)})^*$ such that $\psi|_{\partial S} = \gamma^{(k)}$.

Using the flux of a tiling around a curve, we may define our first notion of twist; the twist for a pair of tilings.

Definition 7.2. *Let R be a cubiculated region. Let tilings t_0 and t_1 be two tilings of R such that* $\text{Flux}(t_0) = \text{Flux}(t_1)$*. Then the twist of t_1 with respect to t_0 is defined as*

$$\text{TW}(t_1; t_0) := \phi(t_1; t_1 - t_0) = \phi(t_0; t_1 - t_0) \in \mathbb{Z}/m\mathbb{Z}.$$

Our larger goal is to define the twist of a single tiling, $\text{Tw}(t)$. In particular, the twist should satisfy $\text{TW}(t_1; t_0) = \text{Tw}(t_1) - \text{Tw}(t_0)$ so that $\text{Tw}(t)$ can be defined using a base tiling and the twist of a pair. To this end, first consider the result of a flip move.

Proposition 7.3. *Let R be a cubiculated region. Let t_0, t_1 and t_2 be tilings of R such that* $\mathrm{Flux}(t_0) = \mathrm{Flux}(t_1)$ *and t_1 and t_2 differ by a flip. Then*

$$\mathrm{TW}(t_1; t_2) = 0 \qquad \text{and} \qquad \mathrm{TW}(t_1; t_0) = \mathrm{TW}(t_2; t_0).$$

Proof. For the first equation simply take the surface for $t_1 - t_2$ to be a single square.

For the second equation, first consider the case where the flip $t_2 - t_1$ is disjoint from the system of cycles $t_1 - t_0$. Then we may take a surface S_1 for $t_1 - t_0$ which is also disjoint from the flip $t_2 - t_1$. Take S_2 to be the disjoint union of S_1 with the square with boundary $t_2 - t_1$. For the general case, we have to consider the position of the single square with respect to the system of cycles $t_1 - t_0$. A case by case analysis shows that suitable surfaces can always be constructed. □

Let γ_1 and γ_2 be disjoint systems of smooth cycles in an oriented 3-manifold R. If $[\gamma_1] = [\gamma_2] = 0 \in H_1(R)$, then there exist Seifert surfaces S_1 and S_2 for γ_1 and γ_2. Classically, the *linking number* $\mathrm{Link}(\gamma_1; \gamma_2) = \mathrm{Link}(\gamma_2; \gamma_1) \in \mathbb{Z}$ of γ_1 and γ_2 is defined as the number of intersections (with sign) between γ_1 and S_2 (or γ_2 and S_1). Furthermore, the linking number is independent of the choice of S_1 and S_2. For our more general spaces, the linking number must be considered with respect to the modulus m of the tiling. Then, the linking number quantifies the difference in twist. Namely, suppose R is a cubiculated region and t_0, t_1, t_2 and t_3 are tilings of R with equal Flux. If the systems of cycles $\gamma_1 = t_1 - t_0 = t_3 - t_2$ and $\gamma_2 = t_2 - t_0 = t_3 - t_1$ are disjoint then

$$\mathrm{TW}(t_3; t_2) - \mathrm{TW}(t_1; t_0) = \mathrm{TW}(t_3; t_1) - \mathrm{TW}(t_2; t_0) = 2\,\mathrm{Link}(\gamma_1; \gamma_2)$$

and $\mathrm{TW}(t_3; t_0) = \mathrm{TW}(t_3; t_2) + \mathrm{TW}(t_2; t_0) = \mathrm{TW}(t_3; t_1) + \mathrm{TW}(t_1; t_0)$. More generally, if $t_2 - t_1$ and $t_1 - t_0$ are not disjoint, then refine and slightly move these systems of cycles using flips (and Proposition 7.3) to obtain disjoint cycles. Together this gives the following.

Proposition 7.4. *Let R be a cubiculated region. Let t_0, t_1 and t_2 be tilings of R with* $\mathrm{Flux}(t_0) = \mathrm{Flux}(t_1) = \mathrm{Flux}(t_2)$. *Then* $\mathrm{TW}(t_2; t_0) = \mathrm{TW}(t_2; t_1) + \mathrm{TW}(t_1; t_0)$.

We are now ready to define the twist of a tiling.

Definition 7.5. *Let R be a cubiculated region. For any possible value Φ of the Flux of a tiling, choose a base tiling t_Φ. For a tiling t of a region R with* $\mathrm{Flux}(t) = \Phi$ *define*

$$\mathrm{Tw}(t) := \phi(t; t - t_\Phi) = \phi(t_\Phi; t - t_\Phi) \in \mathbb{Z}/m\mathbb{Z}.$$

Corollary 7.6. *If t_1 is obtained from t_0 by a positive trit then* $\mathrm{Flux}(t_1) = \mathrm{Flux}(t_0)$ *and* $\mathrm{Tw}(t_1) = \mathrm{Tw}(t_0) + 1$.

8 Height Functions

In this section we consider discrete Seifert surfaces together with the restriction of $\mathcal{G}(R^*)$ to the surface. We prove that, under suitable hypothesis, two tilings of such a surface are connected by flips. The proof relies on a development of *height functions* appropriate to our setting. Height functions are a standard tool in the study of domino tilings, see e.g. [25], [5], [6].

If S is a discrete Seifert surface in R, a tiling of R which is tangent to S restricts to a tiling of S, i.e., a matching of $\mathcal{G}^*$, the restriction of $\mathcal{G}(R^*)$ to S. Let $\mathcal{T} = \mathcal{T}(S, \mathcal{G}^*)$ be the set of all tilings of S.

Just as for 3-dimensional dominoes, if $t_0, t_1 \in \mathcal{T}$, the difference $t_1 - t_0$ can be seen as a system of cycles in S. We interpret $[t_1 - t_0]$ to be an element of $H_1(S)$ and, given a choice of a base tiling $t_\oplus \in \mathcal{T}$, set $\operatorname{Flux}_S(t) = [t - t_\oplus] \in H_1(S)$. Given $a \in H_1$, let $\mathcal{T}_a$ be the equivalence class of tilings of S with the same Flux a. A tiling $t \in \mathcal{T}$ is *stable* if every edge of $\mathcal{G}^*$ belongs to some tiling in the equivalence class $\mathcal{T}_{\operatorname{Flux}_S(t)}$; in this case, we also call the set $\mathcal{T}_{\operatorname{Flux}_S(t)}$ *stable*.

When taking refinements, we may use $t'_\oplus$ as a base tiling of refined S (after a few flips to make sure the tiling is tangent to S). It is then easy to see that $\operatorname{Flux}_S(t') = \operatorname{Flux}_S(t)$ (we will blur the distinction between a surface S and its refinement S'). Furthermore, sufficient refinement makes any class stable: the condition of existence of a suitable tiling becomes easy after refinement.

Let $a = \operatorname{Flux}_S(t_0) = \operatorname{Flux}_S(t_1)$ so that $t_0, t_1 \in \mathcal{T}_a$. Let V be the set of components of $S \smallsetminus \mathcal{G}$ (squares) with one extra object called ∞ corresponding to ∂S. For each oriented edge of $\mathcal{G}$, there is an element $e_l \in V$ to its left and an element $e_r \in V$ to its right; if the edge is contained in ∂S then one of these is ∞. Two elements $v_0, v_1 \in V$ are *neighbors* if there exists an oriented edge e with $v_0 = e_l$ and $v_1 = e_r$.

Let C_2 be the $\mathbb{Z}$-module of functions $w : V \to \mathbb{Z}$. Let C_1 be the $\mathbb{Z}$-module spanned by oriented edges of $\mathcal{G}^*$. Define the boundary map $\partial : C_2 \to C_1$ as follows: given $w \in C_2$ and e an oriented edge of $\mathcal{G}^*$, the coefficient of e in ∂w is $w(e_l) - w(e_r)$. Let $B_1 \subseteq C_1$ be the image of $\partial : C_2 \to C_1$: given $g \in B_1$, there is a unique element $w = \operatorname{wind}(g) \in C_2$ with $w(\infty) = 0$ and $\partial w = g$. We call $\operatorname{wind}(g)$ the *winding* of g.

Given two tilings $t, \tilde{t} \in \mathcal{T}_a$, we have $t - \tilde{t} \in B_1$. The function $w = \operatorname{wind}(t - \tilde{t})$ satisfies $w(\infty) = 0$. Furthermore, given a black-to-white edge e, we have $w(e_l) - w(e_r) = [e \in t] - [e \in \tilde{t}]$ (we use Iverson notation: $[e \in t]$ equals 1 if $e \in t$ and 0 otherwise). In particular, if $v_0, v_1 \in V$ are neighbors then $|w(v_0) - w(v_1)| \leq 1$. We shall associate to each tiling $t \in \mathcal{T}_a$ a *height function* $h_t : V \to \mathbb{R}$ defined by

$$h_t = \frac{1}{|\mathcal{T}_a|} \sum_{\tilde{t} \in \mathcal{T}_a} \operatorname{wind}(t - \tilde{t}).$$

Thus, h_t is the average of the windings $\operatorname{wind}(t - \tilde{t})$ for $\tilde{t} \in \mathcal{T}_a$. Notice that for any $t \in \mathcal{T}_a$ the height function h_t satisfies:

(a) $h_t(\infty) = 0$;

(b) for any $v \in V$, $h_t(v) \equiv h_{t_0}(v) \pmod{\mathbb{Z}}$;

(c) if $v_0, v_1 \in V$ are neighbors then $|h_t(v_0) - h_t(v_1)| < 1$.

Condition (a) follows from the equivalent equation for $\operatorname{wind}(t - \tilde{t})$; (b) follows from $\operatorname{wind}(t - \tilde{t}) = h_t - h_{\tilde{t}}$. Finally, the strict inequality in (c) follows from the hypothesis of t_0 being stable.

Conversely, any function $h : V \to \mathbb{R}$ satisfying conditions (a), (b) and (c) above is the height function of a (unique) tiling $t \in \mathcal{T}_a$. Indeed, take $w = h - h_{t_0}$ (which is integer valued) and $t = \partial h$: conditions (b) and (c) guarantee that for any vertex in $\mathcal{G}^*$ exactly one edge adjacent to it receives the coefficient 1. In particular, the maximum or minimum of two height functions is a height function. For $t, \tilde{t} \in \mathcal{T}_a$, write $t \leq \tilde{t}$ if $h_t(v) \leq h_{\tilde{t}}(v)$ for all $v \in V$. Also, $t < \tilde{t}$ if $t \leq \tilde{t}$ and $t \neq \tilde{t}$.

Two height functions $t, \tilde{t} \in \mathcal{T}_a$ differ by a flip at $v \in V \smallsetminus \{\infty\}$ if $h_t(v) - h_{\tilde{t}}(v) = \pm 1$ and $h_t(\tilde{v}) = h_{\tilde{t}}(\tilde{v})$ for $\tilde{v} \in V \smallsetminus \{v\}$. Conversely, a flip at $v \in V \smallsetminus \{\infty\}$ is allowed from $t \in \mathcal{T}_a$ if and only if v is a local maximum or minimum of h_t.

Theorem 3. *Consider a pair $(S, \mathcal{G}^*)$ and two stable tilings t_0, t_1 of S with $\operatorname{Flux}_S(t_0) = \operatorname{Flux}_S(t_1)$. Then t_0 and t_1 are connected by flips.*

Proof. Assume $t_0 < t_1$. We show that we can perform a flip on t_1 in order to obtain $\tilde{t}$ with $t_0 \leq \tilde{t} < t_1$; by induction, this completes the proof. Let $V_1 \subset V$ be the set of $v \in V$ for which $h_{t_1}(v) - h_{t_0}(v)$ is maximal. Let v_2 be the point of V_1 where h_{t_1} is maximal. We claim that v_2 is a local maximum of h_{t_1}. Indeed, let $\tilde{v}$ be a neighbor of v_2. If $\tilde{v} \in V_1$ then $h_{t_1}(\tilde{v}) < h_{t_1}(v_2)$ by definition of v_2. If $\tilde{v} \notin V_1$ then $h_{t_0}(v_2) < h_{t_0}(\tilde{v}) \leq h_{t_1}(\tilde{v}) < h_{t_1}(v_2)$ by definition of V_1, completing the proof of the claim. Let

$$h_{\tilde{t}}(v) = \begin{cases} h_{t_1}(v), & v \neq v_2; \\ h_{t_1}(v_2) - 1; & v = v_2 : \end{cases}$$

the function $h_{\tilde{t}}$ satisfies conditions (a), (b) and (c) and therefore defines a valid tiling $\tilde{t} \in \mathcal{T}_a$ with the required properties. □

Remark 8.1. The results of this section, including Theorem 3, hold more generally for any *coquadriculated surface*, a pair $(S, \mathcal{G}^*)$ consisting of: an oriented compact connected topological surface S with non-empty boundary ∂S and an embedded connected bipartite graph $\mathcal{G}^* \subset S$ such that every connected component of $S \smallsetminus \mathcal{G}^*$ is a square, that is, is surrounded by a cycle of length 4 in $\mathcal{G}^*$. Namely, it is not necessary to have an ambient manifold to induce the surface being tiled.

9 Connectivity

In this section we prove our main theorem. Let R be a cubiculated region. Let t_0, t_1 be tilings of R with $\mathrm{Flux}(t_0) = \mathrm{Flux}(t_1)$. A discrete Seifert surface (S, ψ) for the pair (t_0, t_1) is:

- *balanced* if the number of black vertices equals the number of white vertices in the interior of S;
- *zero-flux* if $\phi(t_0; S) = \phi(t_1; S) = 0$;
- *tangent* if t_0 and t_1 are tangent to S (including in the interior of the surface).

Notice that a tangent surface is clearly both balanced and zero-flux. The converse implications in general do not hold.

Recall from Lemma 5.3 that if $\mathrm{Flux}(t_0) = \mathrm{Flux}(t_1)$ then there exists a refinement $R^{(k)}$ and a discrete Seifert surface for the pair (t_0, t_1) in this refinement. We are now ready to prove our main theorem by considering when nicer Seifert surfaces can be constructed.

Proof of Theorem 1. We in fact prove a series of results. Each item below could be considered a distinct lemma, but since they build naturally to the main result, we prefer to present this as a single proof.

Consider a cubiculated region R and two distinct tilings t_0 and t_1 of R with $\mathrm{Flux}(t_0) = \mathrm{Flux}(t_1)$.

1. If $\mathrm{Tw}(t_0) = \mathrm{Tw}(t_1)$ then there exists a refinement $R^{(k)}$ and a discrete *zero-flux* Seifert surface for the pair (t_0, t_1) in this refinement.

2. If there exists a discrete zero-flux Seifert surface for the pair (t_0, t_1) then there exists a refinement $R^{(k)}$ and a discrete *balanced zero-flux* Seifert surface for the pair (t_0, t_1) in this refinement.

3. If there exists a discrete balanced zero-flux Seifert surface for the pair (t_0, t_1) then there exists a refinement $R^{(k)}$ and tilings $\tilde{t}_i$ of $R^{(k)}$ with $\tilde{t}_1 - \tilde{t}_0 = (t_1 - t_0)^{(k)}$, $\tilde{t}_i$ obtained from $t_i^{(k)}$ by a sequence of flips, and a *tangent* Seifert surface for the pair $(\tilde{t}_0, \tilde{t}_1)$.

4. If there exists a discrete tangent Seifert surface for the pair (t_0, t_1) then there exists a refinement $R^{(k)}$ for which there exists a sequence of flips taking $t_0^{(k)}$ to $t_1^{(k)}$.

For item (1), start by constructing a smooth Seifert surface S. By the definition of twist, $\phi(t_0; S) = \phi(t_0; S)$ is a multiple of m and therefore an integer linear combination of $\phi(t_0; S_i) = \phi(t_1; S_i)$ where the S_i are smooth closed surfaces:

$$\phi(t_0; S) + \sum c_i \phi(t_0; S_i) = 0.$$

Let $\tilde{S}$ be the union of S with c_i copies of S_i (reverse orientation if $c_i < 0$). Perturb the surfaces to guarantee transversality. Along intersection curves, perform a standard cut-re-glue-smoothen procedure (see e.g.[17]) to obtain a smooth embedded Seifert Surface $\hat{S}$ with $\phi(t_0; \hat{S}) = \phi(t_0; \hat{S})$. As in Lemma 5.3, after refinements, $\hat{S}$ can be approximated by a discrete surface.

For item (2), notice first that the parity of the number of black and white interior vertices is already the same, otherwise $\phi(t_0; S)$ would be odd. In order to increase the black-minus-white difference by two, look for a white vertex in a planar part of S and lift a surrounding square as in Figure 9. Repeat the procedure as needed.

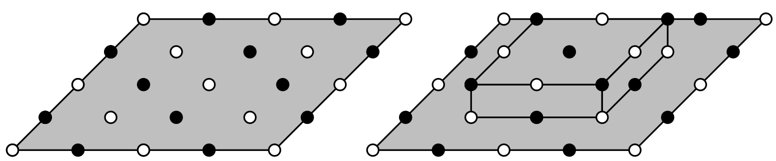

Figure 9: The four center squares on the bottom layer are removed and the surfaces is lifted. This increases the difference between the number of black versus white vertices; there is a net increase of 5 black vertices and 3 white vertices.

For item (3), consider vertices in the interior of S which are matched with points outside S. Classify these vertices as black-above, white-above, black-below and white-below according to their color and the position of their match. It follows from S being balanced and zero-flux that the number of black-above and white-above are equal (and similarly for black-below and white-below). Associate vertices black-above and white-above in pairs (and similarly for black-below and white-below). We show how to perform flips in order to have such pairs cancel out. Starting from a black-above vertex, draw a simple curve γ on S using tangent dominoes until you arrive at its white-above partner. Since the surface has been refined there are large regions of parallel dimers (or dominoes) both on the surface and near S. We may therefore construct a narrow disk S_1 by going "above" γ in the direction normal to S. As in Figure 4, by performing flips and taking refinements, we may assume S_1 to be a surface with an induced tiling by t. By construction, both the black-above and the white-above dominoes belong to $\gamma_1 = \partial S_1$. Again by virtue of refinements we may assume t to be stable (as a tiling of S_1). Apply Theorem 3 (for S_1, not the original S) in order to obtain a sequence of flips whose effect is to rotate γ_1, thus getting rid of both the black-above and the white-above dominoes.

For item (4), we apply Theorem 3 to the surface S. If t_0 (and t_1) are stable, this can be done directly. Otherwise, take refinements: appropriate refinements of t_0 (and t_1) are stable. This completes the proof of item (b) of Theorem 1.

As to item (a), let l be a positive integer such that $\mathrm{Tw}(t_1) = \mathrm{Tw}(t_0) \pm l$. Apply refinements such that $t_0^{(k)}$ and $t_1^{(k)}$ contain at least l boxes of dimension $3 \times 3 \times 2$ tiled by 9 parallel dominoes. Starting from $t_0^{(k)}$, apply flips and trits inside l such boxes so as to increase or decrease the twist by l, thus connecting by flips and trits $t_0^{(k)}$ to a tiling t with $\mathrm{Tw}(t) = \mathrm{Tw}(t_1)$. By item (2), t and t_1 can be connected by flips (possibly after further refinement). □

10 Final Remarks

Theorem 1 is the first positive result concerning connectivity of three-dimensional domino tilings. Given the nature of the result, many natural questions arise. We discuss some of these here, many more constructions, experiments, open problems and conjectures will appear in a forthcoming paper [9].

Our methods rely heavily on refinements but it is important to point out that they are not simply an artifact of our proof techniques. For example, already in the $4 \times 4 \times 4$ box, there are tilings with the same twist which are not flip connected before refinement.

On the other hand, in the general case of boxes, it is not known if refinements are needed: If t_0 and t_1 are two tilings of a fixed box, can t_0 and t_1 be connected by flips and trits? It is known that refinements are neccessary for connectivity of tilings of simple prisms and tori [20]. What is not clear, however, is how many refinements are needed. Our results offer no bounds on the number of times one needs to refine two tilings before they become connected.

It is also unknown how often refinement is neccessary for certain regions. For example, consider the cubical torus given by $\mathbb{Z}^3/(N \cdot \mathbb{Z}^3)$ for some even integer N. Based on computational experiments, we conjecture that the probability that two tilings with the same flux and twist are connected by flips tends to 1 as N tends to infinity, see [9].

More broadly, one asks, what does a 'typical' tiling look like? The twist essentially partitions the space of tilings into flip connected components. For a fixed region, what is the distribution of the twist? For example, consider again the case of boxes. Let the base tiling (twist 0) be one in which all tiles are parallel to a fixed axis. Related results suggest the following question: is the twist normally distributed about 0?

References

[1] Vladimir I Arnold and Boris A Khesin. *Topological methods in hydrodynamics*, volume 125. Springer Science & Business Media, 1999.

[2] Olivier Bodini and Damien Jamet. Tiling a Pyramidal Polycube with Dominoes. *Discrete Mathematics & Theoretical Computer Science*, 9(2), 2007.

[3] Mihai Ciucu. An improved upper bound for the three dimensional dimer problem. *Duke Math. J*, 94:1–11, 1998.

[4] Henry Cohn, Noam Elkies, and James Propp. Local statistics for random domino tilings of the Aztec diamond. *Duke Mathematical Journal*, 85(1):117–166, 1996.

[5] John H Conway and Jeffrey C Lagarias. Tiling with polyominoes and combinatorial group theory. *Journal of combinatorial theory, Series A*, 53(2):183–208, 1990.

[6] Guy David and Carlos Tomei. The problem of the calissons. *The American Mathematical Monthly*, 96(5):429–431, 1989.

[7] Noam Elkies, Greg Kuperberg, Michael Larsen, and James Propp. Alternating-sign matrices and domino tilings (Part II). *Journal of Algebraic Combinatorics*, 1(3):219–234, 1992.

[8] Alberto Enciso, Daniel Peralta-Salas, and Francisco Torres de Lizaur. Helicity is the only integral invariant of volume-preserving transformations. *Proceedings of the National Academy of Sciences*, 113(8):2035–2040, 2016.

[9] Juliana Freire, Caroline Klivans, Pedro Milet, and Nicolau Saldanha. Constructions in three-dimensional tilings. *In preparation.*

[10] Shmuel Friedland and Uri N Peled. Theory of computation of multidimensional entropy with an application to the monomer–dimer problem. *Advances in applied mathematics*, 34(3):486–522, 2005.

[11] J. M. Hammersley. Existence theorems and Monte Carlo methods for the monomer-dimer problem. In F.N. David, editor, *Research Papers in Statistics: Festschrift for J. Neyman*, pages 125–146. Wiley, 1966.

[12] William Jockusch, James Propp, and Peter Shor. Random domino tilings and the arctic circle theorem. *arXiv:math/9801068*, 1995.

[13] P.W. Kasteleyn. The statistics of dimers on a lattice: I. The number of dimer arrangements on a quadratic lattice. *Physica*, 27(12):1209 – 1225, 1961.

[14] Richard Kenyon and Andrei Okounkov. Planar dimers and Harnack curves. *Duke Mathematical Journal*, 131(3):499–524, 2006.

[15] Richard Kenyon, Andrei Okounkov, and Scott Sheffield. Dimers and amoebae. *Annals of mathematics*, pages 1019–1056, 2006.

[16] Boris Khesin and VI Arnold. Topological fluid dynamics. *Notices AMS*, 52(1):9–19, 2005.

[17] WB Raymond Lickorish. *An introduction to knot theory*, volume 175. Springer Science & Business Media, 2012.

[18] Joakim Linde, Cristopher Moore, and Mats G Nordahl. An n-dimensional generalization of the rhombus tiling. *Discrete Mathematics and Theoretical Computer Science*, 23:42, 2001.

[19] Pedro H Milet and Nicolau C Saldanha. Domino tilings of three-dimensional regions: flips, trits and twists. *arxiv preprint arxiv:1410.7693.*

[20] Pedro H Milet and Nicolau C Saldanha. Flip invariance for domino tilings of three-dimensional regions with two floors. *arXiv preprint arXiv:1404.6509*, 2014.

[21] Henry Keith Moffatt. The degree of knottedness of tangled vortex lines. *Journal of Fluid Mechanics*, 35(01):117–129, 1969.

[22] Igor Pak and Jed Yang. The complexity of generalized domino tilings. *The Electronic Journal of Combinatorics*, 20(4):P12, 2013.

[23] Dana Randall and Gary Yngve. Random three-dimensional tilings of Aztec octahedra and tetrahedra: an extension of domino tilings. In *Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*, pages 636–645. Society for Industrial and Applied Mathematics, 2000.

[24] Nicolau C Saldanha, Carlos Tomei, Mario A Casarin Jr, and Domingos Romualdo. Spaces of domino tilings. *Discrete & Computational Geometry*, 14(1):207–233, 1995.

[25] William P. Thurston. Conway's Tiling Groups. *The American Mathematical Monthly*, 97(8):pp. 757–773, 1990.

Departamento de Matemática, PUC-Rio
Rua Marquês de São Vicente, 225, Rio de Janeiro, RJ 22451-900, Brazil
`jufreire@gmail.com`
`milet@mat.puc-rio.br`

saldanha@puc-rio.br

Division of Applied Mathematics and Department of Computer Science
Brown University, Providence, RI, USA
caroline_klivans@brown.edu