Cut-and-paste of quadriculated disks and arithmetic properties of the adjacency matrix

Nicolau C. Saldanha and Carlos Tomei

April 29, 2009

Abstract

We define *cut-and-paste*, a construction which, given a quadriculated disk obtains a disjoint union of quadriculated disks of smaller total area. We provide two examples of the use of this procedure as a recursive step. Tilings of a disk Δ receive a *parity*: we construct a perfect or near-perfect matching of tilings of opposite parities. Let B_Δ be the black-to-white adjacency matrix: we factor $B_\Delta = L\tilde{D}U$, where L and U are lower and upper triangular matrices, $\tilde{D}$ is obtained from a larger identity matrix by removing rows and columns and all entries of L, $\tilde{D}$ and U are equal to 0, 1 or -1.

1 Introduction

In this paper, a *square* is a topological disk with four privileged boundary points, the *vertices*; the boundary of the square consists of four *edges*. A *quadriculated disk* Δ is a closed topological disk formed by the juxtaposition along edges of finitely many squares such that interior vertices belong to precisely four squares: it may be considered as a closed subset of the plane $\mathbb{R}^2$ tiled by quadrilaterals. A simple example is the $n \times m$ rectangle divided into unit squares, another is shown in Figure 1.

Given Δ, we define the planar *dual graph* $\mathcal{G}_\Delta$: vertices of $\mathcal{G}_\Delta$ correspond to squares in Δ and two vertices of $\mathcal{G}_\Delta$ are adjacent if their corresponding squares share an edge. Quadriculated disks are *bi-colored*: the squares are black and white

2000 *Mathematics Subject Classification.* Primary 05B45, 05C70; Secondary 05B20, 05C50.
Keywords and phrases Quadriculated disk, matchings, tilings by dominoes, dimers.

The authors gratefully acknowledge the support of CNPq, Faperj. The first author thanks the kind hospitality of The Ohio State University during part of the time when this paper was written.

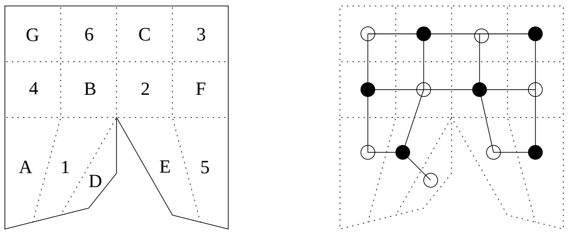

$$B_\Delta = \begin{pmatrix} 1 & 1 & & 1 & & & \\ & 1 & 1 & & 1 & 1 & \\ & & 1 & & & 1 & \\ 1 & 1 & & & & & 1 \\ & & & & 1 & 1 & \\ & 1 & 1 & & & & 1 \end{pmatrix}$$

Figure 1: A quadriculated disk, its dual graph and its black-to-white matrix

in a way that squares with a common edge have opposite colors (equivalently, $\mathcal{G}_\Delta$ is bipartite). Label the black (resp. white) squares of a quadriculated disk Δ by $1, 2, \ldots, b$ (resp. $1, 2, \ldots, w$). The $b \times w$ *black-to-white* (adjacency) matrix B_Δ has (i, j) entry $b_{ij} = 1$ if the i-th black and j-th white squares share an edge and $b_{ij} = 0$ otherwise. Figure 1 is an example of black-to-white matrix; black and white squares are labeled by numbers and letters, respectively. Throughout the paper, blank matrix entries equal 0. For a labeling in which black vertices come first, the adjacency matrix of $\mathcal{G}_\Delta$ is

$$\begin{pmatrix} 0 & B_\Delta \\ B_\Delta^T & 0 \end{pmatrix}.$$

The following result [1] indicates an unexpected spectral rigidity of B_Δ.

Theorem 1 *Let Δ be a quadriculated disk with $b = w$ and black-to-white matrix B_Δ. Then* $\det(B_\Delta)$ *equals* 0, 1 *or* -1.

This result admits a combinatorial interpretation. A *domino tiling* τ of Δ is a decomposition of Δ as a union of *dominos* (i.e., 2×1 rectangles) with disjoint interior. Let $\mathcal{T}_\Delta$ be the set of domino tilings of Δ. There is a natural parity function on $\mathcal{T}_\Delta$ (see Section 4) and the determinant $\det(B_\Delta)$ counts tilings with a sign given by parity. The theorem above thus says that there exists a *quasi-perfect matching* in $\mathcal{T}_\Delta$, i.e., a correspondence between even and odd tilings leaving out at most one element of $\mathcal{T}_\Delta$, the *loner*. We provide a new, (quasi-) bijective proof of Theorem 1 by constructing a quasi-perfect matching in the bipartite set $\mathcal{T}_\Delta$.

We extend Theorem 1 in a different, more algebraic, direction. A rectangular matrix $\tilde{D}$ is a *defective identity* if it can be obtained from the identity matrix by adding rows and columns of zeros. For a $n \times m$ matrix A, an $L\tilde{D}U$ decomposition of A is a factorization $A = L\tilde{D}U$ where L (resp. U) is $n \times n$ (resp. $m \times m$) lower (resp. upper) invertible and $\tilde{D}$ is a defective identity.

Theorem 2 *Let Δ be a quadriculated disk with at least two squares. For an appropriate labeling of its squares, the black-to-white matrix B_Δ admits an $L\tilde{D}U$ decomposition whose factors have all entries equal to* 0, 1 *or* -1.

Thus, for example, the matrix B_Δ in Figure 1 admits the decomposition

$$\begin{pmatrix} 1 & & & & & \\ & 1 & & & & \\ & & 1 & & & \\ 1 & & & -1 & & \\ & & & & 1 & \\ & 1 & & & -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & & & & & & \\ & 1 & & & & & \\ & & 1 & & & & \\ & & & 1 & & & \\ & & & & 1 & 0 & \\ & & & & & & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 & & 1 & & & \\ & 1 & 1 & & 1 & 1 & \\ & & 1 & & & 1 & \\ & & & 1 & & & -1 \\ & & & & 1 & 1 & \\ & & & & & 1 & \\ & & & & & & 1 \end{pmatrix}.$$

Both the construction of the quasi-perfect matching and the proof of Theorem 2 use *cut-and-paste*, a recursive operation on quadriculated disks. A quadriculated disk Δ is *cut* along *diagonals* and *pasted* to obtain a disjoint union of smaller disks $\Delta'_1, \ldots, \Delta'_d$, often with $d = 1$ (Lemma 3.1). Every nontrivial quadriculated disk admits cut-and-paste (Proposition 2.3).

The proof of Theorem 2 relies on a procedure to convert $L\tilde{D}U$ decompositions of $B_{\Delta'_1}, \ldots, B_{\Delta'_d}$ into a similar decomposition of B_Δ (Lemma 5.2). The proof yields a fast algorithm to obtain the appropriate labeling of vertices, the matrices in the factorization, $\det(B_\Delta)$ and $\operatorname{rank}(B_\Delta)$.

In Section 2 we present the facts about *diagonals* of quadriculated disks which will be used in Section 3 to describe *cut-and-paste*. In Section 4 we construct the quasi-perfect matching. The inductive step in the proof of Theorem 2, the algebraic counterpart of cut-and-paste, is the main topic of Section 5. Finally, in Section 6, we study *boards*, quadriculated disks which are subsets of the quadriculated plane $\mathbb{Z}^2 \subset \mathbb{R}^2$; Theorem 3 states that cut-and-paste can be performed within this smaller class.

Counting tilings with sign given by parity (as in Theorem 1) corresponds to the case $q = -1$ of the q-counting of domino tilings with respect to *height* or *volume* as in [7], [2] and [6]. In a similar vein, [5] extends Theorem 1 to quadriculated annuli by introducing a polynomial which counts tilings with respect to yet another integral parameter, the *flux*. It is not clear whether the cut-and-paste procedure can be extended to take such parameters into account.

2 Diagonals

A *corner* of a quadriculated disk Δ is a boundary point which is a vertex of a single square. A *pre-diagonal* of length $k > 0$ of Δ is a sequence of vertices $v_0 v_1 \ldots v_k$ such that

(i) v_0 is a corner, $v_1, v_2, \ldots, v_{k-1}$ are interior vertices;

(ii) consecutive vertices v_i and v_{i+1}, $i = 0, \dots, k-1$, are opposite vertices of a square $s_{i+1/2}$;

(iii) consecutive squares $s_{i-1/2}$ and $s_{i+1/2}$, $i = 1, \dots, k-1$, have a single vertex in common (which is v_i);

(iv) the vertices v_i and the squares $s_{i+1/2}$, $i = 0, \dots, k-1$, are distinct.

A *diagonal* is a maximal pre-diagonal (under inclusion). More geometrically, we may think of a diagonal as a line $\ell = \ell(v_0, s_{1/2}, v_1, \dots, s_{k-1/2}, v_k)$ connecting v_0, the center of $s_{1/2}$, v_1, the center of $s_{3/2}$ and so on up to v_k. The squares $s_{1/2}, \dots, s_{k-1/2}$ are the *squares of the diagonal*. Usually, the vertices $s_{1/2}, \dots, s_{k-1/2}$ form a cut set of the dual graph $\mathcal{G}_\Delta$. Diagonals, being sequences of vertices, are naturally oriented. Figure 2 shows examples of diagonals; vertices and squares of δ_1 are indicated.

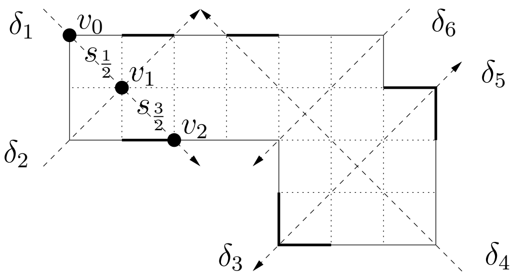

Figure 2: A quadriculated disk and its six diagonals

Proposition 2.1 *Given a corner v_0 of Δ there is a unique diagonal starting at v_0. Furthermore, all diagonals end at boundary points.*

Proof: In principle, there are three types of diagonals: the vertex v_k may coincide with some v_i, $i < k$ (Figure 3, (a)), the square $s_{k+1/2}$ may coincide with some $s_{i+1/2}$, $i < k$ (Figure 3, (b)) or v_k may be a boundary vertex of Δ. Existence and uniqueness of a diagonal δ starting at the corner v_0 follows from finiteness. The reader may check that self-intersection would happen at right angles, as in the figure. Bicoloring of squares and vertices of Δ, as in Figure 3, yields a contradiction in either case. ■

Let δ be a diagonal of a quadriculated disk Δ associated to the line $\ell = \ell(v_0, s_{1/2}, \dots, v_k) \subset \Delta$. Given a vertex v of $\Delta \smallsetminus \ell$, draw a smooth curve $\gamma : [0,1] \to \Delta$, $\gamma(0) = v$, $\gamma(1) \in \ell$, $\gamma(t) \in \Delta \smallsetminus \ell$ for $t < 1$ and $\gamma'(1)$ transversal to ℓ. We say that v is to the *left* (resp. *right*) of δ if $\det(v_1 - v_0, \gamma'(1))$ is negative (resp. positive). The existence of the curve γ follows from the fact that Δ is path-connected. A vertex v is not simultaneously to the left and right of δ: indeed,

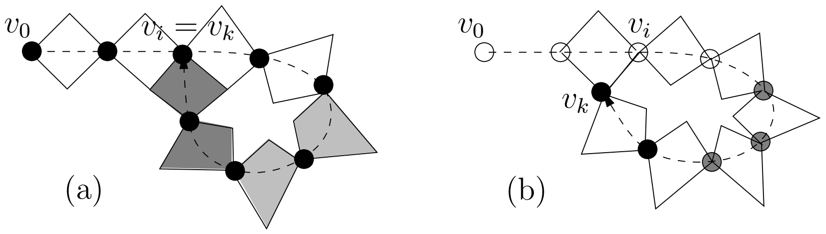

Figure 3: Impossible diagonals

if $\gamma_l, \gamma_r : [0,1] \to \Delta$ satisfy the hypothesis above and $\det(v_1 - v_0, \gamma_l'(1)) < 0 < \det(v_1 - v_0, \gamma_r'(1)$ then juxtaposition of γ_l and time-reversal of γ_r obtains a loop which crosses ℓ exactly once, a contradiction.

In the next section we will use diagonals to cut-and-paste. Not all diagonals are suitable for this construction. Call the two edges of $s_{k-1/2}$ ending at v_k *terminal* edges. A diagonal $v_0 \ldots v_k$ is a *good diagonal* if at least one terminal edge is contained in the boundary of Δ. In Figure 2, δ_6 is the only bad diagonal. A square has four diagonals, all good.

To prove the existence of good diagonals, we use a quadriculated version of the Gauss-Bonnet theorem. Let V be the number of vertices of Δ and write $V = V_I + V_1 + V_2 + \cdots + V_r$ where V_I counts interior vertices and V_r is the number of boundary vertices belonging to exactly r squares. Notice that V_1 is the number of corners of Δ.

Lemma 2.2 $V_1 - V_3 - 2V_4 - \cdots - (r-2)V_r = 4.$

Proof: Let E and F be the number of edges and faces (i.e., squares) of Δ. Write $E = E_I + E_B$, where E_I (resp. E_B) counts interior (resp. boundary) edges. Clearly, $4F = 2E_I + E_B = 2E - E_B$ and therefore $4E = 8F + 2E_B$. Also, $4F = 4V_I + V_1 + 2V_2 + \cdots + rV_r = 4V - (3V_1 + 2V_2 + \cdots + (4-r)V_r)$ and $4V = 4F + (3V_1 + 2V_2 + \cdots + (4-r)V_r)$. By Euler, $4V - 4E + 4F = 4$. Substituting the above formulas and using $E_B = V_1 + V_2 + \cdots + V_r$ we have the desired identity. ∎

Proposition 2.3 *Any quadriculated disk Δ admits at least four good diagonals.*

Proof: Each vertex counted in V_1 is a starting corner for a diagonal: we have to prove that at least four of these V_1 diagonals are good. Each vertex counted in V_3, for example, is the endpoint of at most three diagonals of which only one is declared bad. More generally, we have at most $V_1 - 4 = V_3 + 2V_4 + \cdots + (r-2)V_r$ bad ends and we are done. ∎

3 Geometric cut-and-paste

We are ready to perform *cut-and-paste* along a good diagonal. A good diagonal $v_0 \ldots v_k$ is *balanced* if exactly one terminal edge is contained in the boundary of Δ. Diagonals δ_1, δ_2 and δ_4 in Figure 2 are balanced; δ_3 and δ_5 are unbalanced.

In Figure 4 we illustrate the cut-and-paste procedure ∂_δ on a quadriculated disk Δ and its dual graph $\mathcal{G}_\Delta$, where δ is an unbalanced diagonal. The operation removes the shaded squares and identifies edges to obtain a new quadriculated disk $\Delta' = \partial_\delta(\Delta)$. Another choice of shaded squares for the same good diagonal δ is indicated in the right and obtains the same quadriculated disk Δ'. In the left (resp. right), squares C and D (resp. A and B) take over the space vacated by A and B (resp. C and D).

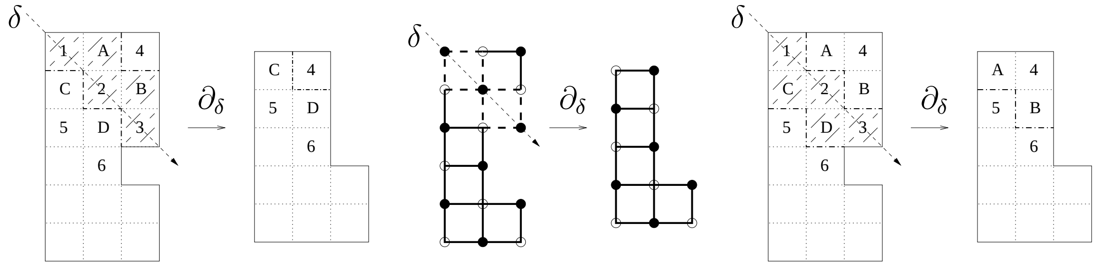

Figure 4: Cut-and-paste along the unbalanced diagonal δ of length $k = 3$

The balanced case shown in Figure 5 is a little different. It turns out that a similar construction with another choice of zig-zag is not appropriate for future purposes.

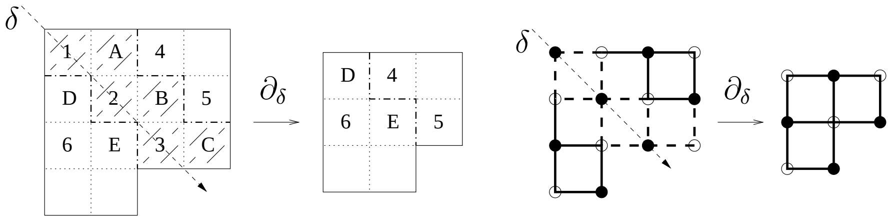

Figure 5: Cut-and-paste along a balanced diagonal, $k = 3$

In the dual graph $\mathcal{G}_\Delta$, cut-and-paste removes the cut set of vertices (of the graph) associated with squares of the diagonal δ and identifies vertices on both sides: vertices left without partners at the end of an unbalanced diagonal are also deleted. This point of view is more symmetric and does not require the specification of zig-zags.

Notice that the extreme vertex v_k of a balanced diagonal may belong to more than two squares, as in Figure 1. This is innocuous, as we shall see.

Cut-and-paste allows for recursive proofs and constructions in the class of finite disjoint unions of quadriculated disks. As we shall prove in Lemma 3.1, given a quadriculated square Δ and a good diagonal δ, cut-and-paste obtains a quadriculated region $\tilde{\Delta}'$ which consists of quadriculated disks $\Delta'_1, \ldots, \Delta'_d$, possibly joined by points. The process of passing from $\tilde{\Delta}'$ to $\Delta' = \Delta'_1 \sqcup \cdots \sqcup \Delta'_d$ is called *detaching*. Clearly, Δ' has fewer squares than Δ. In the two previous examples, $d = 1$; in Figure 6, $d = 3$.

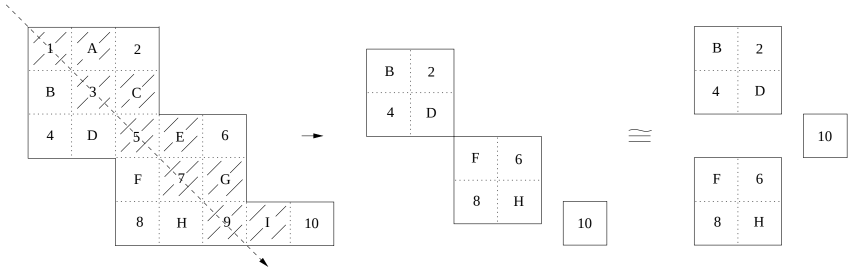

Figure 6: Cut-and-paste may produce a disjoint union of disks

In a somewhat degenerate case, $\Delta' = \varnothing$ if and only if Δ consists of one or two squares. Also, if δ is an unbalanced diagonal of length $k = 1$, the quadriculated disk Δ' is obtained from Δ by deleting two squares.

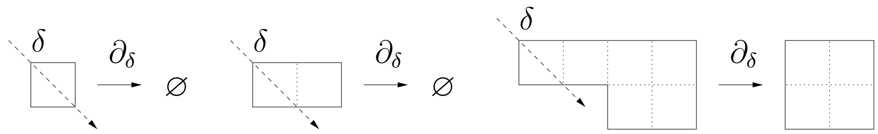

Figure 7: Cut-and-paste in extreme situations

Lemma 3.1 *Let Δ be a quadriculated disk with a good diagonal δ of length $k > 1$. Let Δ' be obtained from Δ by cut-and-paste along δ (and detaching): Δ' is a disjoint union of quadriculated disks.*

We use a notation for vertices and squares near a good diagonal. Squares immediately to the left (resp. right) of the diagonal are labelled $s^l_1, s^l_2, \ldots$ (resp. $s^r_1, s^r_2, \ldots$). Similarly, vertices to the left (resp. right) are labelled $v^l_{1/2}, v^l_{3/2}, \ldots$ (resp. $v^r_{1/2}, v^r_{3/2}, \ldots$). Thus, in Figure 4, $s^l_1 = A$, $s^l_2 = B$, $s^r_1 = C$, $s^r_2 = D$; in Figure 5, $s^l_1 = A$, $s^l_2 = B$, $s^l_3 = C$, $s^r_1 = D$, $s^r_2 = E$. The squares deleted in the cut-and-paste construction (dashed in the figures) are $s_{1/2}, s^x_1, s_{3/2}, \ldots, s^x_{k-1}, s_{k-1/2}$ and, in the balanced case, s^x_k; here $x = l$ or $x = r$. Let Δ^r (resp. Δ^l) be the closed regions to the right (resp. left) of the deleted squares. Attach Δ^l to Δ^r by identifying edges in order to obtain a quadriculated region $\tilde{\Delta}'$.

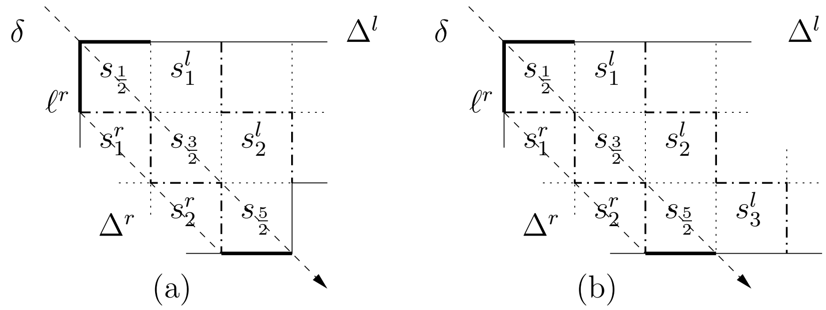

Figure 8: Notation for cut-and-paste; unbalanced and balanced cases

Proof: Assume without loss that cut-and-paste along δ deletes the squares $s_{1/2}, s^l_1, \ldots, s^l_{k-1}, s_{k-1/2}$ and, if δ is unbalanced, s^l_k. We claim that Δ^r is non-empty, path-connected and simply connected. Indeed, the squares $s^r_1, \ldots, s^r_{k-1}$ exist (since $v_1, \ldots, v_{k-1}$ are interior points, $k > 1$). To show that Δ^r is path-connected, it suffices to join by a path in Δ^r any point $x \in \Delta^r$ to the line $\ell^r = (v^r_{1/2}, s^r_1, \ldots, s^r_{k-1}, v^r_{k-1/2})$. Notice that the edges $v_0 v^r_{1/2}$ and $v_k v^r_{k-1/2}$ are in the boundary of Δ. If $x \in \Delta^r$ lies between ℓ and ℓ^r then x belongs to one of the squares $s^r_1, \ldots, s^r_{k-1}$ and the path is easy to construct. Otherwise, take $\gamma : [0,1] \to \Delta$ as in the definition of left and right of δ in Section 2; γ must cross ℓ^r and a restriction of γ yields the required path. As to simple connectivity, take a simple closed curve α contained in Δ^r and therefore in Δ. By Jordan's Theorem, α encloses a disk A. Since Δ is simply connected, $A \subset \Delta$. Also, a path in A from $x \in A$ to α guarantees that x and α are on the same side of δ.

The region Δ^l may be disconnected or even empty. On the other hand, the argument above shows that its connected components are simply connected. Thus, $\tilde{\Delta}'$ is obtained by gluing the simply connected pieces Δ^r and the components of Δ^l: we must now study the gluing process. Let ζ^r and ζ^l be the zig-zag lines $v^r_{1/2} v_1 v^r_{3/2} \ldots v_{k-1} v^r_{k-1/2}$ and $v^l_{1/2} v^{ll}_1 v^l_{3/2} \ldots v^{ll}_{k-1} v^l_{k-1/2}$, where v^{ll}_i is the left-most vertex of s^l_i. Cut-and-paste obtains $\tilde{\Delta}'$ by gluing Δ^r and Δ^l along ζ^r and ζ^l. Notice that ζ^r is contained in the boundary of Δ^r. It is possible, however, that parts of ζ^l are part of the boundary of Δ and not in Δ^l.

We claim that, given a connected component D of Δ^l, its intersection with ζ^l is either empty or path-connected. In other words, for any two points $x_0, x_1 \in D \cap \zeta^l$, the segment $[x_0, x_1] \subset \zeta^l$ between x_0 and x_1 is contained in D. Indeed, there is a curve α in D joining x_0 and x_1. Juxtaposition of α and $[x_0, x_1]$ obtains a closed curve in Δ. As before, simple connectivity of Δ implies that the region surrounded by this closed curve is contained in Δ and therefore in Δ^l and D, completing the proof of the claim.

The claims and Seifert-Van Kampen's Theorem ([3]) imply that each connected component of $\tilde{\Delta}'$ is simply connected. Detaching guarantees that each connected component of Δ' is a simply connected surface with boundary — a disk. ■

4 A bijective proof of Theorem 1

A nonzero entry b_{ij} of the black-to-white matrix B_Δ corresponds to a domino contained in Δ: the indices i and j indicate the black and white squares in the domino and $b_{ij} \neq 0$ when these two squares are adjacent. A *domino tiling* of Δ is a decomposition of Δ as a union of dominos with disjoint interiors; let $\mathcal{T}_\Delta$ be the set of all domino tilings of Δ. A nonzero monomial of the black-to-white matrix B_Δ corresponds to some $\tau \in \mathcal{T}_\Delta$. Indeed, the dominos associated with the entries cover Δ and their interiors are disjoint. Equivalently, for a labeling of black and white squares by $\{1, 2, \ldots, b\}$ and $\{1, 2, \ldots, w\}$, we may consider a tiling τ as a function $\pi : \{1, 2, \ldots, w\} \to \{1, 2, \ldots, b\}$ with $\pi(j) = i$ if and only if the i-th black square and the j-th white square form a domino in τ. With $b = w$, this provides an identification between $\mathcal{T}_\Delta$ and a subset of the symmetric group S_w.

The above identification endows a tiling with parity (or sign). Tilings differing by a *flip* (i.e., by exactly two dominos forming a 2×2 square) have opposite parities: if their corresponding permutations are π_1 and π_2 then $\pi_2^{-1}\pi_1$ is a cycle of length 2, interchanging the two white squares in the flip. The combinatorial interpretation of Theorem 1 is that the number of even and odd tilings in $\mathcal{T}_\Delta$ differ by at most 1. In this section we provide a bijective proof of this statement.

More precisely, we present an algorithm that, given a quadriculated disk Δ, obtains a *quasi-perfect matching* in $\mathcal{T}_\Delta$, i.e., a subset $\mathcal{T}^*_\Delta \subseteq \mathcal{T}_\Delta$ whose complement has at most one element, the *loner*, and an involution $\rho : \mathcal{T}^*_\Delta \to \mathcal{T}^*_\Delta$ (i.e., $\rho^2(\tau) = \tau$) inverting parity. The argument proceeds by induction on the number of squares of Δ. The construction of the quasi-perfect matching is trivial if Δ has fewer than 4 squares.

In general, start with a quadriculated disk Δ with $b = w$ and take a good diagonal δ as in Figure 9. Draw and number *wedges* along δ as in the figure; a tiling *respects* a wedge if no domino in the tiling crosses a leg of the wedge. We define a partition $\mathcal{T}_\Delta = \mathcal{D}_\Delta \sqcup \mathcal{R}_\Delta$: a tiling τ belongs to $\mathcal{D}_\Delta$ if and only if τ *disrespects* at least one of the wedges along δ (see [4] for a similar construction with a different purpose). The loner of the quasi-perfect matching, if it exists, will belong to $\mathcal{R}_\Delta$; the sets $\mathcal{D}_\Delta$ and $\mathcal{R}^*_\Delta = \mathcal{R}_\Delta \cap \mathcal{T}^*_\Delta$ will be invariant by ρ. Equivalently, deletion of the edges of $\mathcal{G}_\Delta$ crossing the wedges obtains a subgraph $\mathcal{G}^{\mathcal{R}}_\Delta$: tilings in $\mathcal{T}_\Delta$ (resp. $\mathcal{R}_\Delta$) correspond to matchings in $\mathcal{G}_\Delta$ (resp. $\mathcal{G}^{\mathcal{R}}_\Delta$).

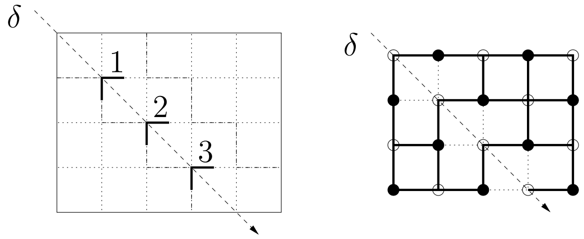

Figure 9: Wedges along a good diagonal and the subgraph $\mathcal{G}_\Delta^{\mathcal{R}}$

We first construct the restriction $\rho|_{\mathcal{D}_\Delta}$. Given $\tau \in \mathcal{D}_\Delta$, assume that the first wedge to be disrespected is the k-th wedge. This means that the first 2×2 square formed by dominos along δ is positioned around that wedge: $\rho(\tau)$ differs from τ by a flip in that square.

There is a natural bijection $\partial : \mathcal{R}_\Delta \to \mathcal{T}_{\Delta'}$, where Δ' is the disjoint union of quadriculated disks obtained from Δ by cut-and-paste along δ. Indeed, for $\tau \in \mathcal{R}_\Delta$, define $\partial(\tau) \in \mathcal{T}_{\Delta'}$ by removing the dominos covering one of the squares $s_{i+1/2}$ along δ and gluing the remaining parts. Given a quasi-perfect matching $\rho' : \mathcal{T}^*_{\Delta'} \to \mathcal{T}^*_{\Delta'}$, define $\mathcal{R}^*_\Delta = \partial^{-1}(\mathcal{T}^*_{\Delta'})$ and $\rho(\tau) = \partial^{-1}(\rho'(\partial(\tau)))$.

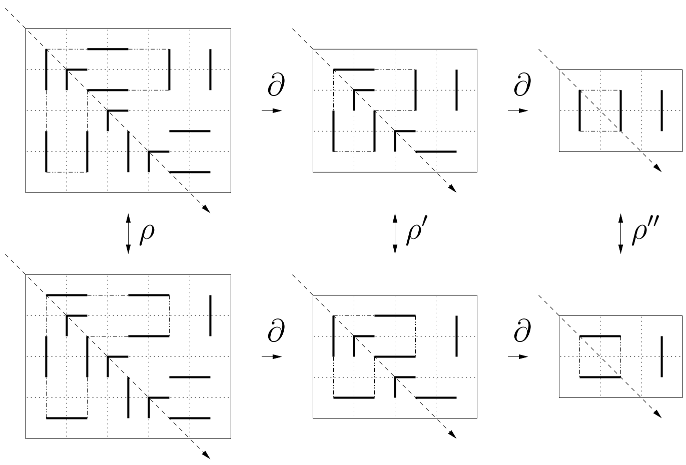

Figure 10: The maps ∂ and ρ

If Δ' is a quadriculated disk, a quasi-perfect matching is obtained by recursion. Otherwise, for the detached collection

$$\Delta' = \Delta'_1 \sqcup \cdots \sqcup \Delta'_d, \quad d > 1,$$

assume (again by recursion) that quasi-perfect matchings $\rho'_i : \mathcal{T}^*_{\Delta'_i} \to \mathcal{T}^*_{\Delta'_i}$ have been obtained for each Δ'_i (possibly with loners). For $\tau' \in \mathcal{T}_{\Delta'}$, let τ'_i be the restriction of τ' to Δ'_i. In order to find $\rho'(\tau')$, search for the smallest i for which $\tau'_i \in \mathcal{T}^*_{\Delta'_i}$ (i.e., τ'_i is *not* a loner); construct $\rho'(\tau')$ by changing τ' in Δ'_i only:

$$\rho'(\tau') = \tau'_1 \sqcup \cdots \sqcup \rho'_i(\tau'_i) \sqcup \cdots \sqcup \tau'_d.$$

A tiling remains unmatched if and only if its restriction to each Δ'_i is a loner: since there is at most one loner in each $\mathcal{T}_{\Delta'_i}$, there is at most one loner in $\mathcal{T}_{\Delta'}$ and ρ' is indeed a quasi-perfect matching.

If the diagonal δ is unbalanced then $b' \neq w'$ and Δ' admits no domino tilings. Consistently, in this case, $\mathcal{R}_\Delta$ is empty: this follows from the impossibility of respecting the last wedge. More generally, if $\Delta' = \Delta'_1 \sqcup \cdots \sqcup \Delta'_d$ and (at least) one of the disks Δ'_i admits no domino tilings then $\mathcal{R}_\Delta$ is empty and we are done.

We must perform a final check: τ and $\rho(\tau)$ are supposed to have opposite parities. This is clear for $\tau \in \mathcal{D}_\Delta$; before we address the issue for $\tau \in \mathcal{R}_\Delta$, we present a few examples.

We follow the construction above in order to compute $\rho(\tau)$ where $\tau \in \mathcal{T}_\Delta$ sits at the upper left hand corner of Figure 10. Recall that the definition of ρ is dependent on a specific choice of good diagonal not only for the original disk Δ but for every disk reached in the process. For δ as indicated, $\tau \in \mathcal{R}_\Delta$. Take $\tau' = \partial(\tau) \in \mathcal{T}_{\Delta'}$ and a good diagonal δ' of Δ'. Again, $\tau' \in \mathcal{R}_{\Delta'}$ so we must go to Δ'' where $\tau'' = \partial(\tau') \in \mathcal{D}_{\Delta''}$. We construct $\rho''(\tau'') \in \mathcal{T}_{\Delta''}$ (vertical arrow) and bring it back to obtain $\rho'(\tau') = \partial^{-1}(\rho''(\tau'')) \in \mathcal{R}_{\Delta'}$ and finally $\rho(\tau) = \partial^{-1}(\rho'(\tau')) \in \mathcal{R}_\Delta$.

In Figure 11, a loner is identified by a sequence of cut-and-paste operations leading to a disk with a unique tiling. In Figure 12 we again compute $\rho(\tau)$ (τ sits on the upper left corner); notice that there is a large region where domino position is forced but the construction still applies.

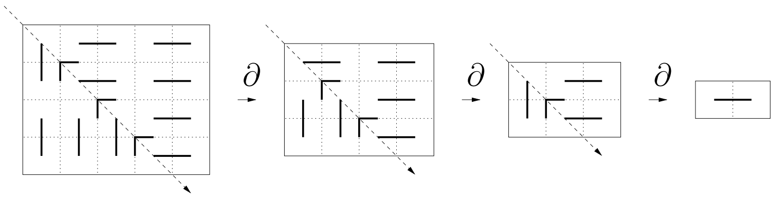

Figure 11: A loner

We recall some well known constructions. The superposition $[\tau_1 - \tau_2]$ of two tilings τ_1 and τ_2 consists of disjoint non-oriented simple closed curves of consecutive dominos (or edges) alternating between τ_1 and τ_2; dominos which are common to τ_1 and τ_2 are discarded. Such curves are cycles (in a different sense) in the dual graph $\mathcal{G}_\Delta$ but we reserve the word for permutation cycles. Consider the bijections $\pi_1, \pi_2 : \{1, 2, \ldots, w\} \to \{1, 2, \ldots, b\}$ associated with the tilings

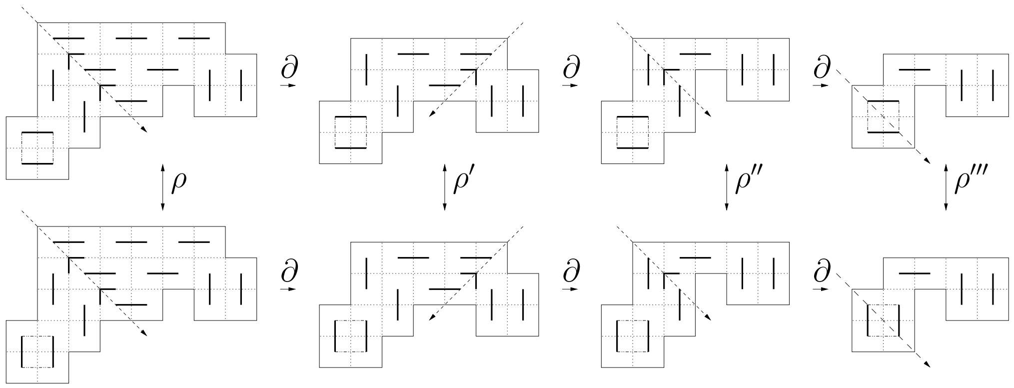

Figure 12: Matching tilings in a more degenerate situation

τ_1, τ_2 and decompose the permutation $\pi_2^{-1}\pi_1 \in S_w$ as a product of disjoint cycles. These cycles correspond to the curves in $[\tau_1 - \tau_2]$ and the length of each curve (defined as the number of edges in $\mathcal{G}_\Delta$) is twice the length of the cycle. The discarded dominos correspond to trivial cycles of length 1 and are irrelevant for parity checks.

If τ_1 and τ_2 differ by a flip then $[\tau_1 - \tau_2]$ is a single curve of length 4 and $\pi_2^{-1}\pi_1$ is a cycle of length 2. More generally, two tilings $\tau_1, \tau_2 \in \mathcal{T}_\Delta$ are *compatible* if $[\tau_1 - \tau_2]$ consists of a single curve whose length is a multiple of 4; we denote compatibility by $\tau_1 \leftrightarrow \tau_2$. If $\tau_1 \leftrightarrow \tau_2$ then $\pi_2^{-1}\pi_1$ is a cycle of even length, an odd permutation, and τ_1 and τ_2 have opposite parities. We claim that, for $\tau_1, \tau_2 \in \mathcal{R}_\Delta$,

$$\tau_1 \leftrightarrow \tau_2 \qquad \Longleftrightarrow \qquad \partial(\tau_1) \leftrightarrow \partial(\tau_2).$$

By the inductive construction of ρ, the claim implies that $\tau_1 \leftrightarrow \rho(\tau_1)$, completing the parity check.

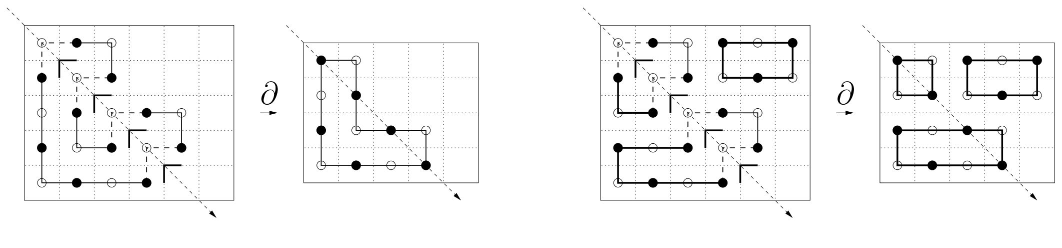

Figure 13: Compatibility is preserved by cut-and-paste

Figure 13 provides two examples of $[\tau_1 - \tau_2]$ and $[\partial(\tau_1) - \partial(\tau_2)]$ for tilings $\tau_i \in \mathcal{R}_\Delta$. The reader should check that in the first example, $\tau_1 \leftrightarrow \tau_2$ and $\partial(\tau_1) \leftrightarrow \partial(\tau_2)$; in the second, $\tau_1 \not\leftrightarrow \tau_2$ and $\partial(\tau_1) \not\leftrightarrow \partial(\tau_2)$. Some vertices of the dual graphs $\mathcal{G}_\Delta$ and $\mathcal{G}_{\Delta'}$ are indicated for clarity.

In general, decompose the curves forming $[\tau_1 - \tau_2]$ into dashed segments through *corridors* between wedges and solid segments on each side of the good diagonal. Cut-and-paste deletes dashed segments and acts on solid segments by translation. Thus, following solid segments yields a natural one-to-one correspondence between curves in $[\tau_1 - \tau_2]$ and curves in $[\partial(\tau_1) - \partial(\tau_2)]$. Furthermore, corresponding curves differ by the deletion of dashed segments of length 2, the passages of the curve through corridors. Since at each such passage the curve goes from one side of the diagonal to the other, the number of passages for each curve is even. Thus, lengths of corresponding curves are congruent $\bmod 4$, proving the claim and completing the proof.

5 Algebraic cut-and-paste

The bulk of this section is dedicated to relating the black-to-white matrices B_Δ and $B_{\Delta'}$ where Δ' is obtained from Δ by cut-and-paste (there is no difficulty in defining black-to-white matrices for bicolored disjoint union of quadriculated disks). More precisely, assume that Δ (resp. Δ') has b (resp. b') black squares and w (resp. w') white squares. Let I_n be the $n \times n$ identity matrix and $I_{n,m}$ be the $n \times m$ defective identity matrix with (i, j) entry equal to 1 if $i = j$ and 0 otherwise. We obtain in Lemma 5.2 a factorization

$$B_\Delta = L_\Delta \begin{pmatrix} I_{b-b',w-w'} & 0 \\ 0 & B_{\Delta'} \end{pmatrix} U_\Delta$$

where L_Δ and U_Δ are very special square triangular matrices. This factorization is the inductive step in the proof of Theorem 2. We first present an example.

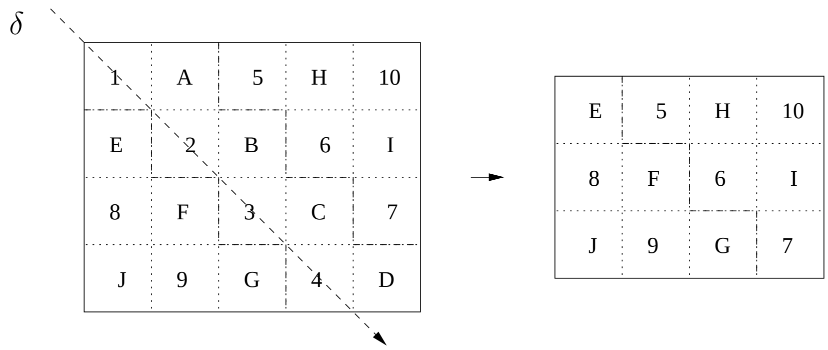

Figure 14: Disks Δ and Δ'

The quadriculated disks shown in Figure 14 have black-to-white matrices

$$B_\Delta = \begin{pmatrix} 1 & & & & 1 & & & & & \\ 1 & 1 & & & 1 & 1 & & & & \\ & 1 & 1 & & & 1 & 1 & & & \\ & & 1 & 1 & & & 1 & & & \\ 1 & 1 & & & & & & 1 & & \\ & 1 & 1 & & & & & 1 & 1 & \\ & & 1 & 1 & & & & & 1 & \\ & & & & 1 & 1 & & & & 1 \\ & & & & & 1 & 1 & & & 1 \\ & & & & & & & 1 & 1 & \end{pmatrix}, \qquad B_{\Delta'} = \begin{pmatrix} 1 & 1 & & 1 & & \\ & 1 & 1 & 1 & 1 & \\ & & 1 & & 1 & \\ 1 & 1 & & & & 1 \\ & 1 & 1 & & & 1 \\ & & & 1 & 1 & \end{pmatrix}.$$

Rows and columns are indexed by numbers and letters respectively in Figure 14. The first four rows and columns of B_Δ correspond to the eight squares removed by cut-and-paste. Partition B_Δ in four blocks so that $B_{11} = B_\delta$ is the black-to-white matrix of the disk around the diagonal δ consisting of squares $1, 2, 3, 4, A, B, C, D$ and B_{22} is the bottom 6×6 principal minor. Notice that B_{22} and $B_{\Delta'}$ are very similar: the difference lies in the top 3×3 principal minor of each matrix. These positions describe adjacencies between squares $5, 6, 7$ and E, F, G.

Elementary operations in rows and columns specified by

$$\tilde{X} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \qquad \tilde{Y} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

can be applied to B_Δ to obtain a block diagonal matrix

$$B_\Delta = \begin{pmatrix} I_4 & 0 \\ \tilde{X} & I_6 \end{pmatrix} \begin{pmatrix} B_\delta & 0 \\ 0 & \tilde{B}_{\Delta'} \end{pmatrix} \begin{pmatrix} I_4 & \tilde{Y} \\ 0 & I_6 \end{pmatrix},$$

where

$$\tilde{B}_{\Delta'} = \begin{pmatrix} -1 & -1 & & 1 & & \\ & -1 & -1 & 1 & 1 & \\ & & -1 & & 1 & \\ 1 & 1 & & & & 1 \\ & 1 & 1 & & & 1 \\ & & & 1 & 1 & \end{pmatrix}$$

is surprisingly similar to $B_{\Delta'}$. More precisely, $\tilde{B}_{\Delta'} = S_{b'} B_\Delta S_{w'}$ where $S_{b'} = \operatorname{diag}(-1, -1, -1, 1, 1, -1)$ and $S_{w'} = \operatorname{diag}(1, 1, 1, -1, -1, 1)$. It is this "coincidence" that allows for this construction to be used as the inductive step in the proof of Theorem 2.

Before discussing the relationship between B_Δ and $B_{\Delta'}$ we present a lemma in linear algebra. The proof is a straightforward computation left to the reader.

Lemma 5.1 *Decompose an $(n+m)\times(n'+m')$ matrix M as*

$$M = \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix},$$

where M_{11} is $n \times n'$. If $n' \le n$ and N is a $n' \times m'$ matrix with $M_{11}N = M_{12}$ then

$$M = \begin{pmatrix} M_{11}I_{n',n} & 0 \\ M_{21}I_{n',n} & I_m \end{pmatrix} \begin{pmatrix} I_{n,n'} & 0 \\ 0 & M_{22}-M_{21}N \end{pmatrix} \begin{pmatrix} I_{n'} & N \\ 0 & I_{m'} \end{pmatrix}.$$

Similarly, if $n' \ge n$ and N is a $m \times n$ matrix with $NM_{11} = M_{21}$ then

$$M = \begin{pmatrix} I_n & 0 \\ N & I_m \end{pmatrix} \begin{pmatrix} I_{n,n'} & 0 \\ 0 & M_{22}-NM_{12} \end{pmatrix} \begin{pmatrix} I_{n',n}M_{11} & I_{n',n}M_{12} \\ 0 & I_{m'} \end{pmatrix}.$$

The next lemma is the inductive step in the proof of Theorem 2.

Lemma 5.2 *Let Δ be a quadriculated disk with b black and w white squares, $b+w > 1$. Let $\Delta' = \Delta'_1 \sqcup \cdots \sqcup \Delta'_d$ (with $b' = b'_1 + \cdots + b'_d$ black and $w' = w'_1 + \cdots + w'_d$ white squares) be obtained from Δ by cut-and-paste along a good diagonal δ. Label black and white squares in Δ so that removed squares come first, in the order prescribed by the good diagonal; label squares in Δ' next. Then the black-to-white matrices B_Δ and $B_{\Delta'}$ satisfy*

$$B_\Delta = \begin{pmatrix} L & 0 \\ X & S_{b'} \end{pmatrix} \begin{pmatrix} I_{b-b',w-w'} & 0 \\ 0 & B_{\Delta'} \end{pmatrix} \begin{pmatrix} U & Y \\ 0 & S_{w'} \end{pmatrix}$$

where L (resp. U) is an invertible lower (resp. upper) square matrix of order $b-b'$ (resp. $w-w'$) and $S_{b'}$ and $S_{w'}$ are square diagonal matrices. Furthermore, all entries of $S_{b'}$, $S_{w'}$, L, U, X and Y equal 0, 1 *or* -1.

The statement above requires clarification in some degenerate cases. If Δ' is empty, $B_{\Delta'}$ collapses and $B_\Delta = LI_{b,w}U$. If instead Δ' is a disjoint union of unit squares, all of the same color, then either $w' = 0$ or $b' = 0$ and

$$B_\Delta = \begin{pmatrix} L & 0 \\ X & S_{b'} \end{pmatrix} \begin{pmatrix} I_{b-b',w} \\ 0 \end{pmatrix} U \quad \text{or} \quad B_\Delta = L \begin{pmatrix} I_{b,w-w'} & 0 \end{pmatrix} \begin{pmatrix} U & Y \\ 0 & S_{w'} \end{pmatrix}.$$

Proof: Assume that the deleted squares are $s_{1/2}, s^l_1, \ldots$ and that the square $s_{1/2}$ is black; thus $k = b-b'$; if $s_{1/2}$ were white all computations would be transposed. Let $j_1, \ldots, j_{k-1}$ be the indices of the white squares $s^r_1, \ldots, s^r_{k-1}$; notice that $j_i > w-w'$. Decompose the matrix B_Δ in four blocks,

$$B_\Delta = \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix},$$

where B_{22} is a $b' \times w'$ matrix. By construction, B_{11} has one of the two forms below, the first case corresponding to balanced good diagonals (i.e., to $b - b' = w - w'$).

$$B_{11} = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ 1 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & 0 & \cdots & 1 & 1 \end{pmatrix}, \quad \text{or} \quad B_{11} = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 1 & 1 & 0 & \cdots & 0 \\ 0 & 1 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}.$$

Let S_b (resp. S_w) be a $b \times b$ (resp. $w \times w$) diagonal matrix with diagonal entries equal to 1 or -1; the i-th entry of S_b (resp. S_w) is -1 if the i-th black (resp. white) square is strictly to the right of δ. Write

$$S_b = \begin{pmatrix} I_{b-b'} & 0 \\ 0 & S_{b'} \end{pmatrix}, \qquad S_w = \begin{pmatrix} I_{w-w'} & 0 \\ 0 & S_{w'} \end{pmatrix}.$$

We have

$$S_b B_\Delta S_w = \begin{pmatrix} B_{11} & -B_{12} \\ B_{21} & B_{22} \end{pmatrix}.$$

The nonzero entries of B_{12} are (i, j_i) and $(i+1, j_i)$ for $i = 1, \ldots, k-1$. Thus, the nonzero columns of B_{12} equal to the first $k-1$ columns of B_{11}. Let N be the $(w-w') \times w'$ matrix with entries 0 or -1, with nonzero entries at $(1, j_1), (2, j_2), \ldots, (k-1, j_{k-1})$. Clearly $B_{11}N = -B_{12}$ and we may apply Lemma 5.1 to write

$$S_b B_\Delta S_w = \begin{pmatrix} B_{11} I_{w-w',b-b'} & 0 \\ B_{21} I_{w-w',b-b'} & I_{b'} \end{pmatrix} \begin{pmatrix} I_{b-b',w-w'} & 0 \\ 0 & B_{22} - B_{21}N \end{pmatrix} \begin{pmatrix} I_{w-w'} & N \\ 0 & I_{w'} \end{pmatrix}.$$

We claim that $B_{\Delta'} = B_{22} - B_{21}N$. The nonzero columns of the matrix $-B_{21}N$ are the columns of B_{21}, except that the first column is moved to position j_1, the second column is moved to j_2 and so on. The k-th column of B_{21}, if it exists, is discarded. These nonzero entries correspond precisely to the identifications which must be performed in order to obtain Δ', i.e., to the ones which must be added to B_{22} in order to obtain $B_{\Delta'}$. Clearing up signs,

$$B_\Delta = \begin{pmatrix} B_{11} I_{w-w',b-b'} & 0 \\ S_{b'} B_{21} I_{w-w',b-b'} & S_{b'} \end{pmatrix} \begin{pmatrix} I_{b-b',w-w'} & 0 \\ 0 & B_{\Delta'} \end{pmatrix} \begin{pmatrix} I_{w-w'} & N S_{w'} \\ 0 & S_{w'} \end{pmatrix}.$$

If the good diagonal is balanced, this finishes the proof. In the unbalanced case, $\tilde{L} = B_{11} I_{w-w',b-b'}$ is not invertible since its last column is zero. Replace the (k,k) entry of $\tilde{L}$ by 1 to obtain a new matrix L: L is clearly invertible and $\tilde{L} I_{b-b',w-w'} = L I_{b-b',w-w'}$. The proof is now complete. ∎

Proof of Theorem 2: The basis of the induction on the number of squares of Δ consists of checking that the theorem holds for disks with at most two squares. Notice that if the disk consists of a single square then $b = 0$ or $w = 0$ and the matrices are degenerate.

Let Δ be a quadriculated disk and $\Delta' = \Delta'_1 \sqcup \cdots \sqcup \Delta'_d$ be obtained from Δ by cut-and-paste. By induction on the number of squares the theorem may be assumed to hold for each Δ'_k and we therefore write $B_{\Delta'} = L_{\Delta'}\tilde{D}_{\Delta'}U_{\Delta'}$. From the induction step, Lemma 5.2, write

$$\begin{aligned} B_\Delta &= \begin{pmatrix} L_{\text{step}} & 0 \\ X_{\text{step}} & S_{b'} \end{pmatrix} \begin{pmatrix} I_{b-b',w-w'} & 0 \\ 0 & B_{\Delta'} \end{pmatrix} \begin{pmatrix} U_{\text{step}} & Y_{\text{step}} \\ 0 & S_{w'} \end{pmatrix} \\ &= L_\Delta \tilde{D}_\Delta U_\Delta. \end{aligned}$$

where

$$L_\Delta = \begin{pmatrix} L_{\text{step}} & 0 \\ X_{\text{step}} & S_{b'} \end{pmatrix} \begin{pmatrix} I_{b-b'} & 0 \\ 0 & L_{\Delta'} \end{pmatrix}, \quad U_\Delta = \begin{pmatrix} I_{w-w'} & 0 \\ 0 & U_{\Delta'} \end{pmatrix} \begin{pmatrix} U_{\text{step}} & Y_{\text{step}} \\ 0 & S_{w'} \end{pmatrix}.$$

The theorem now follows from observing that each nonzero entry of L_Δ (resp. U_Δ) is, up to sign, copied from either $L_{\Delta'}$, L_{step} or X_{step} (resp. $U_{\Delta'}$, U_{step} or Y_{step}) and is therefore equal to 1 or -1. ■

We present a direct consequence of Theorem 2.

Corollary 5.3 *Let Δ be a quadriculated disk with black-to-white matrix B_Δ. If v has integer entries and the system $B_\Delta x = v$ admits a rational solution then the system admits an integer solution.*

This corollary may be interpreted as saying that the co-kernel $\mathbb{Z}^b/B_\Delta(\mathbb{Z}^w)$ of $B_\Delta : \mathbb{Z}^w \to \mathbb{Z}^b$ is a free abelian group. From Theorem 2, the rank r of B_Δ is the same in $\mathbb{Q}$ as in $\mathbb{Z}_p$ for any prime number p. Notice that the proof of Theorem 1 in [1] is based on this fact for $p = 2$.

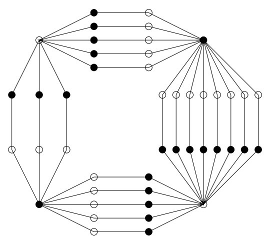

Figure 15: Determinant 1 does not imply $L\tilde{D}U$ decomposition

The example in Figure 15 is instructive: the B_G matrix of this planar graph G has determinant 1 but admits no $L\tilde{D}U$ decomposition where the matrices have integer coefficients since the removal of any two vertices of opposite colors from G yields a graph whose determinant has absolute value greater than 1.

6 Boards

Topological subdisks of $\mathbb{R}^2$ consisting of unit squares with vertices in $\mathbb{Z}^2$ are *boards*. In other words, a board is a topological subdisk of $\mathbb{R}^2$ whose boundary is a polygonal curve consisting of segments of length 1 joining points in $\mathbb{Z}^2$. The quadriculated disk in Figure 1 is not a board. The class of boards is not closed under cut-and-paste: in Figure 16, the two enhanced segments on the boundary would be superimposed by cut-and-paste along the good diagonal on the left. Cut-and-paste along the good diagonal indicated on the right, however, yields a smaller board. The main result of this section is that, given a board Δ, it is always possible to choose a good diagonal δ such that $\Delta' = \partial_\delta(\Delta)$ is a disjoint union of boards.

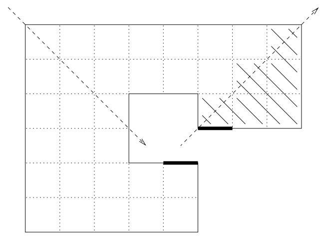

Figure 16: A board and two good diagonals, one excellent.

Orient the boundary of a board Δ counterclockwise, so that Δ lies to the left of the boundary. Consider boundary vertices which are local extrema for the restriction of $x+y$ to the boundary: as in Figure 17, call such vertices *positive* if they are corners (equivalently, if they are local extrema for the restriction of $x+y$ to Δ) and *negative* otherwise. Let $V_{B,+}$ (resp. $V_{B,-}$) be the number of positive (resp. negative) boundary vertices.

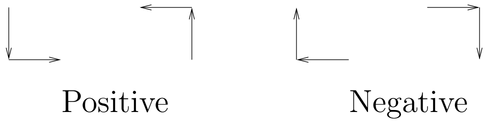

Figure 17: Positive and negative boundary vertices

Lemma 6.1 $V_{B,+} - V_{B,-} = 2$.

Proof: Define F, E, E_I, E_B and V_I as in Lemma 2.2. The number of boundary vertices is $V_B = V_{B,+} + V_{B,-} + V_{B,0}$ where $V_{B,0}$ is the number of boundary vertices which are neither positive nor negative. For each square, consider its NW and SE vertices: interior vertices and negative vertices are counted twice, positive vertices are not counted and other boundary vertices are counted once and therefore $2F = 2V_I + 2V_{B,-} + V_{B,0} = 2V_I + V_B - (V_{B,+} - V_{B,-})$. Recall that $4F = 2E - E_B$ (Lemma 2.2) and $E = V + F - 1$ (Euler) and therefore $2F = 2V_I + V_B - 2$, completing the proof. ∎

Theorem 3 *It is always possible to cut-and-paste a given board Δ to obtain a disjoint union of boards Δ'.*

Proof: A diagonal is *excellent* if the x and y coordinates are both monotonic along one of the two boundary arcs between v_0 and v_k; without loss, let this arc lie to the right of the diagonal. Excellent diagonals are good: the vertex $v^r_{k-1/2}$ is on the boundary. We interpret cut-and-paste along an excellent diagonal as leaving Δ^l fixed and moving Δ^r. In this way, Δ' becomes a subset of Δ and is therefore a disjoint union of boards. We are left with proving that any board admits excellent diagonals. Each diagonal defines two boundary arcs: order these arcs by inclusion. We claim that a diagonal defining a minimal arc is excellent.

Let $\delta^m = (v^m_0 v^m_1 \ldots v^m_k)$ be a diagonal inducing a minimal arc α: assume without loss of generality that $v^m_i = (a + i, b + i)$ for integers a and b. Consider the set $\tilde{\Delta}$ (dashed in Figure 16) consisting of the squares totally or partially surrounded by α and δ^m. It is easy to verify that $\tilde{\Delta}$ is a legitimate board with boundary consisting of α and ζ, where ζ is the zig-zag line next to δ^m. Thus, the last edge of ζ can not overlap with α without contradicting the fact that the boundary point v^m_k of Δ is surrounded by at most three squares in Δ.

By Lemma 6.1, the board $\tilde{\Delta}$ has at least two positive boundary points. We claim that the existence of a positive boundary point distinct from v^m_0 and v^m_k contradicts minimality. Notice that at this point it is clear that v^m_0 is positive; the status of v^m_k as a positive boundary point will only follow from the claim. Indeed, such a positive point $\hat{v}$ can not belong to the zig-zag line ζ and must therefore belong to α. Draw a diagonal $\hat{\delta}$ starting at $\hat{v}$: being parallel to ζ, $\hat{\delta}$ must intersect the boundary of $\tilde{\Delta}$ in α and therefore defines a smaller arc $\hat{\alpha}$, contradicting minimality and proving the claim. Again by Lemma 6.1, there are no negative boundary vertices. In particular, there are no positive or negative boundary vertices in α and we are done. ∎

References

[1] Deift, P. A. and Tomei, C., *On the determinant of the adjacency matrix for a planar sublattice*, J. Combin. Theory Ser. B, **35**, 278–289 (1983).

[2] Elkies, N., Kuperberg, G., Larsen, M. and Propp, J., *Alternating sign matrices and domino tilings*, J. Alg. Combin. **1**, 112–132 and 219–239 (1992).

[3] Massey, W. S., *A basic course in algebraic topology*, GTM 127, Springer-Verlag, New York (1991).

[4] Pachter, L., *Combinatorial approaches and conjectures for* 2*-divisibility problems concerning domino tilings of polyominoes*, The Eletronic Journal of Combinatorics, 4, # R29 (1997).

[5] Saldanha, N. C. and Tomei, C., *Tilings of quadriculated annuli*, Journal of Combinatorial Theory, Series B 88, 153-183 (2003).

[6] Saldanha, N. C., Tomei, C., Casarin, M. A. and Romualdo D., *Spaces of domino tilings*, Discrete Comput. Geom. **14**, 207–233 (1995).

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

Nicolau C. Saldanha and Carlos Tomei
Departamento de Matemática, PUC-Rio
R. Marquês de S. Vicente 225, Rio de Janeiro, RJ 22453-900, Brazil

nicolau@mat.puc-rio.br; http://www.mat.puc-rio.br/∼nicolau/
tomei@mat.puc-rio.br