Tilings of rectangles with T-tetrominoes

Michael Korn and Igor Pak

Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA, 02139

mikekorn@mit.edu, pak@math.mit.edu

August 26, 2003

Abstract

We prove that any two tilings of a rectangular region by T-tetrominoes are connected by moves involving only 2 and 4 tiles. We also show that the number of such tilings is an evaluation of the Tutte polynomial. The results are extended to a more general class of regions.

1 Introduction

The subject of tilings is a wonderful story that started as a collection of amateur problems (cf. [6]) and has now become an area of study in its own right, with numerous connections and applications to other fields: from group theory to topology, from enumerative combinatorics to probability. In the last decade various advanced methods have been developed which allowed some hard questions to be answered. This resulted in a structural approach to the study of tilings, which was presented in a recent survey [13] by the second author. The current paper carries out this approach to the very end for a special set of tiles. An unexpected combinatorial connection to the Tutte polynomial is a bonus and a delightful surprise.

In this paper we consider the set of T-tetrominoes, which has been studied earlier by Walkup in his curious paper [21]. His main result states that only rectangles of the form $4m \times 4n$ are tileable by T-tetrominoes. The proof in [21] is interesting but rather ad hoc. We explore further the structure of these tilings, combining Walkup's approach with several new direct bijections.

Our main result is local move connectivity of T-tetromino tilings for rectangular regions, resolving a conjecture in [13] in this case. This is done by introducing a new type of height function, and relating it to the tiling by means of two bijections. The height function technique for domino tilings was discovered by Thurston [19], and was used extensively in the recent literature to prove the local move connectivity for various sets of tiles (see e.g. [2, 8, 12, 17]).

Along the way we show that our new height functions enable us to define a lattice structure on all T-tetromino tilings of a rectangle. While this may seem a theoretical curiosity, in fact there are important applications of this result. There is a natural definition of a Markov chain on tilings (perform "random local moves") which translates into a nice Markov chain on height functions. Using the "coupling from the past" technique of Propp and Wilson and the fact that the height functions form a lattice (see [15, 14]) one can sample random T-tetromino tilings from an *exactly* uniform distribution (as opposed to a "nearly uniform" distribution usually obtained by the Markov chain approach). We refer to [16] for a full discussion of this approach and other examples of lattice structures on tilings.

Another classical problem for general tilings is their enumeration. Unfortunately, there seems to be little hope for a closed formula for the number of T-tetromino tilings of rectangles[1]. We show, however, that the

[1]For example, the 8×12 rectangle has an unpromising $1182 = 2 \cdot 3 \cdot 197$ tilings by T-tetrominoes.

number is an evaluation $T(3,3)$ of the *Tutte polynomial*, well-studied in the literature. Let us mention here that the Tutte polynomial is a fundamental invariant of graphs, and is related to a number of problems in discrete mathematics, computer science and several models in statistical mechanics. It has been extensively studied for various series of graphs as well as from a computational point of view (approximating its values is one of the challenges in the field). We refer to [22] for a beautifully written survey and a starting point for countless references.

Given the lack of a closed formula, one can ask whether the Markov chain approach can be used to efficiently *approximate* the number of tilings. In the past decade this idea has been developed into an important technique, first in the graph theoretic setting [7] and later in the tilings literature [9]. Without going into the details, the technique is based on showing rapid mixing of a Markov chain, and establishing a so-called self-reducibility property, the latter being usually much simpler. Roughly, one uses a Markov chain to sample a number of uniform tilings, collect statistics for certain patterns among the sampled tilings, and reduce the problem to a smaller similar problem. The self-reducibility allows such a reduction and keeps the errors relatively small. Unfortunately, we are unable to carry out this approach in full. We show the self-reducibility, but rapid mixing of our Markov chain goes beyond the scope of this work and is stated as a conjecture.

Returning back to combinatorics of the T-tetromino tilings, we answer a question as to what extent our results can be generalized from rectangular to other regions. We show that in fact all the bijective proofs go through for quadruplicated simply connected regions. As a corollary, we have the local move connectivity for such regions. We conclude by showing that if either condition on the regions is dropped, there is no local move property.

After the results of this paper were obtained, we learned of an alternative approach to the problem presented in a recent manuscript by Konstantin and Yuri Makarychev [10]. The authors showed that one can prove local move connectivity of T-tetromino tilings for rectangular regions by means of the so-called *ice graphs* and by using Eloranta's theorem (see [4]). We discovered that this approach can be combined with ours and one can define a height function, similar to ours. We present our findings in the appendix. For completeness and clarity of the exposition we start by recalling Makarychev's results (with independent proofs), define a new height function, and then proceed to prove the local move connectivity by a height function.

A few words about the structure of the paper. We start with definitions and basic results. In section 3 we state Walkup's result about the structure of T-tetromino tilings of a rectangle which will be an important technical tool for the rest of the paper. Then, in section 4, we define a new notion of chain graphs and show that they are in one-to-one correspondence with T-tetromino tilings. In section 5 we introduce the height functions, which we use in section 6 to prove local connectivity. We introduce a lattice structure on height functions in section 7. In section 8 we consider tilings of non-rectangular regions, proving local connectivity for quadruplicated simply connected regions. In section 9 we define two planar graphs which correspond to chain graphs, which enables us to obtain an enumerative formula for the number of T-tetromino tilings. Section 10 contains the description of the Markov chain $\mathcal{M}$ on T-tetromino tilings and proves the self-reducibility. We conclude with final remarks and the appendix outlining an alternative approach to local connectivity.

2 Main results

A *T-tetromino* is the figure formed by four unit squares arranged as shown in Figure 1. We make no distinction between the four possible orientations of the T-tetromino. A *tiling* of a region Γ with T-tetrominoes is an arrangement of T-tetrominoes which covers every square of Γ exactly once. An individual T-tetromino in such a tiling is called a *tile*. Let $\mathcal{T}_\Gamma$ denote the set of all possible tilings of Γ by T-tetrominoes.

In this paper, we will work exclusively with tilings by T-tetrominoes. Every reference to tilings or tiles refers to T-tetrominoes, even if this is not explicitly stated. We will chiefly be interested in tiling rectangular regions, although in section 8 we will see that much of what we prove also holds for a somewhat more general

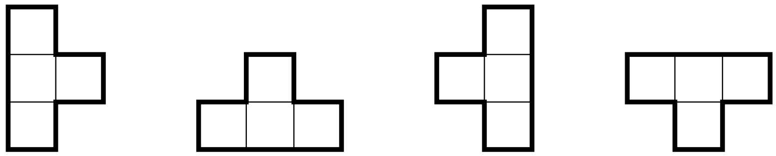

Figure 1: T-tetrominoes.

class of shapes.

Given a tiling of a region Γ, one can transform it into a new tiling by performing a "local move". A local move consists of picking up some (small) number of tiles, then re-filling that area with tiles in a different way. Two natural local moves for T-tetrominoes are shown in Figure 2. We call them the 2-move and the 4-move.

Figure 2: Local 2-move and local 4-move.

Suppose we are given a region Γ and a collection $\mathcal{L}$ of allowable local moves, and suppose that τ_1 and τ_2 are two different tilings of Γ. Let us say that τ_1 and τ_2 are *local-move equivalent* with respect to $\mathcal{L}$ if it is possible to transform τ_1 into τ_2 by performing a sequence of local moves from $\mathcal{L}$. This is an equivalence relation on the set of all tilings of Γ. A natural question to ask is whether all tilings of Γ are local-move equivalent. If so, we say that the region Γ has *local connectivity* with respect to $\mathcal{L}$. If $\mathcal{R}$ is a set of regions (the set of all rectangles, or the set of all simply-connected regions, for example), we say that $\mathcal{R}$ has a *local-move property* if there exists a finite set $\mathcal{L}$ of local moves such that all regions $\Gamma \in \mathcal{R}$ have local connectivity with respect to $\mathcal{L}$.

One of our main results is the following.

Theorem 1 *For tilings by T-tetrominoes, the set of all rectangles has a local-move property. Specifically, every rectangle Γ has local connectivity with respect to the 2-move and 4-move.*

This result was conjectured in [13] to hold for all simply connected regions. Later on, in section 8, we extend this theorem to a more general class of regions and show that the conjecture does not hold in full generality.

3 Tiling rectangles with T-tetrominoes

Without loss of generality, let Γ be a rectangle which is situated in the first quadrant of the Cartesian plane, with one corner at $(0, 0)$. Let a *type-A* point be a point whose coordinates are congruent mod 4 to (0,0) or (2,2), and let a *type-B* point be a point whose coordinates are congruent mod 4 to (0,2) or (2,0). A segment of length 1 is called a *cut* if there is no valid tiling of Γ in which a tile crosses that segment. A point is called *cornerless* if there is no valid tiling of Γ in which that point is one of the eight corners of a tile.

In [21], Walkup proves the following property of T-tetromino tilings of rectangles (see Figure 3).

Theorem 2 (Walkup) *If an $m \times n$ rectangle can be tiled by T-tetrominoes, then both m and n must be divisible by 4. Furthermore, all segments incident to type-A points are cuts, and all type-B points are cornerless.*

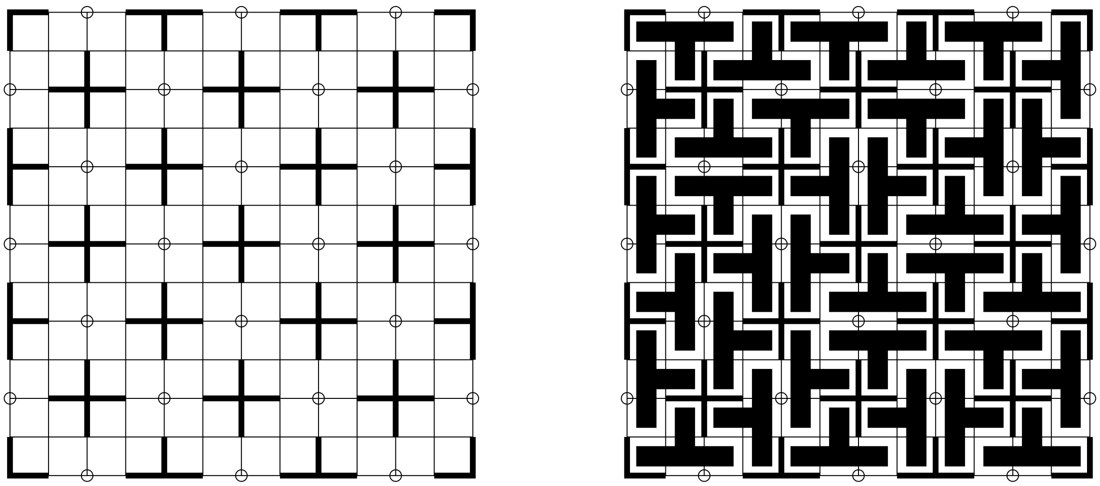

Figure 3: The dark lines are cuts. Circles are cornerless points.

From now on, we will only be concerned with rectangles having sides divisible by 4, since all other rectangles are untileable.

Define a *block* to be a 2×2 square whose corners have even coordinates. The following lemma is immediate by inspection from the structure of cuts and cornerless points.

Lemma 3 *In any tiling of a rectangle by T-tetrominoes, each tile contains three squares from one block and one square from an adjacent block. Similarly, each block contains three squares from one tile and one square from another tile.*

4 Chain graphs

Define an *antiblock* to be a 2×2 square whose corners have odd coordinates. Color the antiblocks white and gray in checkerboard fashion, so that antiblocks centered at type-A points are gray and those centered at type-B points are white.

For a $4m \times 4n$ rectangle Γ, let V_Γ be the set of points in Γ which have odd coordinates. Say that a directed graph on the vertices V_Γ is a *chain graph* if it satisfies the following properties:

- every edge connects vertices that are two units apart (either vertically or horizontally),
- every vertex has indegree 1 and outdegree 1, and
- every white antiblock contained in Γ borders exactly two edges of the graph, and these edges are non-adjacent.

Let $\mathcal{C}_\Gamma$ denote the set of all chain graphs of a region Γ.

Theorem 4 *For any $4m \times 4n$ rectangle Γ, we have $|\mathcal{C}_\Gamma| = |\mathcal{T}_\Gamma|$.*

The proof is based on an explicit bijection $\phi : \mathcal{T}_\Gamma \to \mathcal{C}_\Gamma$ defined as follows.

Let $\tau \in \mathcal{T}_\Gamma$ be a tiling. Notice that each vertex in V_Γ lies in the middle of some block. By Lemma 3, each tile in τ contains three squares from one block and one square from an adjacent block. Call these blocks the *primary* and *secondary* blocks of the tile respectively. For each tile, draw a directed edge from its primary block to its secondary block, and define $\phi(\tau)$ to be the directed graph which results (see Figure 4).

Theorem 4 follows immediately from the following Lemma.

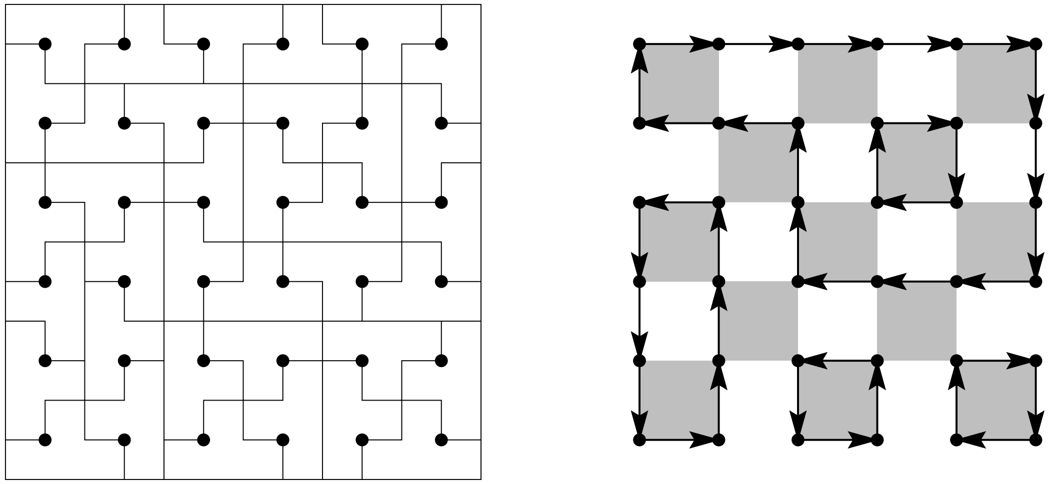

Figure 4: A tiling τ, and the chain graph $\phi(\tau)$.

Lemma 5 *For any $4m \times 4n$ rectangle Γ, the map ϕ defined above is a bijection between $\mathcal{T}_\Gamma$ and $\mathcal{C}_\Gamma$.*

Proof: First let us show that $\phi(\tau)$ is a chain graph. It is clear from the definition, and from Lemma 3, that every vertex will have indegree 1 and outdegree 1, and that edges will only connect vertices which are two units apart. As for the third restriction, consider a type-B point not on the boundary. Up to rotations and reflections, the tiles surrounding it must look like one of the two possibilities shown in Figure 5. Thus there will be exactly two edges bordering the associated white antiblock, and they will be non-adjacent. Hence $\phi(\tau)$ is a chain graph for all τ.

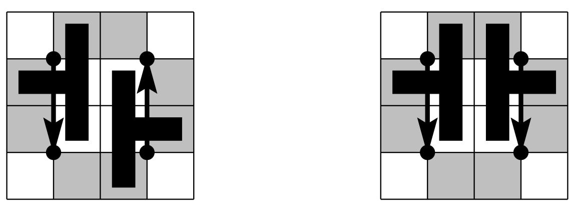

Figure 5: The two possibilities for a type-B point.

Notice that each tile corresponds to an edge in this graph. For each edge, there is only one possible tile placement which yields that edge and is consistent with the cuts and cornerless points. Hence the map ϕ is injective.

What remains to be shown is that every chain graph is equal to $\phi(\tau)$ for some tiling τ. As we just observed, for each edge there is only one possible tile placement that can yield that edge. So any chain graph will yield a collection of tile placements. It remains to be checked that these tiles cover all of Γ and do not overlap. Since each vertex has outdegree 1, the number of edges equals the number of blocks, so the total area of the tiles will equal the area of Γ. Thus it will be sufficient to verify that the tiles do not overlap.

Assume there are two tiles which overlap. Let us assume the overlap occurs in the block containing the squares A, B, D, and E (see Figure 6). Without loss of generality, we may take one of the tiles to be the one covering squares B, D, E, and F. Since each vertex has indegree 1 and outdegree 1, the tile which overlaps this one must contain only one square from this block, hence the overlap must occur at E. There are two possible tiles which cover E. First there is the tile which covers C, E, F, and G. If we have this, then the graph must contain both an edge and its opposite. This violates the rule about what a white antiblock may border. The other possibility is the tile which covers E, H, I, and J. In this case, the graph must contain two adjacent edges both on the same white antiblock, which again violates the constraint. Thus there can be no overlaps, which proves that every chain graph is $\phi(\tau)$ for some tiling τ. ■

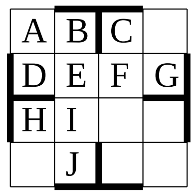

Figure 6: How an overlap may occur.

5 Height functions

Let us call a point having coordinates congruent mod 4 to (0,0) a *type-A0* point. Similarly, a point congruent to (2,2) will be called a *type-A1* point. (Points congruent to either (0,2) or (2,0) will still be called *type-B* points.)

For a $4m \times 4n$ rectangle Γ, let W_Γ be the set of points in Γ which have even coordinates. Let $\partial\Gamma$ denote the set of boundary point of Γ. Say that a function $f : W_\Gamma \to \mathbb{Z}$ is a *height function* if it satisfies the following properties:

- $f(x) = 0$ for all $x \in \partial\Gamma$,
- $f(x)$ is an even integer for all type-A0 points x,
- $f(x)$ is an odd integer for all type-A1 points x, and
- $|f(x) - f(y)| \leq 1$ whenever x and y are adjacent (at a distance of two units).

Let $\mathcal{H}_\Gamma$ denote the set of all height functions of a region Γ.

Theorem 6 *For any $4m \times 4n$ rectangle Γ, we have $|\mathcal{H}_\Gamma| = |\mathcal{T}_\Gamma|$.*

We define a map $\psi : \mathcal{C}_\Gamma \to \mathcal{H}_\Gamma$ as follows. Let $C \in \mathcal{C}_\Gamma$ be a chain graph. Define a function f° on the faces of C by the following rules. Let f° have the value 0 on the unbounded face of C. As we pass an edge of the graph, if the edge points to the right, let the value of f° increase by 1. (Similarly, if the edge points to the left, let the value of f° decrease by 1.) Now define $f : W_\Gamma \to \mathbb{Z}$ by letting $f(x)$ equal the value of f° on the face in which x lies (see Figure 7). Define $\psi(C)$ to be this function f.

Theorem 6 follows immediately from Theorem 4 and the following lemma.

Lemma 7 *For any $4m \times 4n$ rectangle Γ, the map ψ defined above is a bijection between $\mathcal{C}_\Gamma$ and $\mathcal{H}_\Gamma$.*

Proof: Let C be a chain graph, and let f be $\psi(C)$ as defined above. Let us first show that the function f is well-defined. If it is not, then there must exist some closed path through the faces of the graph such that the net change in the value of f° is non-zero. This means that upon going around this path counterclockwise, we cross more right-pointing edges than left-pointing edges, say. Therefore more edges leave the area enclosed by the path than enter that area. But this is impossible since every vertex has equal indegree and outdegree, so the net flow out of any region must be zero. Hence f is a well-defined function on W_Γ.

Next, let us verify that f is a valid height function. Points $x \in \partial\Gamma$ lie in the unbounded face of C, hence $f(x) = 0$ for such points. And if x and y are adjacent points, then they lie either in the same face of C or in adjacent faces of C, hence the difference between $f(x)$ and $f(y)$ is at most 1. Now let us verify the other two statements. As one travels from a type-A0 point x to another type-A0 point y which is 4 units away, one

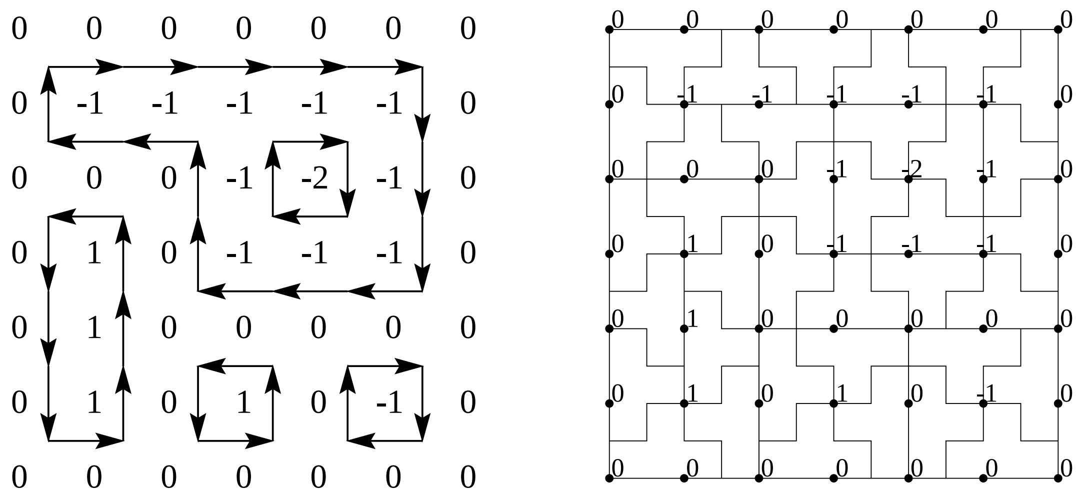

Figure 7: A chain graph C, and the function $f = \psi(C)$.

passes through the middle of a white antiblock (see Figure 8). In doing so, one crosses either 0 or 2 edges of C, hence the value of f° will have changed twice, or not at all, so $f(x)$ and $f(y)$ will have the same parity. Since (0,0) is a type-A0 point, and $f((0,0)) = 0$, it follows that $f(x)$ will be even for all type-A0 points x. By the same argument, all type-A1 points must have the same parity as each other. And (2,2) is a type-A1 point with $f((2,2)) = \pm 1$, so $f(x)$ will be odd for all type-A1 points x. Thus f is in fact a height function.

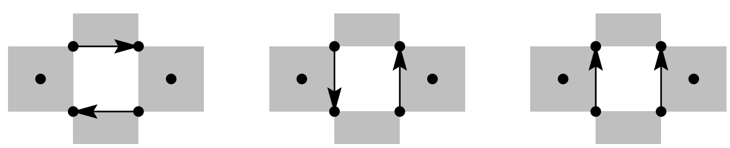

Figure 8: Two type-A0 points, and what might lie between them.

Given a height function $f = \psi(C)$, one can uniquely reconstruct the chain graph C by inserting directed edges in the places where the value of f increases or decreases. Hence ψ is an injective map. It remains to be shown that every height function f is equal to $\psi(C)$ for some valid chain graph C.

Take a height function f, and insert directed edges along the boundaries where the value of f increases or decreases. Call this graph C. Consider a vertex of C. To one corner of it, there is a type-A0 point x_0, on the opposite corner is a type-A1 point x_1, and the remaining two corners are type-B points y_0 and y_1. Since $f(x_0)$ is even, and $f(x_1)$ is odd, these values must differ by exactly 1. Without loss of generality, assume $f(x_0) = h$ and $f(x_1) = h + 1$. Then both $f(y_0)$ and $f(y_1)$ must be h or $h + 1$ as well. Up to rotations, the situation must look like one of the possibilities in Figure 9. Thus the vertex in question will have indegree 1 and outdegree 1.

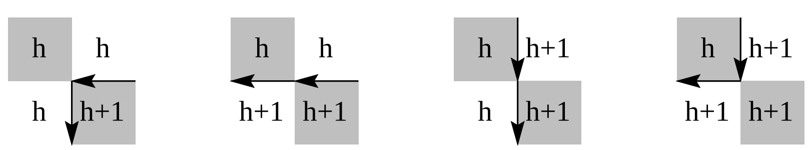

Figure 9: The possibilities for a vertex of C.

Now consider a type-B point y, which corresponds to a white antiblock. Let $f(y) = h$, and assume without loss of generality that h is even. If z_1 and z_2 are the two type-A0 points adjacent to y, then we must have $f(z_1) = f(z_2) = h$. If z_3 and z_4 are the two type-A1 points adjacent to y, then we must have $f(z_3) = h \pm 1$ and $f(z_4) = h \pm 1$, not necessarily the same (see Figure 10). So this white antiblock will border exactly two non-adjacent edges of C.

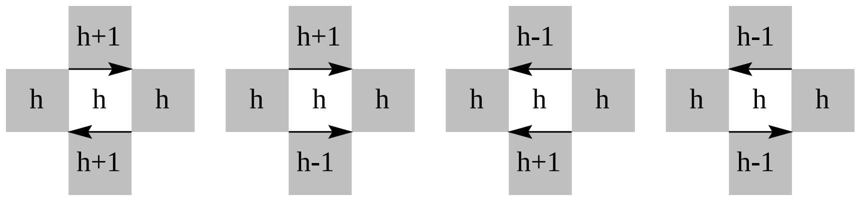

Figure 10: The possibilities for a white antiblock.

Hence the graph C constructed in this way from a height function f is indeed a chain graph, and $\psi(C) = f$. This completes the proof. ■

For ease of notation, define $\zeta(\tau) = \psi(\phi(\tau))$. For a $4m \times 4n$ rectangle Γ, the map ζ is the canonical bijection between $\mathcal{T}_\Gamma$ and $\mathcal{H}_\Gamma$.

Lemma 8 *Let Γ be a $4m \times 4n$ rectangle and let $\tau_1, \tau_2 \in \mathcal{T}_\Gamma$ be tilings of Γ. The tilings τ_1 and τ_2 differ by a 2-move if and only if the height functions $\zeta(\tau_1)$ and $\zeta(\tau_2)$ differ by 1 on some type-B point, and are the same everywhere else. The tilings τ_1 and τ_2 differ by a 4-move if and only if the height functions $\zeta(\tau_1)$ and $\zeta(\tau_2)$ differ by 2 on some type-A point, and are the same everywhere else.*

Proof: By inspection of the structure of cuts and cornerless points, one sees that the 2-move must be centered at a type-B point, and the 4-move must be centered at a type-A point. From Figure 11, one can see that if τ_1 and τ_2 differ by a 2-move, then the height functions $\zeta(\tau_1)$ and $\zeta(\tau_2)$ differ by 1 in their values on the corresponding type-B point. Similarly, if τ_1 and τ_2 differ by a 4-move, then the height functions $\zeta(\tau_1)$ and $\zeta(\tau_2)$ differ by 2 in their values on the corresponding type-A point.

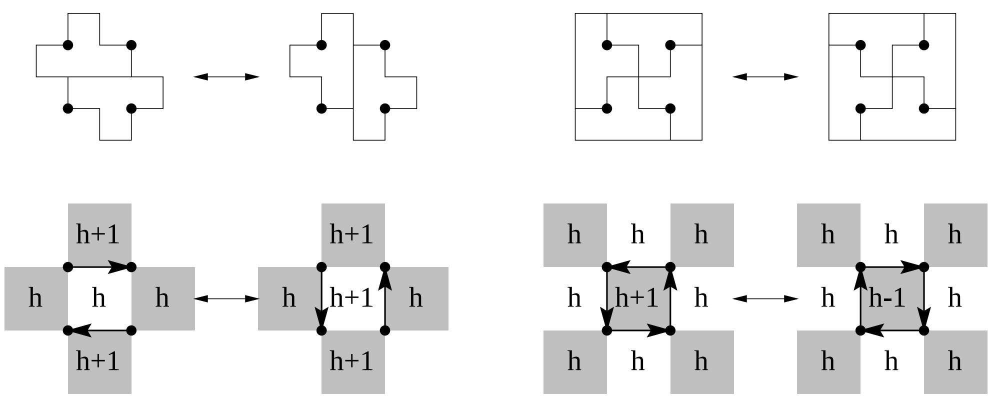

Figure 11: The 2-move and 4-move, and their effect on $\zeta(\tau)$.

As for the converse, suppose there are height functions f_1 and f_2 which are identical everywhere, except $f_1(y) = h$ and $f_2(y) = h + 1$ for some type-B point y. Thus the value of f_1 (or f_2) on the neighbors of y must be $h, h + 1, h$, and $h + 1$ (since they must alternate even and odd). Hence the picture must look like the bottom left of Figure 11, possibly rotated. Going backwards, we see what the chain graph and the tiling must then look like, and that in fact, $\zeta^{-1}(f_1)$ and $\zeta^{-1}(f_2)$ differ by a 2-move.

Similarly, suppose there are height functions f_1 and f_2 which are identical everywhere, except $f_1(x) = h+1$ and $f_2(x) = h - 1$ for some type-A point x. Thus $f_1(y) = f_2(y) = h$ for all neighbors y of x. Hence the picture must look like the bottom right of Figure 11. Going backwards, we see what the chain graph and the tiling must then look like, and that in fact, $\zeta^{-1}(f_1)$ and $\zeta^{-1}(f_2)$ differ by a 4-move. ■

For height functions $f_1, f_2 \in \mathcal{H}_\Gamma$, say that f_1 and f_2 differ by a 2-move (or 4-move) if the tilings $\zeta^{-1}(f_1)$ and $\zeta^{-1}(f_2)$ differ by a 2-move (or 4-move). By the previous Lemma, one can see that performing a 2-move on a height function f is equivalent to increasing or decreasing its value by 1 at some type-B point. Similarly, performing a 4-move is equivalent to increasing or decreasing the value of f by 2 at some type-A point. Of course, such moves may only be applied if the function that results is a valid height function.

6 Local connectivity from height functions

Theorem 1 will easily follow from the following lemma.

Lemma 9 *Let Γ be a $4m \times 4n$ rectangle, and let $f_1, f_2 \in \mathcal{H}_\Gamma$ be height functions. It is always possible to convert f_1 into f_2 by performing a sequence of 2-moves and 4-moves.*

Proof: For a $4m \times 4n$ rectangle Γ, let f_0 be the height function which is 1 on the type-A1 points of Γ, and 0 everywhere else. We would like to show that every height function f can be transformed into f_0. If every height function can be transformed into f_0, it follows that any height function can be transformed into any other. Suppose $f(x) > 1$ for some x. Let h be the largest value that f attains. Suppose there is a type-B point y which attains this value. Then f must take the values $h, h-1, h$, and $h-1$ on the neighbors of y. So we can perform a 2-move to change $f(y)$ to $h-1$ and still have a valid height function. We do this for all type-B points at which f attains the value h. Now look at any remaining (type A0 or A1) point x having $f(x) = h$. We must have $f(z) = h-1$ for the neighbors z of x, since there are no type-B points remaining for which $f(z) = h$. So we can perform a 4-move to change $f(x)$ to $h-2$. We do this for every point where f attains the value h. Now the largest value which appears is at most $h-1$, and we repeat the procedure until we have $f(x) \leq 1$ for all x.

We do a similar thing for points where $f(x) < 0$, increasing them until $f(x) \geq 0$ for all x. At this point, all points will have the value 0 or 1 (in particular, $f(x) = 0$ for all type-A0 points x, and $f(x) = 1$ for all type-A1 points). It just remains to set $f(y) = 0$ for all type-B points y, which can be done by a sequence of 2-moves. This finishes the procedure, proving the lemma. ■

7 The lattice structure on height functions

There is a natural partial order on $\mathcal{H}_\Gamma$. If $f_1, f_2 \in \mathcal{H}_\Gamma$ are height functions, we say $f_1 \leq f_2$ iff $f_1(x) \leq f_2(x)$ for all points x. This partial order can be extended to tilings—say $\tau_1 \leq \tau_2$ if $\zeta(\tau_1) \leq \zeta(\tau_2)$.

Theorem 10 *For any $4m \times 4n$ rectangle Γ, the poset P_Γ consisting of all tilings of Γ, with this order relation, is a distributive lattice.*

Proof: In order to prove that P_Γ is a lattice, we need to show that for height functions f_1 and f_2, there exists a unique greatest lower bound ("meet") α and least upper bound ("join") β. We define $\alpha(x) = \min\{f_1(x), f_2(x)\}$ and $\beta(x) = \max\{f_1(x), f_2(x)\}$, for all x. Clearly $\alpha \leq f_1$ and $\alpha \leq f_2$, and all other lower bounds are less than α. It just remains to be shown that α is a valid height function. Clearly the values of α on the boundary will be 0, and the type-A0 points will be even and the type-A1 points will be odd, because these properties hold for f_1 and f_2. As for adjacent values differing by at most 1, suppose x and y are adjacent points, and $\alpha(x) \geq \alpha(y) + 2$. Without loss of generality, assume $\alpha(y) = f_1(y)$. Then it would follow that $f_1(x) \geq \alpha(x) \geq \alpha(y) + 2 = f_1(y) + 2$, a contradiction. Therefore, α is a valid height function. The proof for β is analogous.

To prove that P_Γ is a distributive lattice, we need to verify the distributive laws: For height functions f, g, and h,

$$(f \vee g) \wedge (f \vee h) = f \vee (g \wedge h) \qquad \text{and} \qquad (f \wedge g) \vee (f \wedge h) = f \wedge (g \vee h).$$

For any x we have:

$$((f \vee g) \wedge (f \vee h))(x) = \min(\max(f(x), g(x)), \max(f(x), h(x)).$$

The functions min and max satisfy the distributive laws, so we have

$$\min(\max(f(x), g(x)), \max(f(x), h(x)) = \max(f(x), \min(g(x), h(x)) = (f \vee (g \wedge h))(x).$$

Hence

$$(f \vee g) \wedge (f \vee h) = f \vee (g \wedge h),$$

as desired. Note that changing the sign of functions switches the role of $\vee$ and $\wedge$, which implies the second distributive law. Therefore, P_Γ is a distributive lattice. ■

8 Non-rectangular regions

A quadruplicated simply connected region is a region which is formed by taking a simply-connected union of grid squares and dilating the figure by 4 in each direction. Let $\mathcal{Q}$ denote the set of all such regions. As we did for rectangles, we will assume that the corners of such a shape have coordinates which are congruent to (0,0) mod 4. Notice that $\mathcal{Q}$ contains all $4m \times 4n$ rectangles.

Theorem 11 *The second part of Theorem 2 holds for all regions* $\Gamma \in \mathcal{Q}$.

Proof: Suppose there exists a region $\Gamma \in \mathcal{Q}$ which can be tiled in a way which violates some of the supposed cuts and cornerless points. Let Γ' be the smallest $4m \times 4n$ rectangle which contains Γ. We can extend the tiling of Γ to a tiling of Γ' by adding tiled 4×4 squares to the part of Γ' which is not in Γ. This gives a tiling of Γ' which violates the necessary cuts and cornerless points, which contradicts Theorem 2. ■

As a result of this, all the above results for rectangles are also true for all $\Gamma \in \mathcal{Q}$. The proofs are the same as before.

The results do not hold if we drop the condition of being simply-connected. (Notice that the correspondence between chain graphs and height functions breaks down if the region is not simply connected, because points on the boundary of the region need not be on the unbounded face of the chain graph, so they may have nonzero height.) For example, Figure 12 shows a tiling of a non-simply connected region where neither the 2-move nor the 4-move can be applied.

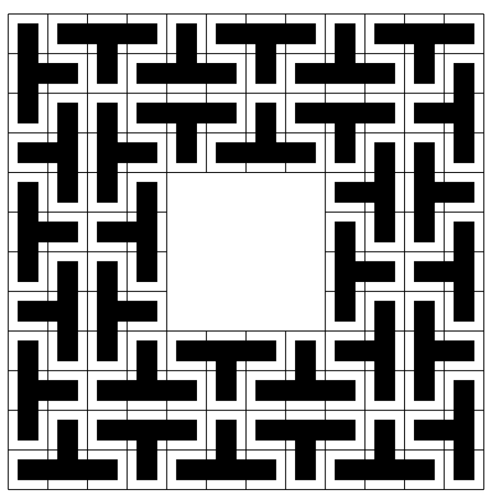

Figure 12: Tiling of a non-simply connected region.

Theorem 12 *Let $\mathcal{S}$ denote the set of all simply-connected regions. For tilings by T-tetrominoes, the set $\mathcal{S}$ does not have a local-move property.*

Proof: Let Δ_1 denote the region shown in Figure 13. It is straightforward to see that this region can be tiled in only two ways, namely the way shown and its mirror image. Since there are no intermediate tilings, and no tile is in the same place in both tilings, the only way for local connectivity to hold for this region is if we declare this entire transformation to be one local move.

In fact, we can generate infinitely many regions which admit only two tilings. Let Δ_k denote the region in Figure 14, where the total length of the region is $8k + 2$. As before, it can only be tiled in two ways, so in order to have local connectivity, the entire region must be considered to be a local move. No finite set of local moves can contain all of these, hence any finite set of local moves is insufficient to give local connectivity for these regions. ■

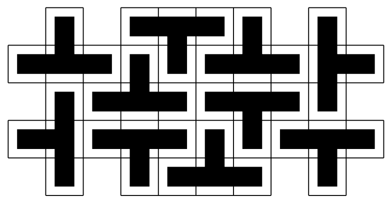

Figure 13: The region Δ_1.

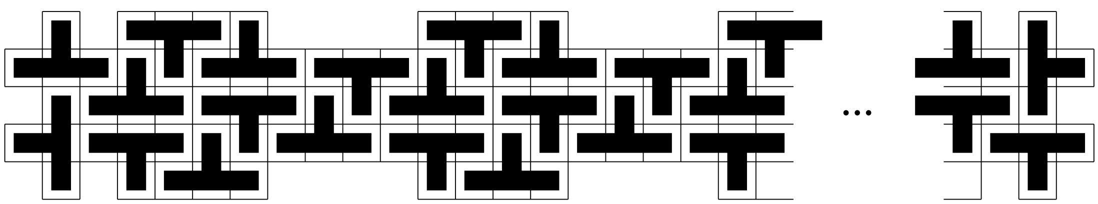

Figure 14: The region Δ_k.

9 Enumeration of tilings and the Tutte polynomial

For a region $\Gamma \in \mathcal{Q}$, define the graph G_Γ as follows. Include a vertex for each type-A1 point, and connect two vertices with an undirected edge if they are 4 units apart (vertically or horizontally). Similarly, define G^*_Γ by including a vertex for every type-A0 point, and again connecting those vertices which are 4 units apart. Note that when Γ is a $4m \times 4n$ rectangle, the graphs G_Γ and G^*_Γ are isomorphic to the $m \times n$ and $(m + 1) \times (n + 1)$ rectangular shape subgraphs of the square grid.

For a graph G, we let $V(G)$ and $E(G)$ denote the set of vertices and edges of G respectively. Let $c(G)$ denote the number of connected components of G. If $e \in E(G)$, let $G\backslash e$ be the graph formed by deleting e from G. Similarly, let G/e be the graph formed by contracting e in G.

The Tutte polynomial $T(G; x, y)$ is a polynomial in the variables x and y which is defined for undirected graphs G. Typically it is defined in terms of the following recursive formulas (see [22]):

- $T(G; x, y) = 1$ if G has no edges,
- $T(G; x, y) = y \cdot T(G\backslash e; x, y)$ if e is a loop,
- $T(G; x, y) = x \cdot T(G/e; x, y)$ if e is a cutedge,
- $T(G; x, y) = T(G\backslash e; x, y) + T(G/e; x, y)$ if e is neither a loop nor a cutedge.

Another equivalent definition of $T(G; x, y)$ is as follows. Let H be a spanning subgraph of G (that is, a subgraph of G which contains all the vertices of G). Then

$$T(G; x, y) = \sum_{H \subset G} (x-1)^{c(H)-c(G)}(y-1)^{c(H)+|E(H)|-|V(G)|}$$

where the sum is over all spanning subgraphs $H \subset G$.

Theorem 13 *For every $\Gamma \in \mathcal{Q}$, the number of T-tetromino tilings of Γ is equal to $2 \cdot T(G_\Gamma; 3, 3)$.*

To prove this, we introduce a few lemmas about spanning subgraphs of G_Γ and G^*_Γ.

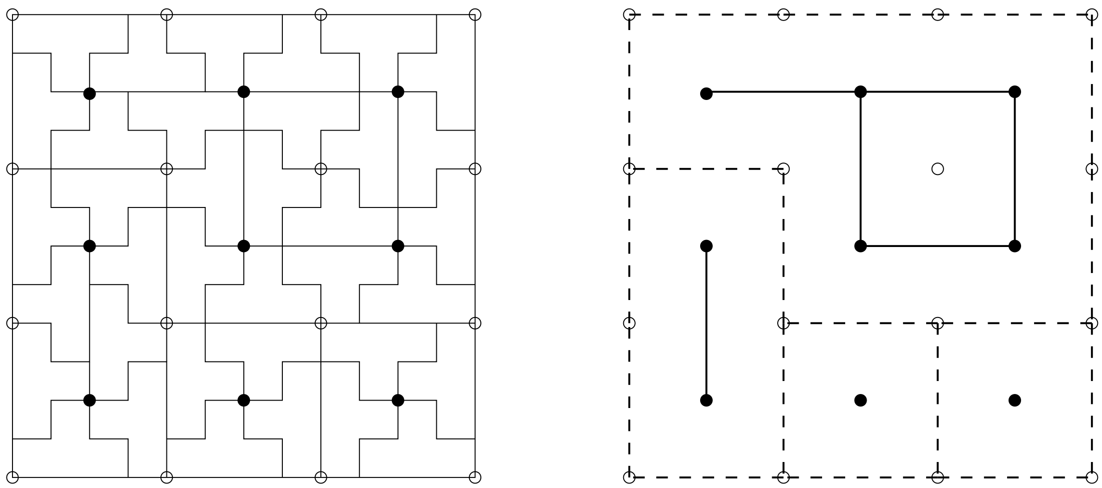

Figure 15: A tiling τ, and the graphs $\sigma(\tau)$ (solid lines) and $\sigma^*(\tau)$ (dotted lines).

Given a tiling τ of Γ, define $\sigma(\tau)$ to be the spanning subgraph of G_Γ which includes those edges which do not cross any tile. Similarly, define $\sigma^*(\tau)$ to be the spanning subgraph of G^*_Γ which includes those edges which do not cross any tile (see Figure 15).

Suppose H is a spanning subgraph of G_Γ. Define $\omega(H)$ to be the spanning subgraph of G^*_Γ consisting of those edges which do not cross any edge of H.

Lemma 14 *Fix $\Gamma \in \mathcal{Q}$ and a tiling $\tau \in \mathcal{T}_\Gamma$. Then $\omega(\sigma(\tau)) = \sigma^*(\tau)$. Furthermore, no edge of the chain graph $\phi(\tau)$ crosses an edge of either $\sigma(\tau)$ or $\sigma^*(\tau)$. Conversely, any edge of G_Γ or G^*_Γ which does not cross any edge of $\phi(\tau)$ is an edge of $\sigma(\tau)$ or $\sigma^*(\tau)$.*

Proof: Notice that the points where an edge of G_Γ and an edge of G^*_Γ intersect are precisely the type-B points in the interior of Γ. Consider any such point. Recalling Figure 5, observe that exactly one of the two edges which meet there will avoid crossing tiles of τ. Hence each such point is on an edge of either $\sigma(\tau)$ or $\sigma^*(\tau)$, but not both. So an edge of G^*_Γ is in $\sigma^*(\tau)$ if and only if no edge of $\sigma(\tau)$ crosses it. Hence $\sigma^*(\tau) = \omega(\sigma(\tau))$.

Recall that in $\phi(\tau)$, each edge corresponds to a tile; the edge connects the two blocks in which the tile lies. Edges of $\sigma(\tau)$ and $\sigma^*(\tau)$ run along block boundaries; an edge is present in these graphs if and only if no tile crosses that boundary. If no tile crosses that boundary, then no edge of $\phi(\tau)$ will either. Conversely, if no edge of $\phi(\tau)$ crosses a block boundary, then no tile crosses that boundary, hence that boundary will be an edge of $\sigma(\tau)$ or $\sigma^*(\tau)$. (See Figure 16.) ■

Corollary 15 *Suppose a region $\Gamma \in \mathcal{Q}$ and tilings $\tau_1, \tau_2 \in \mathcal{T}_\Gamma$ satisfy $\sigma(\tau_1) = \sigma(\tau_2)$. Then $\phi(\tau_1)$ and $\phi(\tau_2)$ are identical up to the orientation of the edges.*

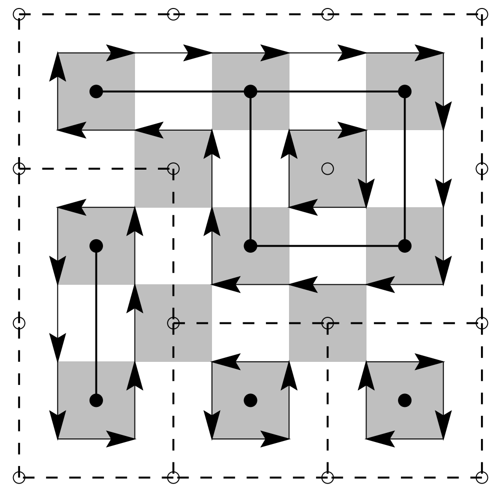

Figure 16: The graphs $\sigma(\tau)$ and $\sigma^*(\tau)$, and the chain graph $\phi(\tau)$.

Proof: Let $H = \sigma(\tau_1) = \sigma(\tau_2)$. For each white antiblock, there is exactly one edge of G_Γ which crosses it. The presence or absence of that edge in H determines which pair of edges along the white antiblock must be included in the corresponding chain graphs. This gives all the edges of the chain graphs, except those which do not border a complete white antiblock (ones near the boundary of the region). By inspection, one can see that all those edges must be included in order to have total degree 2 at each vertex of the chain graphs. ■

Lemma 16 *Let $\Gamma \in \mathcal{Q}$, and let H be a spanning subgraph of G_Γ. Then*

$$c(\omega(H)) \;=\; c(H) + |E(H)| - |V(G_\Gamma)| + 1.$$

Proof: We fix Γ and prove this by induction on the number of edges in H. If H has no edges, then $c(H) = |V(G_\Gamma)|$, so $c(H) + |E(H)| - |V(G_\Gamma)| + 1 = 1$, which is equal to $c(\omega(H))$, as required. Now assume that the result holds for all subgraphs $H \subset G_\Gamma$ with $|E(H)| < k$.

Consider a subgraph H with $|E(H)| = k$, and let $e \in E(H)$. First, suppose that e is a cutedge of H. Then $c(H\backslash e) = c(H) + 1$, $|E(H\backslash e)| = |E(H)| - 1$, and $c(\omega(H\backslash e)) = c(\omega(H))$. We conclude:

$$c(\omega(H)) = c(\omega(H\backslash e)) = c(H\backslash e) + |E(H\backslash e)| - |V(G_\Gamma)| + 1 = c(H) + |E(H)| - |V(G_\Gamma)| + 1.$$

Now suppose that e is not a cutedge of H. Then $c(H\backslash e) = c(H)$, $|E(H\backslash e)| = |E(H)| - 1$, and $c(\omega(H\backslash e)) = c(\omega(H)) - 1$. We have

$$c(\omega(H)) = c(\omega(H\backslash e)) + 1 = c(H\backslash e) + |E(H\backslash e)| - |V(G_\Gamma)| + 2 = c(H) + |E(H)| - |V(G_\Gamma)| + 1,$$

as desired. Therefore $c(\omega(H)) = c(H) + |E(H)| - |V(G_\Gamma)| + 1$ holds for all subgraphs $H \subset G_\Gamma$. ■

Suppose H is a spanning subgraph of G_Γ. Define $a(H) = 2c(H) + |E(H)| - |V(G_\Gamma)|$. Theorem 13 now follows from the following lemma.

Lemma 17 *Let Γ be a region in $\mathcal{Q}$. For every spanning subgraph $H \subset G_\Gamma$, there are exactly $2^{a(H)}$ tilings τ for which $\sigma(\tau) = H$.*

Proof: We need to show that for every spanning subgraph $H \subset G_\Gamma$, the corresponding (undirected) chain graph consists of $a(H)$ cycles. Each cycle can be oriented in two ways, hence we will get $2^{a(H)}$ valid chain graphs which correspond to H. Since chain graphs are in one-to-one correspondence with tilings, the result will follow.

Let C be a chain graph which corresponds to H. If C consists of k cycles, then it divides the plane into $k+1$ zones (possibly having holes). Each such zone is a maximal connected region on which the height function f is constant. Each zone must contain at least one type-A point, and thus must contain at least one vertex of H or $\omega(H)$. It cannot contain points from both H and $\omega(H)$, since the value of f is odd on the vertices of H and it is even on the vertices of $\omega(H)$. Observe that all vertices of H or $\omega(H)$ which live in the same zone are connected. Hence H and $\omega(H)$ have a total of $k+1$ connected components. Then $k = c(H) + c(\omega(H)) - 1 = 2c(H) + |E(H)| - |V(G_\Gamma)| = a(H)$, so the number of cycles in C is equal to $a(H)$, which proves the lemma. ∎

10 Sampling of tilings

Let $\Gamma \in \mathcal{Q}$ be a quadruplicated simply-connected region. Define a Markov chain $\mathcal{M}$ whose states are T-tetromino tilings of Γ. Allow a transition from τ_1 to τ_2 if τ_1 and τ_2 differ by a 2-move or 4-move, with the probability of such a transition being $1/N$, where $N = |\Gamma|$ is the area of Γ. Observe that $N/2$ is larger than the maximum number of different local moves which can be applied to any one tiling. Now, let the probability of staying put in the state τ_1 be $1 - k/N \geq 1/2$, where k is the number of different local moves which can be applied to τ_1.

Observe that $\mathcal{M}$ is symmetric, and aperiodic since the probability of staying put is always $\geq 1/2$. Therefore, by Theorem 1, the Markov chain $\mathcal{M}$ is ergodic and converges to the uniform distribution on $\mathcal{T}_\Gamma$. The mixing time of $\mathcal{M}$ remains open, but we would like to make the following conjecture:

Conjecture 18 *The mixing time of the Markov chain $\mathcal{M}$ is polynomial in the area of Γ.*

We refer the reader to [1] for the various definitions of the mixing time of Markov chains and related results. Now, if the conjecture is true, we can use the Markov chain $\mathcal{M}$ to sample tilings $\tau \in \mathcal{T}_\Gamma$ from a nearly uniform distribution. Using the notion of self-reducibility (see introduction, [18]), we can use sampling to approximate $|\mathcal{T}_\Gamma|$. The self-reducibility of tilings follows from the following lemma.

Lemma 19 *Let $\Gamma \in \mathcal{Q}$, and consider a tiling $\tau \in \mathcal{T}_\Gamma$ chosen uniformly at random. Let S be the leftmost 4-by-4 square in the top row of Γ. Unless S is all of Γ, the probability that S is isolated in τ (covered by exactly 4 tiles) is at least 1/3 and at most 2/3.*

Proof: The 4-by-4 square S corresponds to a vertex s in G_Γ. Notice that because there is nothing to the left of S or above it, the vertex s must have degree 1 or 2 in G_Γ. The square S will be isolated if and only if no edge of $\sigma(\tau)$ is incident to s.

Case 1: Suppose s has degree 1 in G_Γ. Let e be the edge of G_Γ incident to s. Let H be a spanning subgraph of $G_\Gamma - \{s\}$. Let H_0 be the spanning subgraph of G_Γ which consists of just those edges in H, and let H_1 be the spanning subgraph of G_Γ which consists of those edges in H, plus e. Consider all tilings τ such that $\sigma(\tau)$ is either H_0 or H_1. We want to know what proportion of these tilings have $\sigma(\tau) = H_0$. Notice that $|E(H_0)| = |E(H_1)| - 1$, and $c(H_0) = c(H_1) + 1$. It follows that $a(H_0) = a(H_1) + 1$. So by Lemma 17, there will be twice as many tilings with $\sigma(\tau) = H_0$ as there are with $\sigma(\tau) = H_1$. This is true for any $H \subset G_\Gamma - \{s\}$. So upon picking a random tiling τ, the probability that e is present in $\sigma(\tau)$ is 1/3. So in this case, S is isolated with probability 2/3.

Case 2: Suppose s has degree 2 in G_Γ. Let e_1 and e_2 be the edges of G_Γ incident to s, and let t_1 and t_2 be the vertices adjacent to s along edges e_1 and e_2 respectively. Let H be a spanning subgraph of $G_\Gamma - \{s\}$. Let H_0 be the spanning subgraph of G_Γ which consists of just those edges in H, let H_1 be the graph which includes the edges of H plus e_1, let H_2 include the edges of H plus e_2, and let H_3 include the edges of H plus e_1 and e_2. Consider two subcases.

Subcase 2a: Suppose t_1 and t_2 are in different components of H. Notice that $|E(H_0)| = |E(H_1)| - 1 = |E(H_2)| - 1 = |E(H_3)| - 2$, and $c(H_0) = c(H_1) + 1 = c(H_2) + 1 = c(H_3) + 2$. So $a(H_0) = a(H_1) + 1 =$

$a(H_2) + 1 = a(H_3) + 2$. So among all tilings τ which come from one of these graphs, 4/9 of them will have $\sigma(\tau) = H_0$, 2/9 of them will have $\sigma(\tau) = H_1$, 2/9 of them will have $\sigma(\tau) = H_2$, and 1/9 of them will have $\sigma(\tau) = H_3$.

Subcase 2b: Suppose t_1 and t_2 are in the same component of H. In this case, $|E(H_0)| = |E(H_1)| - 1 = |E(H_2)| - 1 = |E(H_3)| - 2$, and $c(H_0) = c(H_1) + 1 = c(H_2) + 1 = c(H_3) + 1$. So $a(H_0) = a(H_1) + 1 = a(H_2) + 1 = a(H_3)$. So among all tilings τ which come from one of these graphs, 1/3 of them will have $\sigma(\tau) = H_0$, 1/6 of them will have $\sigma(\tau) = H_1$, 1/6 of them will have $\sigma(\tau) = H_2$, and 1/3 of them will have $\sigma(\tau) = H_3$.

Combining subcases 2a and 2b, we get the following. For any H, either 1/3 or 4/9 of the tilings which correspond to H will have S isolated. Hence when we sum over all possible graphs H, we find that between 1/3 and 4/9 of all tilings of Γ have S isolated, when s has degree 2 in G_Γ.

This proves the lemma. ■

11 Final remarks

We should mention that our chain graphs seem to be well known in the Statistical Physics literature under a name "fully-packed loop model on the square lattice"; in this case all loops have fugacity 2. We refer to [**?**] for an appearance of this model in Combinatorics literature, exact terminology and further references.

A number of questions remain for future study. First and foremost, it would be interesting to show that the mixing time of $\mathcal{M}$ is polynomial, resolving Conjecture 18. If the proof goes along similar lines as that in [9], it should lead to new interesting combinatorial notions of the "intermediate" height functions between the smallest and the largest (of a fixed region).

A related question would be to show hardness of approximation of the number of T-tetromino tilings of regions $\Gamma \in \mathcal{Q}$. For general regions and for planar bipartite regions #P results have been obtained for various evaluations of the Tutte polynomial (see [20, 22]), but for regions on a square grid much work is yet to be done (cf. [5]).

In a different direction, are there other tiling type problems which lead to, perhaps other, evaluations of the Tutte polynomial? Can the present construction be extended to coverings of certain perhaps complicated graphs with copies of $K_{1,3}$, so that their number is equal to $T(G; 3, 3)$ for general graphs G?

A more philosophical (and thus more difficult) question is to explain the meaning behind T-tetrominoes. What geometric properties of the T-tetrominoes force the rigid structure discovered by Walkup? Do all such structures imply the existence of height functions?

In general, are there other simple collections of tiles so that a certain rich collection of regions has a rigid structure of tilings, while general simply connected regions do not? Philosophically, this amounts to understanding the extent the boundary conditions control the structure of tilings in the middle. The example of domino tilings of Aztec diamonds comes to mind [3].

Finally, what can be said about the asymptotic behavior of the number a_n of T-tetromino tilings of a $4n \times 4n$ square? It is not hard to show that there exists a limit $c = \lim \log a_n/n^2$ as $n \to \infty$. It would be interesting to find upper and lower bounds on c similar to that in [11, 5].

Acknowledgements

We are grateful to the Makarychev brothers for sending us their preprint [10], and to David Karger, Jim Propp, and Van H. Vu for interesting conversations and helpful remarks. We are also thankful to referees for their comments and a suggestion to strengthen Theorem 10 to show that P_Γ is distributive.

The first author was supported by a National Defense Science and Engineering Graduate Fellowship. The second author was partially supported by NSA and NSF.

12 Appendix

12.1 The ice graph

Ice graphs are another type of directed graph which can be associated with a tiling. These graphs, and their associated height functions, provide another means of proving local connectivity for regions $\Gamma \in \mathcal{Q}$.

For a region $\Gamma \in \mathcal{Q}$, let B_Γ be the set of type-B points in Γ or $\partial\Gamma$. A directed graph on B_Γ is called an *ice graph* if it satisfies the following conditions:

- every two points which lie at opposite corners of the same block of Γ are connected with an edge, either one direction or the other, but not both, and
- every vertex has equal indegree and outdegree.

This notion has been explored by Eloranta [4] and others.

Let $\mathcal{I}_\Gamma$ denote the set of all ice graphs of a region Γ. Call a vertex *alternating* if it is incident to four edges which are oriented "in, out, in, out", in alternating order. Let $z(G)$ be the number of alternating vertices in an ice graph G.

In [10] the Makarychev brothers constructed a map $\mu : \mathcal{T}_\Gamma \to \mathcal{I}_\Gamma$ as follows.

For a tiling $\tau \in \mathcal{T}_\Gamma$, define a directed graph on B_Γ as follows. Observe that within each block, three squares belong to one T-tetromino, while one square, call it the *oddball*, belongs to a different T-tetromino. By inspection, we see that the oddball must be incident to a type-B point, rather than a type-A point. For each block, include a directed edge from the point next to the oddball square to the opposite corner of the block (see Figure 17). Define $\mu(\tau)$ to be the directed graph which results.

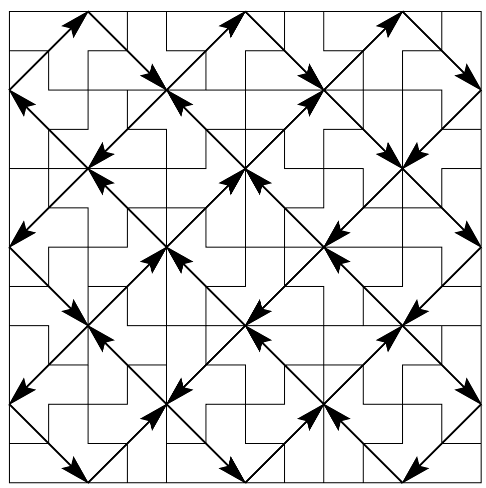

Figure 17: A tiling τ, and the ice graph $\mu(\tau)$.

Lemma 20 (K. and Y. Makarychev) *For any region $\Gamma \in \mathcal{Q}$, the map μ is a surjection from $\mathcal{T}_\Gamma$ to $\mathcal{I}_\Gamma$, in which every ice graph G is the image of $2^{z(G)}$ tilings.*

Sketch of proof: First let us show that $\mu(\tau)$ is an ice graph. Every edge connects two opposite corners of some block, so this graph will have edges in the correct places. Notice that each type-B point is adjacent to exactly two oddballs (recall Figure 5), unless the point is on $\partial\Gamma$, in which case it is adjacent to only one. Therefore, every vertex has equal indegree and outdegree. So $\mu(\tau)$ is in fact an ice graph.

Now we just need to show that every ice graph G comes from exactly $2^{z(G)}$ tilings. Take a vertex of G. If the vertex is on $\partial\Gamma$, there is only one way to place the tile which touches this vertex (see Figure 18).

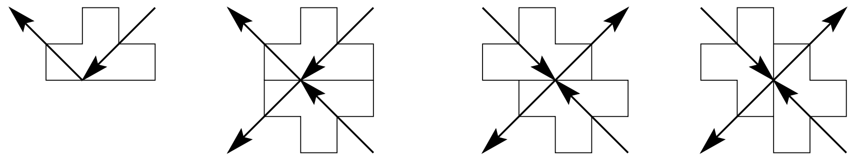

Figure 18: A boundary vertex, a nonalternating vertex, and the two options for an alternating vertex.

Similarly, if the vertex is not on the boundary, and not alternating, there is only one way to place the two tiles which touch this vertex. However, if the vertex is alternating, there are two ways to place the tiles around the vertex. The squares covered by the two tiles are the same in either case, so the decision of which one to use does not affect the rest of the tiling. Hence there are $2^{z(G)}$ ways to convert an ice graph G into a tiling. ■

Lemma 21 *If $\tau_1, \tau_2 \in \mathcal{T}_\Gamma$ are tilings such that $\mu(\tau_1) = \mu(\tau_2)$, then τ_1 and τ_2 are local-move equivalent.*

Sketch of proof: As we just saw, the only way in which these tilings may differ is in the way the tiles next to alternating points are arranged. Converting one such configuration into the other is done by performing a 2-move. Each tile is adjacent to only one type-B point, so these moves are disjoint and can be done independently of each other. So one can convert any such tiling into any other by a sequence of 2-moves. ■

12.2 Height on the ice graph

For a region $\Gamma \in \mathcal{Q}$, let A_Γ be the set of type-A points in Γ or $\partial\Gamma$. Say that a function $f : A_\Gamma \to \mathbb{Z}$ is an *ice-height function* if it satisfies the following conditions:

- $f(x) = 0$ for all points $x \in \partial\Gamma$, and
- $|f(x) - f(y)| = 1$ whenever x and y are adjacent (differ by 2 in each coordinate).

Let $\mathcal{J}_\Gamma$ denote the set of all ice-height functions of a region Γ.

Theorem 22 *For any region $\Gamma \in \mathcal{Q}$, we have $|\mathcal{J}_\Gamma| = |\mathcal{I}_\Gamma|$.*

We define a map $\nu : \mathcal{I}_\Gamma \to \mathcal{J}_\Gamma$ as follows. Let $G \in \mathcal{I}_\Gamma$ be an ice graph. Define a function f° on the faces of G by the following rules. Let f° have the value 0 on the unbounded face of G. As we pass an edge of the graph, if the edge is oriented left-to-right as we pass it, let the value of f° increase by 1. (Similarly, if the edge is oriented right-to-left, let the value of f° decrease by 1.) Now define $f : A_\Gamma \to \mathbb{Z}$ by letting $f(x)$ equal the value of f° on the face in which x lies (see Figure 19). Define $\nu(G)$ to be this function f.

Theorem 22 will follow from the following lemma.

Lemma 23 *For any region $\Gamma \in \mathcal{Q}$, the map ν is a bijection between $\mathcal{I}_\Gamma$ and $\mathcal{J}_\Gamma$.*

Proof: Let G be an ice graph, and let f be $\nu(C)$. The function f is well-defined for the same reason that the height function for the chain graph is well-defined—because every vertex has equal indegree and outdegree. It is clear that such a function meets the criteria for being an ice-height function.

From an ice-height function f, one can reconstruct the ice graph $G = \nu^{-1}(f)$ by directing every edge so the face with greater height is on the left. Since the net change in height going around any vertex is 0, every vertex will have equal indegree and outdegree, thus the graph so constructed will be a valid ice graph. ■

For ease of notation, define $\xi(\tau) = \nu(\mu(\tau))$. For a region $\Gamma \in \mathcal{Q}$, the map ξ is the canonical bijection between $\mathcal{T}_\Gamma$ and $\mathcal{J}_\Gamma$.

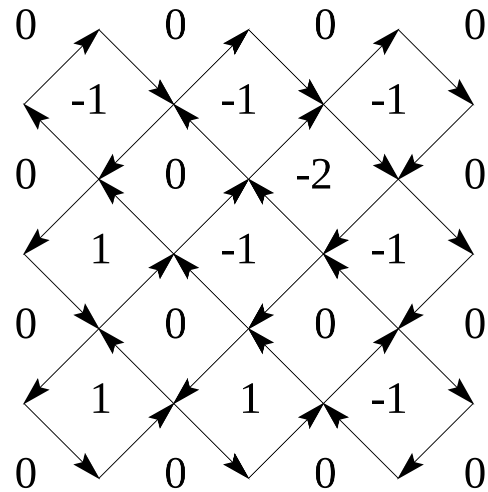

Figure 19: An ice graph G, and the function $f = \nu(G)$.

Lemma 24 *Let $\Gamma \in \mathcal{Q}$ and let $\tau_1, \tau_2 \in \mathcal{T}_\Gamma$ be tilings of Γ. If the tilings τ_1 and τ_2 differ by a 2-move, then $\xi(\tau_1) = \xi(\tau_2)$. If the tilings τ_1 and τ_2 differ by a 4-move, then $\xi(\tau_1)$ and $\xi(\tau_2)$ differ by 2 on some point, and are the same everywhere else. If f_1 and f_2 are ice-height functions which differ by 2 on some point and are the same everywhere else, then there exist tilings τ_1 and τ_2 such that $\xi(\tau_1) = f_1, \xi(\tau_2) = f_2$, and τ_1 and τ_2 differ by a 4-move.*

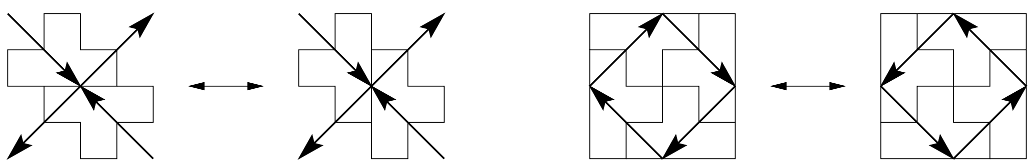

Figure 20: The effect of local moves on the ice graph.

Sketch of proof: A 2-move can only occur at an alternating type-B point, so if τ_1 and τ_2 differ by a 2-move, then $\mu(\tau_1) = \mu(\tau_2)$, so $\xi(\tau_1) = \xi(\tau_2)$ (see Figure 20).

If τ_1 and τ_2 differ by a 4-move, then $\mu(\tau_1)$ and $\mu(\tau_2)$ differ by the reversal of a directed 4-cycle, thus $\xi(\tau_1)$ and $\xi(\tau_2)$ will differ by 2 on the point inside that 4-cycle, and be the same everywhere else.

Now suppose f_1 and f_2 are ice-height functions such that $f_1(x) = h + 1$ and $f_2(x) = h - 1$, but $f_1 = f_2$ everywhere else. We must then have $f_1(y) = f_2(y) = h$ for the neighbors y of x. So x will be surrounded by a counterclockwise directed 4-cycle in the ice graph corresponding to f_1, and a clockwise directed 4-cycle in the ice graph corresponding to f_2. The problem is that a tiling which corresponds to f_1 may look like the left side of Figure 21. However, in such a case, there is always another tiling (which differs from the original by some 2-moves) such that a 4-move can be applied. ■

For ice-height functions $f_1, f_2 \in \mathcal{J}_\Gamma$, say that f_1 and f_2 differ by a 4-move if there exist tilings $\tau_1, \tau_2 \in \mathcal{T}_\Gamma$ which differ by a 4-move such that $\xi(\tau_1) = f_1$ and $\xi(\tau_2) = f_2$. By the previous Lemma, one can see that performing a 4-move on an ice-height function f is equivalent to increasing or decreasing its value by 2 at some point. Of course, such a move may only be applied if the function that results is a valid ice-height function.

Notice that for tilings τ_1 and τ_2, having $\xi(\tau_1)$ and $\xi(\tau_2)$ differ by a 4-move does not imply that τ_1 and τ_2 differ by a 4-move. However, it does imply that there exist tilings τ_1' and τ_2' which differ by a 4-move such

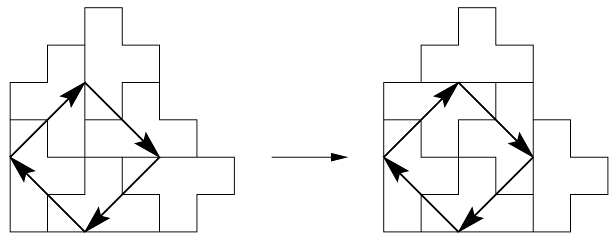

Figure 21: A tiling where a 4-move cannot be applied, and one where it can.

that $\xi(\tau_1') = \xi(\tau_1)$ and $\xi(\tau_2') = \xi(\tau_2)$. It then follows, by Lemmas 21 and 23, that τ_1 is local-move equivalent to τ_1' and τ_2 is local-move equivalent to τ_2'. Hence τ_1 and τ_2 will be local-move equivalent whenever $\xi(\tau_1)$ and $\xi(\tau_2)$ differ by a 4-move, or more generally, by a sequence of 4-moves.

Theorem 1 will now easily follow from the following Lemma.

Lemma 25 *Let $\Gamma \in \mathcal{Q}$, and let $f_1, f_2 \in \mathcal{J}_\Gamma$ be ice-height functions. It is always possible to convert f_1 into f_2 by performing a sequence of 4-moves.*

Proof: For any region, there will be a unique ice-height function f_0 whose value at each point is either 0 or 1. (Each face is either "even" or "odd", depending on how many steps from the exterior it is, thus each even face will have the value 0, and each odd face will have the value 1.) It will be sufficient to show that any ice-height function f can be transformed into f_0. Suppose $f(x) > 1$ for some point x. Let x be the point where f attains its largest value, call it h (if there are several possible points, choose any one). We must then have $f(y) = h - 1$ for the neighbors y of x. Thus we can perform a 4-move, and decrease $f(x)$ to $h - 2$. Repeat this process until f attains no values greater than 1. Now if there are points x where $f(x) < 0$, find the one where f attains its minimum. We can perform a 4-move to increase $f(x)$ by 2. We repeat this until f attains no values less than 0. Now $0 \leq f(x) \leq 1$ for all x, so we are done. ∎

References

[1] D. Aldous and J. Fill. *Reversible Markov Chains and Random Walks on Graphs.* monograph in preparation.

[2] T. Chaboud. Domino tiling in planar graphs with regular and bipartite dual. *Theoret. Comput. Sci.*, 159(1):137–142, 1996. Selected papers from the "GASCOM '94" (Talence, 1994) and the "Polyominoes and Tilings" (Toulouse, 1994) Workshops.

[3] H. Cohn, N. Elkies, and J. Propp. Local statistics for random domino tilings of the Aztec diamond. *Duke Math. J.*, 85(1):117–166, 1996.

[4] K. Eloranta. Diamond ice. *J. Statist. Phys.*, 96(5-6):1091–1109, 1999.

[5] M. E. Gegúndez, A. Márquez, and P. Revuelta. The Tutte polynomial on the square lattice. (preprint).

[6] S. W. Golomb. *Polyominoes.* Georges Allen and Unwin Ltd, London, 1966.

[7] M. Jerrum and A. Sinclair. Approximating the permanent. *SIAM J. Comput.*, 18(6):1149–1178, 1989.

[8] C. Kenyon and R. Kenyon. Tiling a polygon with rectangles. *Proc. 33-rd FOCS*, pages 610–619, 1992.

[9] M. Luby, D. Randall, and A. Sinclair. Markov chain algorithms for planar lattice structures. *SIAM J. Comput.*, 31(1):167–192 (electronic), 2001.

[10] K. Makarychev and Y. Makarychev. Proof of Pak's conjecture on tilings by T-tetrominoes (in Russian), `http://home.earthlink.net/~makarychev/tetromino/tetro.htm`.

[11] C. Merino and D. J. A. Welsh. Forests, colorings and acyclic orientations of the square lattice. *Ann. Comb.*, 3(2-4):417–429, 1999.

[12] C. Moore, I. Rapaport, and E. Remila. Tiling groups for wang tiles. *Proc. 13-th SODA*, pages 402–411, 2002.

[13] I. Pak. Tile invariants: New horizons. *Theor. Comp. Sci.* (to appear).

[14] J. Propp. Generating random elements of finite distributive lattices. *Electron. J. Combin.*, 4(2):Research Paper 15, approx. 12 pp. (electronic), 1997. The Wilf Festschrift (Philadelphia, PA, 1996).

[15] J. Propp and D. Wilson. Exact sampling with coupled Markov chains and applications to statistical mechanics. In *Proceedings of the Seventh International Conference on Random Structures and Algorithms (Atlanta, GA, 1995)*, volume 9, pages 223–252, 1996.

[16] J. Propp and D. Wilson. Coupling from the past: a user's guide. In *Microsurveys in discrete probability (Princeton, NJ, 1997)*, volume 41 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 181–192. Amer. Math. Soc., Providence, RI, 1998.

[17] E. Rémila. Tiling groups: new applications in the triangular lattice. *Discrete Comput. Geom.*, 20(2):189–204, 1998.

[18] Alistair Sinclair. *Algorithms for random generation and counting.* Progress in Theoretical Computer Science. Birkhäuser Boston Inc., Boston, MA, 1993.

[19] W. Thurston. Conway's tiling groups. *Amer. Math. Monthly*, 97(8):757–773, 1990.

[20] D. L. Vertigan and D. J. A. Welsh. The computational complexity of the Tutte plane: the bipartite case. *Combin. Probab. Comput.*, 1(2):181–187, 1992.

[21] D. W. Walkup. Covering a rectangle with T-tetrominoes. *Amer. Math. Monthly*, 72:986–988, 1965.

[22] D. J. A. Welsh. *Complexity: knots, colourings and counting*, volume 186 of *London Mathematical Society Lecture Note Series.* Cambridge University Press, Cambridge, 1993.

Note:

After the paper was finished there has been several subsequent developments. First, Michel Las Vergnas established a connection between our Theorem 13 and his results in the paper *"On the evaluation at* $(3,3)$ *of the Tutte polynomial of a graph"*, J. Combin. Theory Ser. B, vol. 45 (1988), 367–372. The authors were not aware of this paper, but the connections is in the same spirit as the appendix.

The authors later generalized and extended Theorem 13 in this paper to plane and ribbon graphs in the recent preprint *"Combinatorial evaluations of the Tutte polynomial"*, which is available from:
`http://www-math.mit.edu/~pak/research.html`