A Note on Domino Shuffling

É. Janvresse, T. de la Rue, Y. Velenik
Laboratoire de Mathématiques Raphaël Salem
UMR 6085 CNRS – Université de Rouen
76801 Saint Etienne du Rouvray, France
Elise.Janvresse@univ-rouen.fr,
Thierry.de-la-Rue@univ-rouen.fr,
Yvan.Velenik@univ-rouen.fr

Submitted: Aug 22, 2005; Accepted: Mar 27, 2006; Published: Apr 4, 2006
Mathematics Subject Classifications: 05B45, 05C70, 52C20, 68R10

Abstract

We present a variation of James Propp's generalized domino shuffling, which provides an efficient way to obtain perfect matchings of weighted Aztec diamonds. Our modification is specially tailored to deal with cases when some of the weights are zero. This allows us to tile efficiently a large class of planar graphs, by embedding them in a large enough Aztec diamond. We also give a sufficient condition on the size of the latter diamond for the algorithm to succeed.

1 Introduction

The domino-shuffling algorithm was originally introduced in [3] for counting the perfect matchings (henceforth, we only write matching) of the Aztec diamond of order n (see Figure 1). It was soon realized that this algorithm could also be used to generate a uniform sampling of such matchings (see [4]).

In [7], James Propp introduced a generalized version of the shuffling algorithm allowing to efficiently generate random matchings of Aztec diamonds with weighted edges. However his method does not always apply when some of the weights are set to zero. Actually, Propp mentions a way to use his algorithm to compute various quantities when such a problem occurs (by setting the corresponding weights to $\varepsilon > 0$ and keeping only the terms of lowest order in $\varepsilon > 0$), but although this might be suitable when performing some analytical computations, it is not clear how it should be efficiently implemented on a computer. To be able to treat the case of zero-weights is important, since this allows one to forbid some edges, and consequently to obtain matchings of a great variety of planar graphs, not only Aztec diamonds.

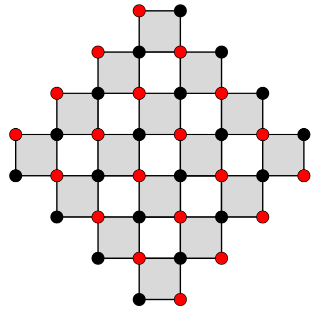

Figure 1: The Aztec diamond of order 4.

We propose in this paper a slight variant of Propp's generalized shuffling algorithm, which is adapted to situations where edges of zero-weight are present. It allows one to generate random matchings of a large class of weighted subgraphs of the square lattice, by embedding them in a large enough Aztec diamond. We also provide (reasonably good) bounds on the minimal size of the Aztec diamond in which the graph should be embedded, which shows that the computational cost for constructing matchings of these graphs is of the same order as for the original generalized shuffling.

Other efficient algorithms applying to various subgraphs of the square lattice exist. The most efficient one uses Wilson's algorithm to generate random spanning trees [8] and a mapping from spanning trees to matchings originally introduced by Temperley, and extended to weighted graphs in [5]. However the latter mapping imposes some constraints on the graph and does not seem to apply to graphs such as the semi-regular 4-6-12 tesselation. This case and many others can be treated using a combination of generalized shuffling and urban renewal, see [6] for examples of implementations.

1.1 Definitions and notations

The Aztec Diamond of order $n \geq 1$ is a planar graph, which can be seen as a subset of the square lattice; its set of vertices is defined as follows

$$A_n = \left\{ (x, y) \in \mathbb{Z}^2 : \left|x - \frac{1}{2}\right| + \left|y - \frac{1}{2}\right| \leq n \right\},$$

and there is one edge connecting each pair of vertices at Euclidean distance 1.

A (perfect) matching of a graph is a subset of edges of the graph such that each vertex of the graph belongs to exactly one edge of the subset.

Associating to each edge e of a graph G a weight $w(e)$, we can define a probability measure on the set of all matchings of G, by setting the probability of a given matching π to be proportional to the product of the weights of the edges in π.

To each perfect matching of a graph corresponds a unique tiling of its dual. In the sequel, we will sometimes use the dual terminology even when talking about graphs; for example, we will say that two vertices are covered by a domino when they are matched, and say that a graph is tileable when it admits a perfect matching.

1.2 Roadmap for the paper

The algorithm is described in Section 2, while the proof that it actually does what is claimed is relegated to Section 3. A lower bound on the size of the Aztec diamond in which the graph should be embedded is given in Section 4, as well as some examples of planar graphs that can be tiled with this algorithm (including the 6-6-6, 4-8-8 and 4-6-12 (semi-)regular tesselations).

2 Description of the shuffling algorithm

Let G be a finite weighted subgraph of the square lattice, with non-negative edge-weights $w(e)$. We make the following assumption:

Assumption 1. *There exists a large Aztec diamond A_n containing G, such that the complement of G in A_n is tileable.*

Sufficient conditions on G for Assumption 1 to hold will be discussed in Section 4. To generate a matching of G, it is sufficient to generate a matching of A_n in which all edges connecting G to its complement are forbidden. We can therefore restrict our attention to matchings of A_n with weighted edges, some of the weights being possibly zero. The procedure we are going to describe is a slight variant of Propp's generalized shuffling algorithm. It consists in recursively computing weights $w_m(e)$ of edges in each Aztec diamond A_m for $n \geq m \geq 1$, then generating matchings of these weighted graphs, from A_1 up to A_n. This procedure will halt prematurely if and only if the graph is not tileable.

Let us start by defining the weights $w_n(e)$ for each edge e of A_n.

$$w_n(e) := \begin{cases} w(e) & \text{if } e \text{ is an edge in } G \\ 0 & \text{if } e \text{ connects } G \text{ to its complement in } A_n \\ 1 & \text{otherwise.} \end{cases} \tag{1}$$

We now define recursively the rules to compute the weights $w_{m-1}(e)$ of A_{m-1} from the weights $w_m(e)$ of A_m. This will be achieved in two steps. Let us suppose that the faces of the square lattice are colored as a chessboard. Notice that all faces of the inner boundary of A_m are of the same color. During the computation of w_{m-1}, we will call *active* faces the faces of A_m of this color. Each edge e is on the boundary of a unique active face $C = C(e)$. We write $e' = e'(e)$ the opposite edge (i.e., the edge on the boundary of the same active face as e, which is not adjacent to e). To each active face C, we associate a number $\mathrm{DP}_m(C)$, the *double product of* C, defined as follows: $\mathrm{DP}_m(C) = w_m(e_1)w_m(e_1') + w_m(e_2)w_m(e_2')$, where e_1 and e_2 are two adjacent edges on the boundary of C.

Step 1: Pre-computation of the weights w_{m-1}

The first step is described in Figure 2[1], where we denote by α, β, γ, δ the weights w_m of the edges of the boundary of an active face of A_m, and indicate the weights w_{m-1} of the same edges in A_{m-1}. We consider three different cases to deal with edges with zero weight, which cover all possibilities up to obvious rotations.

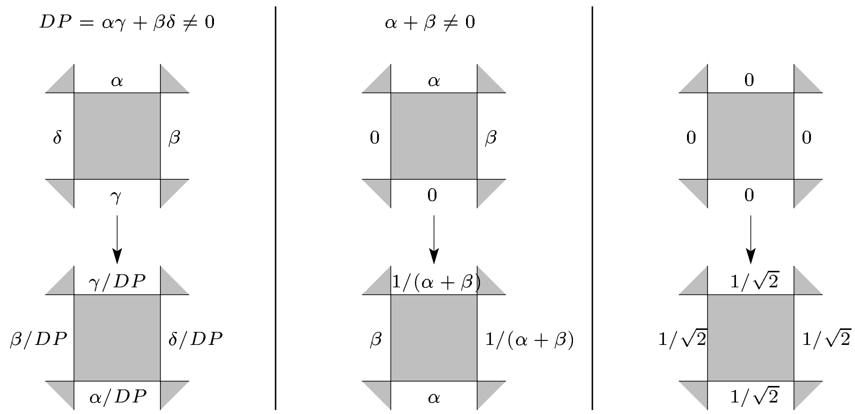

Figure 2: Pre-computation of the weights w_{m-1} from the weights w_m.

Step 2: Edge-erasing procedure

In this step, we apply a special treatment each time we find two adjacent edges of an active cell with zero w_m-weight.

Stop condition 1. *If we find two adjacent edges of an active cell with zero w_m-weight lying on the boundary of A_m, the procedure stops: The graph G is not tileable.*

If these two edges do not lie on the boundary of A_m, we denote by V the vertex where these two edges meet, and replace by 0 the computed weights w_{m-1} of all other edges of A_{m-1} incident on V, as indicated on Figure 3.

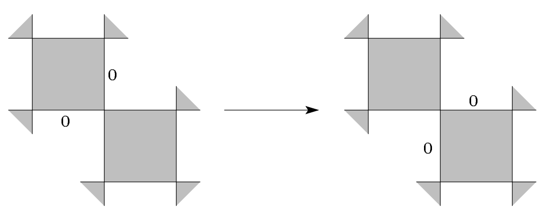

Figure 3: Edge-erasing procedure: On the left part of the figure, we indicate edges whose w_m-weight is zero, and on the right part the edges whose w_{m-1}-weight is set to zero.

[1] In all the pictures, the described rules are invariant by rotation, and the active faces are shaded.

If a face happens to have more than one pair of adjacent edges with zero w_m-weight, the procedure is applied more than one time for that face.

We will now generate random matchings of A_m, $m = 1, \ldots, n$ distributed according to the weights w_m. Assume we have already constructed a matching of A_m; we describe how to obtain a matching of A_{m+1}. The active faces during this step are those colored like the inner boundary of A_{m+1}. There are three possible cases depending on the number of dominoes on the boundary of each active face (see the picture below): If there is a single domino, along edge e, we move it to e'; if there are two dominoes, we remove them; if there is no domino, we place two new dominoes on opposite edges, choosing their orientation with probability depending on the weights w_{m+1}.

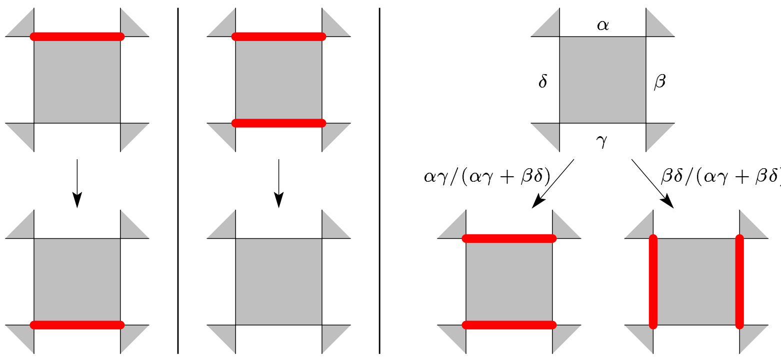

Figure 4: Shuffling: A matching of A_m (top) is mapped to a matching of A_{m+1} (bottom). The mapping is random in the case of an empty face.

To initiate the process, *i.e.* to tile A_1, we simply apply the third rule to choose the orientation of the two dominoes.

Stop condition 2. *If the probabilities used to fill in the empty faces are not well defined, i.e. if the involved double product $\alpha\gamma + \beta\delta$ vanishes, then the procedure stops: The graph G is not tileable.*

Remark 2.1. *It turns out that if the probabilities involved when constructing a matching of A_1 are well defined, then all subsequent probabilities are also well defined and G is tileable. This is a corollary of Lemma 2 below.*

Remark 2.2. *The fact that we may have to consider situations where double products $\mathrm{DP}_m(C)$ vanish during the computation of the weights does not mean that G is not tileable. Indeed, the corresponding active faces may never be empty during the shuffling process.*

Theorem 1. *Under Assumption 1, either G is tileable and the procedure described above generates a random matching of G according to the weights w, or G is not tileable and the procedure stops prematurely.*

3 Proof that it works

3.1 Equivalence classes of matchings

We first define an equivalence relation on all matchings of A_m, when faces of a given color are active. Let π be a matching of A_m, and C be an active face. We say that C is *empty* in π if there is no domino on its boundary; we say that C is *full* in π if there are two dominoes on its boundary; and we say that an edge e is *lonely* in π if it carries the only domino on the boundary of the adjacent active face. Let π_1 and π_2 be two matchings of A_m. We say that they are equivalent when they have the same empty faces, the same full faces and the same lonely edges[2]. We denote by π^* an equivalence class of matchings of A_m, and define its weight as follows:

$$\begin{aligned} w_m(\pi^*) &:= \sum_{\pi \in \pi^*} \prod_{e \in \pi} w_m(e) \\ &= \prod_{e \text{ lonely in } \pi^*} w_m(e) \prod_{C \text{ full in } \pi^*} \mathrm{DP}_m(C). \end{aligned} \tag{2}$$

In some matching π of A_m, we say that a vertex v is *tied* to the active face F if the domino which covers v lies on the boundary of F. Every vertex is tied to exactly one of its two adjacent active faces. We can observe that the equivalence class of a given matching is characterized by stating to which active face is tied each vertex.

3.2 Shuffling on equivalence classes

Let π^* be an equivalence class of matchings of A_m, where we have chosen the active faces to be colored like the inner boundary of A_m. We define $\sigma(\pi^*)$ as the equivalence class of matchings of A_{m-1} obtained using the following rules: Faces which were full in π^* are empty in $\sigma(\pi^*)$; faces which were empty in π^* are full in $\sigma(\pi^*)$; if e was lonely in π^*, e' is lonely in $\sigma(\pi^*)$.

The fact that we get an equivalence class of matchings of A_{m-1} is ensured by the following observation: Each vertex in A_{m-1} is adjacent to exactly two active faces in A_m, and the transformation from π^* to $\sigma(\pi^*)$ simply exchanges the active face to which the vertex is tied (Figure 5), which shows that such a vertex is covered by exactly one domino in $\sigma(\pi^*)$. The vertices in $A_m \setminus A_{m-1}$ were necessarily tied in π^* to an active face belonging to the inner boundary of A_m, and are not covered anymore by a domino in $\sigma(\pi^*)$.

The following lemma is the analog of the general complementation theorem (Theorem 2.3 in [2]).

Lemma 2. *Assuming that we cannot find two adjacent edges of an active cell with zero w_m-weight lying on the boundary of A_m, there exists a constant D_m such that, for any*

[2]We can point out that being equivalent in this sense corresponds to having the same alternating sign pattern as defined in [1]

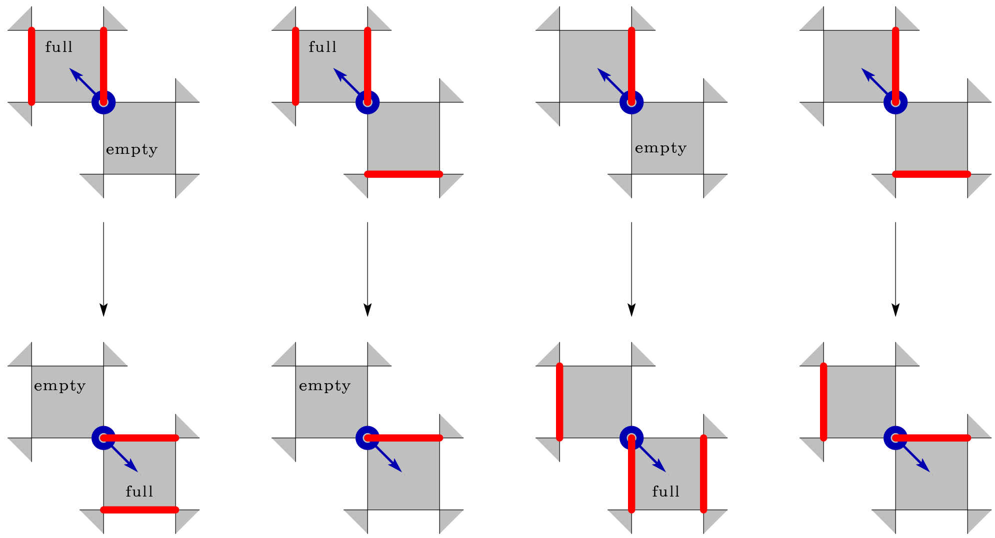

Figure 5: The transformation from π^* to $\sigma(\pi^*)$ simply exchanges the active face to which the vertex is tied.

equivalence class π^* *of matchings in* A_m*, we have*

$$w_m(\pi^*) = D_m\, w_{m-1}(\sigma(\pi^*)). \tag{3}$$

Proof. Let us first consider the case $w_m(\pi^*) \neq 0$. We introduce $\mathrm{DP}^\star_m(C) = \mathrm{DP}_m(C)$ if $\mathrm{DP}_m(C) \neq 0$ and $\mathrm{DP}^\star_m(C) = 1$ otherwise. Then, using (2) and observing that each cell C which is full in π^* satisfies $\mathrm{DP}_m(C) \neq 0$, we get

$$w_m(\pi^*) = \prod_{e \text{ lonely in } \pi^*} \frac{w_m(e)}{\mathrm{DP}^\star_m(C(e))} \prod_{C \text{ empty in } \pi^*} \frac{1}{\mathrm{DP}^\star_m(C)} \prod_C \mathrm{DP}^\star_m(C).$$

Setting $D_m = \prod_C \mathrm{DP}^\star_m(C)$, we can rewrite the above expression as

$$w_m(\pi^*) = D_m \prod_{e \text{ lonely in } \sigma(\pi^*)} w_{m-1}(e) \prod_{C \text{ full in } \sigma(\pi^*)} \mathrm{DP}_{m-1}(C)\,.$$

Indeed, it is enough to check that this identity holds in all the situations described in Section 2, which we proceed to do now.

Observe that we have chosen the weights w_{m-1} in such a way that $\mathrm{DP}_{m-1}(C)$ is always equal to $1/\mathrm{DP}^\star_m(C)$ before the edge-erasing procedure is applied. Now, observe that the edge-erasing procedure has no effect on $\mathrm{DP}_{m-1}(C)$ when C is full in $\sigma(\pi^*)$. Indeed, if the latter procedure results in setting the weight of some edges e_1 to zero, then the corresponding cell $C(e)$ cannot be full in $\sigma(\pi^*)$ under the condition $w_m(\pi^*) \neq 0$ (see Figure 6).

We must now take care of the lonely edges in $\sigma(\pi^*)$. An edge e can be lonely in $\sigma(\pi^*)$ only if e' was lonely in π^*, which can only happen if $w_m(e') \neq 0$; among the

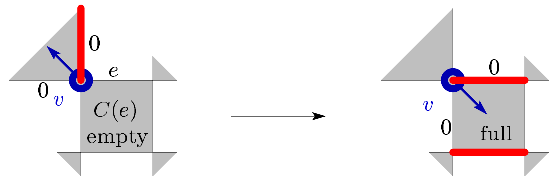

Figure 6: Let e be an edge whose weight has been set to zero by the edge-erasing procedure. If $C(e)$ is full in $\sigma(\pi^*)$, the vertex v is tied to $C(e)$ in $\sigma(\pi^*)$, and therefore is tied to the adjacent active cell in π^*. Thus a zero-w_m-weight edge must be occupied in π^*.

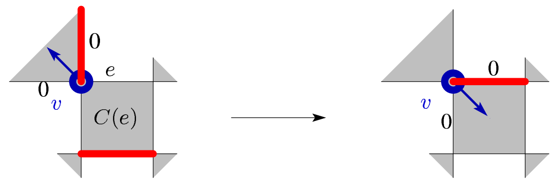

Figure 7: Let e be an edge whose weight has been set to zero by the edge-erasing procedure. If e is lonely in $\sigma(\pi^*)$, the vertex v is tied to $C(e)$ in $\sigma(\pi^*)$, and therefore is tied to the adjacent active cell in π^*. Thus a zero-w_m-weight edge must be occupied in π^*.

three cases depicted in Fig. 2, this only leaves the first two, for which it is obvious that $w_{m-1}(e) = w_m(e')/\mathrm{DP}^\star_m(C(e'))$. Now, the same argument as above proves that the edge-erasing procedure does not affect $w_{m-1}(e)$ when e is lonely in $\sigma(\pi^*)$ (see Figure 7). Notice that in π^* one of the two edges e_1 or e_2 *must* have been covered. Now, either the opposite edge was also covered, and the cell was then full in π^* and therefore empty in $\sigma(\pi^*)$, or e_1 (or e_2) was lonely in π^* and therefore neither e_1 nor e_2 are covered in $\sigma(\pi^*)$. We thus conclude that e_1 and e_2 *must* both be empty in $\sigma(\pi^*)$, and thus setting the weight of these edges to zero does not affect the above expression.

Let us now turn to the case $w_m(\pi^*) = 0$. This can happen only if either there is a full cell C in π^* with $\mathrm{DP}_m(C) = 0$, or a lonely edge e with $w_m(e) = 0$. In the case of a full cell C, Figure 8, left, shows that the resulting weight $w_{m-1}(\sigma(\pi^*))$ is also zero. We can thus turn our attention to the case of a lonely edge e with $w_m(e) = 0$. If $\mathrm{DP}_m(C(e)) \neq 0$, the weight is moved together with the lonely edge across C, and therefore the resulting weight is also zero. If $\mathrm{DP}_m(C(e)) = 0$, we conclude as before (see Figure 8, right).

□

The previous lemma immediately implies the claim made in Remark 2.1.

Lemma 3. *If there exists $m \geq 1$ such that two adjacent edges of an active cell with zero w_m-weight lie on the boundary of A_m, then G is not tileable.*

Proof. We may assume that m is the greatest integer satisfying the assumption of the lemma. Let π be a matching of A_n and π^* its equivalence class. Then, by several applica-

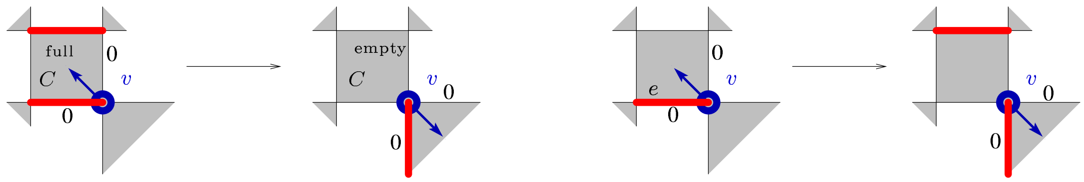

Figure 8: Left: case of a full cell. Right: case of a lonely edge. Notice that the vertex v where the two zero-w_m-weight edges meet belongs to A_{m-1} by virtue of the assumption made in the lemma.

tions of Lemma 2, $w_m(\sigma^{n-m}(\pi^*))$ is proportional to $w_n(\pi^*)$. The hypothesis of the lemma yields $w_m(\sigma^{n-m}(\pi^*)) = 0$, hence $w_n(\pi^*) = 0$. □

Proof of Theorem 1. Let us introduce the notation

$$Z_m := \sum_{\pi \in \Pi(A_m)} w_m(\pi).$$

Lemma 2 shows that if G is not tileable and the procedure has not stopped before constructing the weights w_1 (see Stop condition 1), then $Z_1 \propto Z_n = 0$. This means that A_1 is not tileable, hence $\mathrm{DP}_1(C) = 0$ for the unique cell C of A_1. The procedure stops.

If G is tileable, we prove by induction that for all $m = 1, \ldots, n$, the generated random matching π_m of A_m is distributed according to the weights w_m. This is obvious for $m = 1$. Let's suppose it is true for some $1 \leq m < n$. For any equivalence class π^* of matching of A_{m+1}, Lemma 2 yields

$$\frac{w_{m+1}(\pi^*)}{Z_{m+1}} = \frac{w_m(\sigma(\pi^*))}{Z_m}.$$

By induction hypothesis, the probability that $\pi_m \in \sigma(\pi^*)$ is given by $\frac{w_m(\sigma(\pi^*))}{Z_m}$. But $\pi_m \in \sigma(\pi^*)$ if and only if $\pi_{m+1} \in \pi^*$. Therefore, the probability that $\pi_{m+1} \in \pi^*$ is $\frac{w_{m+1}(\pi^*)}{Z_{m+1}}$.

It remains to check that the probability to get π when $\pi_{m+1} \in \pi^*$ is given by $\frac{w_{m+1}(\pi)}{w_{m+1}(\pi^*)}$. This follows from the fact that the only freedom in selecting π in π^* comes from the choice of the orientation of the edges in full active cells. This is done independently for each of the latter, with probability given in Figure 4. This concludes the proof. □

4 Embedding a graph into an Aztec diamond

4.1 How large should the containing diamond be?

In Section 2, we have proposed an algorithm to tile a weighted subgraph G of the square lattice by embedding it in a sufficiently large Aztec diamond. In this section, we give a sufficient condition on G for this algorithm to be applicable, and provide an explicit

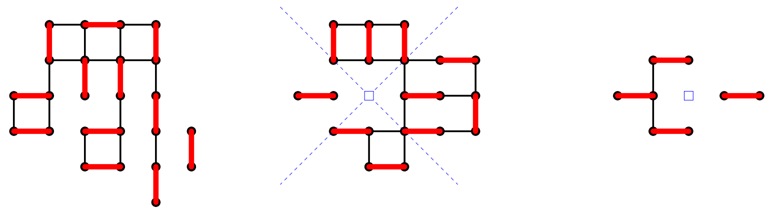

Figure 9: Three tileable graphs G_1, G_2 and G_3. G_1 belongs to $\mathbf{G}$, but neither G_2 (because it contains two diagonal slices with gaps), nor G_3 (because it is not connected) belongs to G.

criterion to determine the size of the Aztec diamond in which G should be embedded. In this part, "vertex" means "vertex in $\mathbb{Z}^2$" (or point $(x, y) \in \mathbb{Z}^2$).

It is convenient to define cardinal points on the plane, in such a way that the vector $(1, 1)$ points towards north. An S-N diagonal of $\mathbb{Z}^2$ is a set of vertices of the form $\{(x, y) \in \mathbb{Z}^2 : x - y = k\}$. Similarly, a W-E diagonal of $\mathbb{Z}^2$ is a set of vertices of the form

$$D^k := \{(x, y) \in \mathbb{Z}^2 : x + y = k\}.$$

Definition 4.1. *We say that a subset $V \subset \mathbb{Z}^2$ is* gapless *if, for any pair of vertices $v_1, v_2 \in V$ lying on the same diagonal, any other vertex v lying on the same diagonal between v_1 and v_2 also belongs to V.*

Definition 4.2. *We say that a subset $V \subset \mathbb{Z}^2$ is* connected *if, for any pair of vertices $v, w \in V$, there exists a path $v_0 = v, v_1, \ldots, v_{\ell-1}, v_\ell = w$ of vertices in V such that $\|v_{k+1} - v_k\| = 1$.*

Let us denote by $\mathbf{G}$ the set of all finite, tileable subgraphs G of the square lattice, the vertices of which form a gapless and connected set, see Figure 9. Observe that except for tileability, there is no condition on the edges of G. We do not even demand that G be a connected graph. Given $G \in \mathbf{G}$, we consider the smallest rectangle containing G with sides parallel to the S-N and W-E-diagonals, and we define by $H_{\text{S-N}}$ (resp. $H_{\text{W-E}}$) the maximum number of vertices of $\mathbb{Z}^2$ on S-N (resp. W-E) diagonal slices of this rectangle, see Figure 10.

Let $b = \min\{k : G \cap D^k \neq \emptyset\}$ and $t = \max\{k : G \cap D^k \neq \emptyset\}$. Let also ℓ (resp. r) be the index of one of the diagonals containing the western-most (resp. eastern-most) vertices of G. (See Figure 10.)

Theorem 4. *Let $G \in \mathbf{G}$. Then for all*

$$n \geq H_{W\text{-}E} + 2\max\left\{\lceil\frac{r-b}{2}\rceil, \lceil\frac{t-r}{2}\rceil\right\} + 2\max\left\{\lceil\frac{\ell-b}{2}\rceil, \lceil\frac{t-\ell}{2}\rceil\right\}, \tag{4}$$

there exists a translate of A_n containing G such that the complement of G in A_n is tileable. In particular, this is true if

$$n \geq H_{W\text{-}E} + 4H_{S\text{-}N}.$$

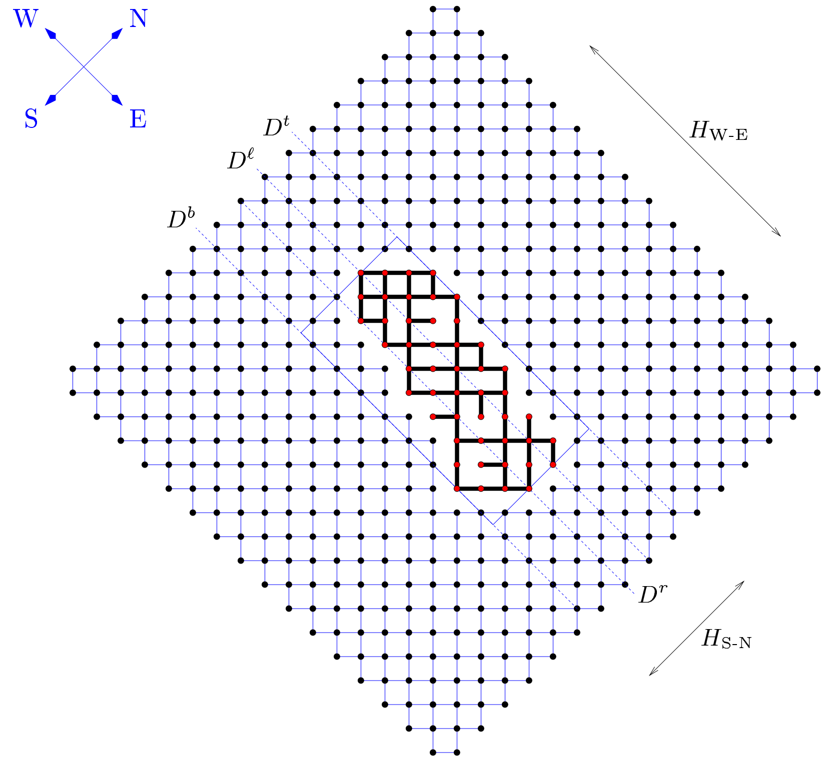

Figure 10: A graph $G \in \mathbf{G}$ (drawn with bold edges), embedded in an Aztec diamond. In this picture, $H_{\text{S-N}} = 5$ and $H_{\text{W-E}} = 9$.

Remark 4.1. *The second and third examples in Figure 9 show that Theorem 4 can fail when G is not connected, or not gapless.*

Proof. Let us start by introducing some terminology (see Figure 11).

Definition 4.3. *Let $n \geq 1$. A* W-E Aztec tube of size n *is a translate of a subset of the form*

$$T_n^{(b,t)} := \{(x, y) \in \mathbb{Z}^2 : b \leq x + y \leq t \text{ and } |x - y| \leq n\},$$

where $b \leq t$ are two integers.

Note that such a tube is the union of W-E diagonal slices

$$D_n^k := \{(x, y) : x + y = k \text{ and } |x - y| \leq n\}, \; b \leq k \leq t.$$

Definition 4.4. *A* quasi-matching *of $T_n^{(b,t)}$ is a collection of edges with endpoints in $T_n^{(b,t)}$, such that all vertices of $T_n^{(b,t)}$, except possibly those on $D_n^b \cup D_n^t$, belong to exactly one edge.*

Given a finite subset of vertices $V \subset D^k$ and a set of edges E, we denote by $|V|$ the cardinality of V, and by $\#_N(V)$ (resp. $\#_S(V)$) the number of edges in E with one endpoint in V and the other endpoint in D^{k+1} (resp. D^{k-1}).

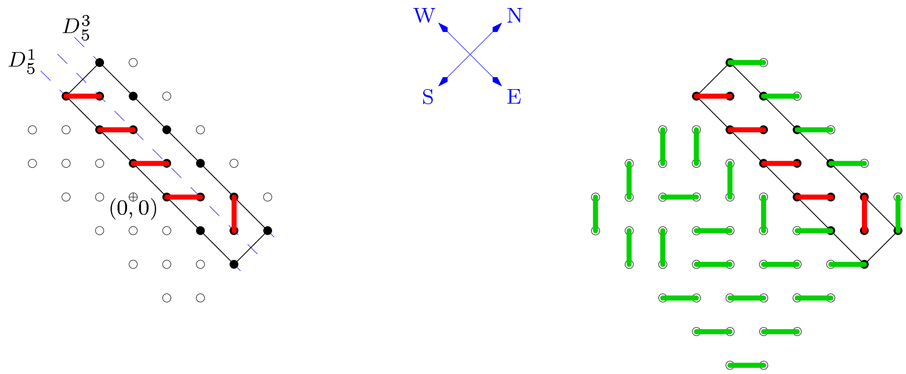

Figure 11: Left: the W-E Aztec tube of size 5 with $b = 1$ and $t = 3$, and a quasi-matching. Right: extension of this quasi-matching to a matching of an Aztec diamond of size 5.

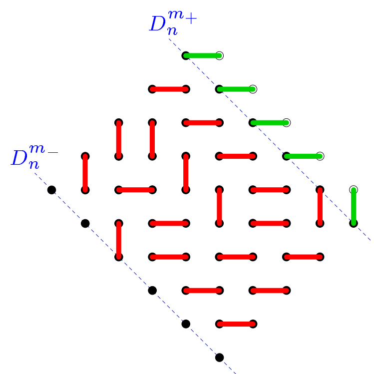

Figure 12: Extension (northward) of a quasi-matching of $T_n^{(m_-,m_+)}$. Note that the resulting quasi-matching of $T_n^{(m_-,m_++1)}$ cannot be extended northward anymore.

Lemma 5. *Any quasi-matching of $T_n^{(b,t)}$ can be extended to a matching of some Aztec diamond A_n of size n containing $T_n^{(b,t)}$.*

Proof of Lemma 5. We are going to extend the quasi-matching of $T_n^{(b,t)}$ to larger and larger tubes of size n as long as we can, then check that what we obtain is a matching of an Aztec diamond of size n.

Notice that the W-E diagonals of an Aztec diamond of size n contain alternatively n and $n+1$ vertices, and that the same is true for W-E Aztec tubes of size n. We call the former *short diagonals* and the latter *long diagonals*. We extend this quasi-matching using the following procedure, see Figure 12.

Step 0. Let $m_+ = t$ and $m_- = b$.
Step 1. If all the vertices of $D_n^{m_+}$ are already covered in the quasi-matching, go to Step 2; observe that this can only happen when $D_n^{m_+}$ is short, since each vertex has to be linked with a vertex on $D_n^{m_+-1}$. If $D_n^{m_+}$ contains at least one uncovered vertex, we claim that it is always possible to extend it one step further. Indeed, there are at most n edges to add,

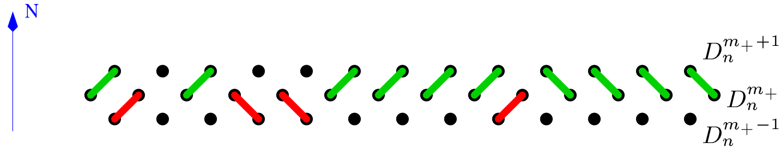

Figure 13: The uncovered vertices of $D_n^{m_+}$ are naturally partitionned by the covered ones. For each group, we choose all dominoes from $D_n^{m_+}$ to $D_n^{m_++1}$ with the same orientation. The only constraints we have are at the boundary of $D_n^{m_+}$ when it is a long diagonal; in that case the orientation is forced. If we have at least one covered vertex in $D_n^{m_+}$, then these constraints do not matter.

and there are at least n vertices on $D_n^{m_++1}$, see Figure 13. Choose one such extension, set $m_+ \leftarrow m_+ + 1$ and return to Step 1.

Step 2. If all the vertices of $D_n^{m_-}$ are already covered in the quasi-matching, stop the procedure; again, this can only happen when $D_n^{m_-}$ is short. If $D_n^{m_-}$ contains at least one uncovered vertex, it is again possible to extend it one step further. Choose one such extension, set $m_- \leftarrow m_- - 1$ and return to Step 2.

We explain now why this procedure will stop after a finite number of steps.

Suppose that D_n^k is a short diagonal, and that the quasi-matching covers D_n^{k+2}. Let $j = \#_N(D_n^k)$, then among the $n+1$ vertices of D_n^{k+1}, exactly j are connected to D_n^k, and therefore $\#_N(D_n^{k+1}) = n+1-j$. But this implies that among the n vertices of D_n^{k+2}, exactly $j-1$ are connected to D_n^{k+3}, *i.e.* $\#_N(D_n^{k+2}) = j-1$. This implies that the procedure stops with a finite value of m_+, and similarly with a finite value of m_-. We thus have obtained a matching of a tube of size n, $T_n^{(m_-,m_+)}$, extending the quasi-matching of $T_n^{(b,t)}$. There only remains to check that $T_n^{(m_-,m_+)}$ is actually an Aztec diamond of size n, *i.e.* that $m_+ - m_- = 2n$.

By construction, $\#_N(D_n^{m_+}) = 0$, and the above argument shows that

$$\#_N(D_n^{m_+-k}) = \begin{cases} \frac{k}{2} & \text{if } k \text{ is even} \\ n+1-\frac{k+1}{2} & \text{if } k \text{ is odd} \end{cases}. \tag{5}$$

In particular, $\#_N(D_n^{m_+-2n}) = n$. But this implies that $m_- = m_+ - 2n$ and the proof of the lemma is over. $\square$

Remark 4.2. *It follows from the proof of the preceding lemma that, for a quasi-matching of $T_n^{(b,t)}$ to exist, it is necessary that $n \geq (t-b)/2$.*

We now start with a matching of G, and show that, for n satisfying (4), we can extend it to a quasi-matching of a translate of the Aztec tube of size n $T_n^{(b,t)}$, for $n \geq \ldots$, which will conclude the proof of Theorem 4 by virtue of Lemma 5. We fix such a translate which contains G, and denote by $L_b, \ldots, L_t$ its W-E diagonals. We partition each line L_k, $k = b, \ldots, t$, into three gapless sets P_k^W, P_k^G, P_k^E, where $P_k^G = L_k \cap G$.

We will need the following lemma.

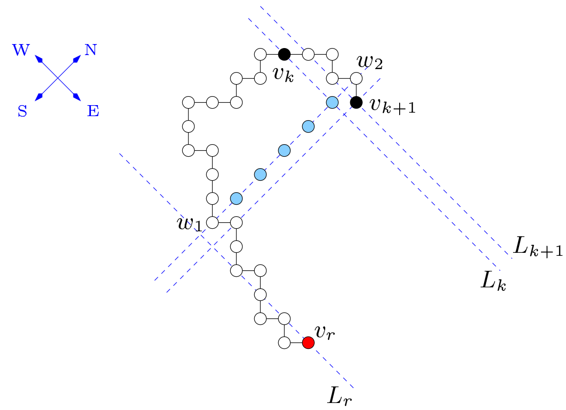

Figure 14: The construction in the proof of Lemma 6. Observe that the edges drawn on this picture are not necessarily edges of G.

Lemma 6. *For all k, $r \leq k \leq t-1$, we have*

$$|P_{k+1}^E| \geq |P_k^E| - 1, \tag{6}$$

and for all k, $r \leq k \leq t-2$, we have

$$|P_{k+2}^E| \geq |P_k^E|. \tag{7}$$

Proof of Lemma 6. Suppose that (6) does not hold for some $r \leq k \leq t-1$. Then the two eastern-most vertices v_k and v_{k+1} of G located on the two consecutive W-E diagonals L_k and L_{k+1} are not nearest neighbours and v_{k+1} is further east than v_k (see Figure 14). Since the set of vertices of G is connected, we can find two paths of vertices of G (white in the picture), one connecting v_k and the eastern-most vertex v_r of G on L_r, the other connecting v_k and v_{k+1}. These paths must cross the S-N diagonal located one-step west of that containing v_{k+1} at two vertices w_1 and w_2. Since the set of vertices of G is gapless, all vertices located between w_1 and w_2 on their common S-N diagonal (drawn in blue in the picture) must also belong to G. This contradicts the fact that v_k was the eastern-most vertex of G on L_k, and therefore (6) holds. The proof of (7) is similar. □

We are going to investigate the required conditions in order to extend the matching of G to a quasi-matching of the tube. We first fill in dominoes with at least one endpoint in P_r^E, then fill all layers P_k^E, $k = r+1, \ldots, t-1$, one after the other. We then do the same thing proceding southward, filling all P_k^E, $k = r-1, \ldots, b+1$, and then repeat this procedure with P_k^W, starting with P_l^W and filling first northward, and then southward.

So we first need to understand under which conditions we can fill the layer P_{k+1}^E once all the layers P_m^E, $m = r, \ldots, k$ have been filled. A similar argument as in Figure 13 together with (6) shows that this is possible as soon as

$$\#_N(P_k^E) \geq 1. \tag{8}$$

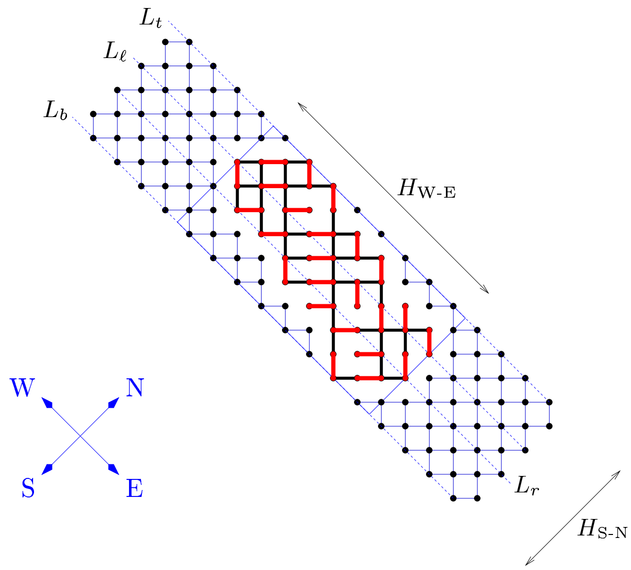

Figure 15: The graph G embedded in a tube. We want to extend the matching of G to a quasi-matching of the tube.

Using (6) and $\#_N(P_m^E) = P_m^E - \#_N(P_{m-1}^E)$, we have

$$\#_N(P_{k+2}^E) \geq \#_N(P_k^E) - 1. \tag{9}$$

We thus see that, in order to be able to extend the construction up to level $t-1$, it is sufficient that, at level r, we have

$$\#_N(P_r^E) \geq \lfloor \frac{t-r}{2} \rfloor, \tag{10}$$

$$\#_S(P_r^E) \geq \lceil \frac{t-r}{2} \rceil. \tag{11}$$

Indeed, the inequality (11) together with (6) implies that

$$\begin{aligned} \#_N(P_{r+1}^E) &= |P_{r+1}^E| - \big(|P_r^E| - \#_S(P_r^E)\big) \\ &= \#_S(P_r^E) + \big(|P_{r+1}^E| - |P_r^E|\big) \\ &\geq \#_S(P_r^E) - 1 \\ &\geq \lceil \frac{t-r}{2} \rceil - 1. \end{aligned} \tag{12}$$

Inequalities (9), (10) and (12) imply that (8) is satisfied for all k, $r \leq k \leq t-1$.

We get a similar condition to fill in layers P_k^E, $k = r-1, \ldots, b+1$: it is possible as soon as

$$\#_S(P_r^E) \geq \lfloor \frac{r-b}{2} \rfloor, \tag{13}$$

$$\#_N(P_r^E) \geq \lceil \frac{r-b}{2} \rceil. \tag{14}$$

Hence, the eastern side of the tube can be filled if the following condition is satisfied:

$$|P_r^E| \geq 2\max\left\{\lceil\frac{r-b}{2}\rceil, \lceil\frac{t-r}{2}\rceil\right\}. \tag{15}$$

Similarly, the western side of the tube can be filled if the following condition is satisfied:

$$|P_\ell^W| \geq 2\max\left\{\lceil\frac{\ell-b}{2}\rceil, \lceil\frac{t-\ell}{2}\rceil\right\}. \tag{16}$$

□

4.2 The case of a diamond embedded in a diamond

We now show on a specific example that Theorem 4 is essentially optimal in general. The case we consider is when G is itself an Aztec diamond, $G = A_m$. Let $0 < \rho_{\mathrm{i}} < \rho_{\mathrm{e}}$; we call *Aztec ring* of internal size ρ_{i} and external size ρ_{e} the graph obtained by removing from an aztec diamond $A_{\rho_{\mathrm{e}}}$ all the vertices of the concentric Aztec diamond $A_{\rho_{\mathrm{i}}}$; the *thickness* of the ring is the difference $\rho_{\mathrm{e}} - \rho_{\mathrm{i}}$. Choosing the best possible values for r and ℓ in that case, Theorem 4 tells us that a sufficient condition for an Aztec ring to be tileable is $\rho_{\mathrm{e}} \geq 3\rho_{\mathrm{i}} + 1$ if ρ_{i} is even, and $\rho_{\mathrm{e}} \geq 3\rho_{\mathrm{i}} + 3$ otherwise. The main result of this section is the following proposition, which shows that this estimate is essentially optimal.

Proposition 1. *The Aztec ring of internal size ρ_{i} and external size ρ_{e} is tileable if and only if at least one of the following conditions holds*

$$\rho_{\mathrm{e}} - \rho_{\mathrm{i}} \text{ is even,} \tag{17}$$

$$\rho_{\mathrm{e}} \geq 3\rho_{\mathrm{i}} - 1. \tag{18}$$

Proof. Let us first prove that under (17), the Aztec ring is always tileable. It is sufficient to notice that the Aztec diamond $A_{\rho_{\mathrm{e}}}$ can always be tiled using only horizontal edges; since in this matching there are no edges between vertices of $A_{\rho_{\mathrm{e}}}$ and vertices of $A_{\rho_{\mathrm{e}}} \setminus A_{\rho_{\mathrm{i}}}$, this provides a matching of the ring (see Figure 16 a).

Let us now show that there is always a matching of an Aztec ring of odd thickness, as soon as condition (18) is verified. By the previous result it is enough to construct a matching when $\rho_{\mathrm{e}} = 3\rho_{\mathrm{i}} - 1$. Indeed, this matching can then be extended to any Aztec ring of odd thickness, same internal size, and larger external size, since the added vertices form an Aztec ring with even thickness. When $\rho_{\mathrm{e}} = 3\rho_{\mathrm{i}} - 1$, we can partition the ring in such a way that we obtain eight subgraphs, which are all tileable. Four of the latter are actually Aztec diamond, while the remaining four can be tiled using parallel edges (two of them with horizontal edges, and the remaining two with vertical edges); see Figure 16 c.

Let us finally prove that it is never possible to tile an Aztec ring with $\rho_{\mathrm{e}} - \rho_{\mathrm{i}}$ odd and $\rho_{\mathrm{e}} < 3\rho_{\mathrm{i}} - 1$. Indeed suppose that we have constructed a matching of the Aztec ring. Together with a matching of the inner (concentric) Aztec diamond of size ρ_{i}, this provides a matching of an Aztec diamond of size ρ_{e}. As we already observed when deriving (5), all ρ_{e}

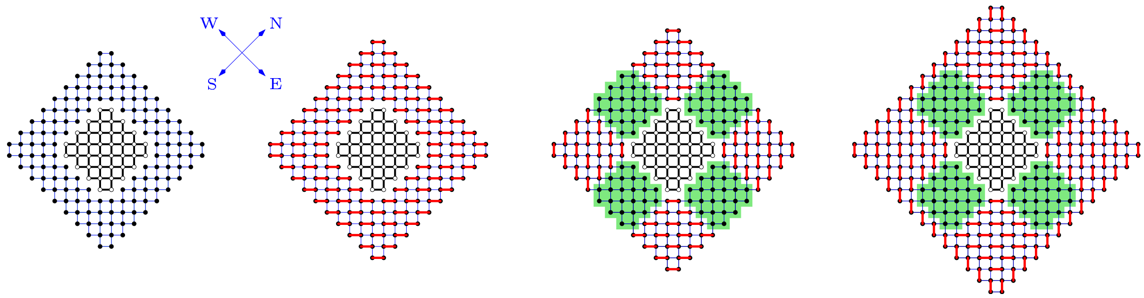

Figure 16: Four Aztec rings with $\rho_\mathrm{i} = 4$ and $\rho_\mathrm{e} = 8, 9, 11$ and 13. The second one is not tileable, while the other three are, since a ring is tileable if and only if either $\rho_\mathrm{e} - \rho_\mathrm{i}$ is even or $\rho_\mathrm{e} \geq 3\rho_\mathrm{i} - 1$.

dominoes on the northern-most W-E diagonal of A_{ρ_e} are oriented southwards, therefore the number of dominoes oriented southward on the W-E diagonal just north of the concentric Aztec diamond A_{ρ_i} is $\rho_\mathrm{e} - (\rho_\mathrm{e} - (\rho_\mathrm{i} + 1))/2$. However, since the matching should not connect vertices from the ring with vertices from A_{ρ_i}, at most $\rho_\mathrm{e} - \rho_\mathrm{i} + 1$ edges on this diagonal can be oriented southwards. But this implies that $\rho_\mathrm{e} \geq 3\rho_\mathrm{i} - 1$. □

4.3 Tiling other graphs

Our algorithm provides an efficient way of obtaining random matchings of various planar graphs, provided they can be embedded in $\mathbb{Z}^2$, in such a way that their complement in some Aztec diamond is tileable. To illustrate this, we show in Figure 17 how the 6-6-6, 4-8-8 and 4-6-12 (semi-)regular tesselations can be embedded (obviously, the remaing 7 (semi-)regular tesselations cannot be embedded in $\mathbb{Z}^2$ as they all contain triangles). Notice that the first two of the latter embeddings do belong to $\mathbf{G}$, but that this is not the case of the last one, which is not gapless. It is easy to remedy this by completing the embedding with new edges, as illustrated in Figure 18. In Figures 20 and 21, we show the dual random tilings for these three graphs, as described in Figure 19.

These three examples had already been generated using tricks such as urban renewal (see [7], [2]). On these examples, our algorithm works without practicing such surgery on the graph. It could be worth seeing whether such tricks combined with our method open up the possibility to cover a much larger class of graphs than the ones presented in this paper.

Acknowledgments

We are grateful to James Propp for his helpful comments on a first draft of this paper. We also thank the referee for useful suggestions.

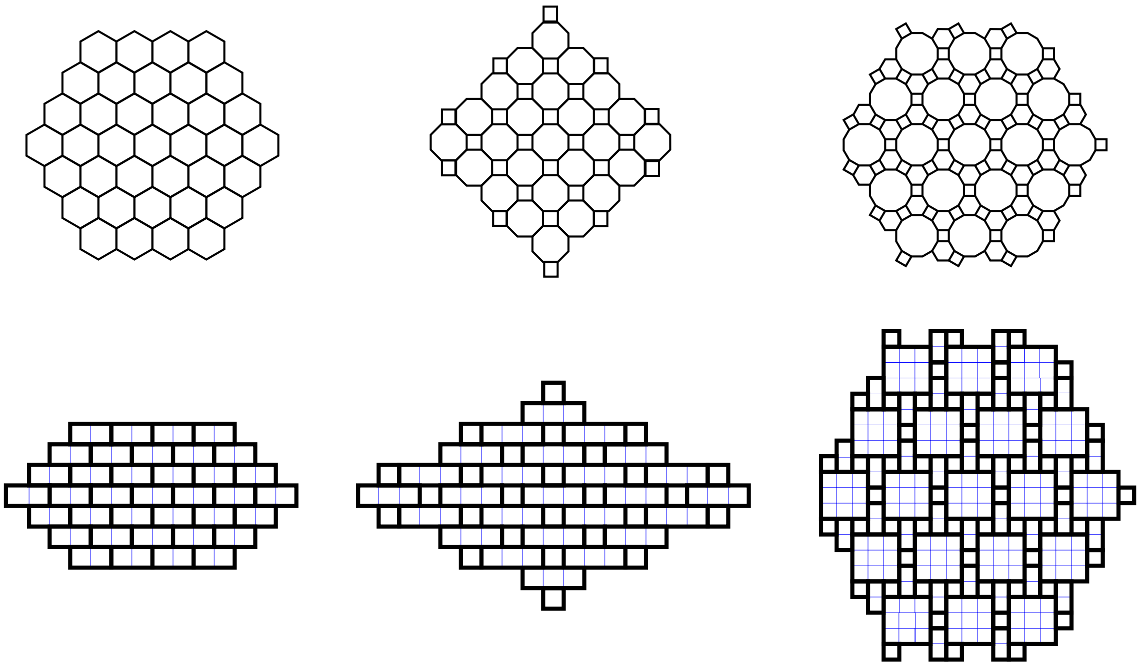

Figure 17: The 6-6-6, 4-8-8 and 4-6-12 (semi-)regular tesselations and their embedding in $\mathbb{Z}^2$. The shapes chosen in the last two cases are known as the fortress and the dungeon.

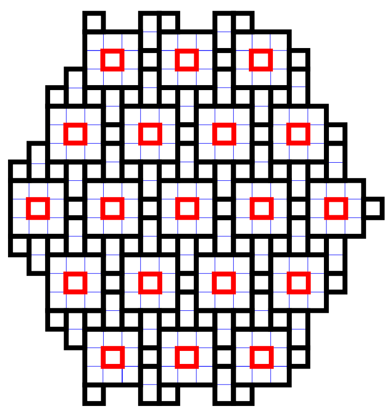

Figure 18: Extending the embedding of the 4-6-12 (semi-)regular tesselation in order to make it gapless.

References

[1] Mihai Ciucu. Perfect matchings of cellular graphs. *J. Algebraic Combin.*, 5(2):87–103, 1996.

[2] Mihai Ciucu. Perfect matchings and perfect powers. *J. Algebraic Combin.*, 17(3):335–375, 2003.

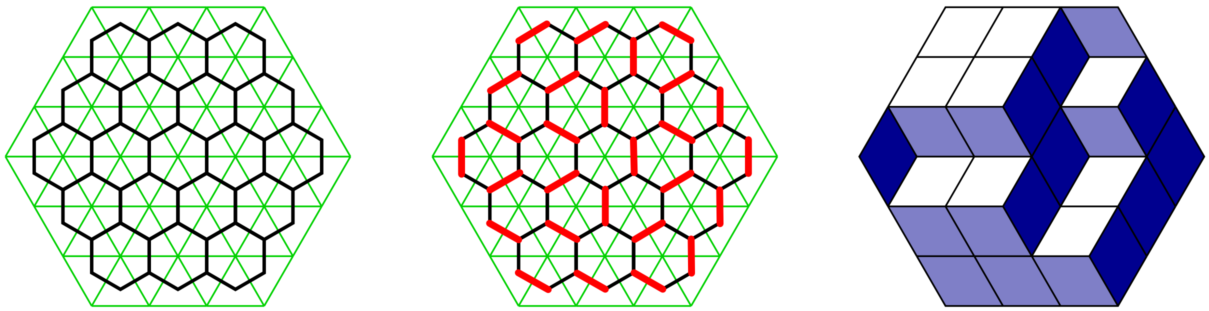

Figure 19: Left: The 6-6-6 tesselation and its dual. Middle: a matching of the graph. Right: The corresponding tiling of the hexagon with dual lozenges.

[3] Noam Elkies, Greg Kuperberg, Michael Larsen, and James Propp. Alternating-sign matrices and domino tilings. *J. Algebraic Combin.*, 1:111–132, 219–234, 1992.

[4] William Jockusch, James Propp, and Peter Shor. Random domino tilings and the arctic circle theorem. preprint, 1995; math.CO/9801068.

[5] Richard Kenyon, James Propp, and David Wilson. Trees and matchings. *Electron. J. Combin.*, 7:Research Paper 25, 34 pp. (electronic), 2000.

[6] Web page of the Undergraduate Research Project in Random Tilings. www.math.wisc.edu/∼propp/tiling/.

[7] James Propp. Generalized domino-shuffling. *Theoret. Comput. Sci.*, 303(2-3):267–301, 2003. Tilings of the plane.

[8] James Propp and David Wilson. How to get a perfectly random sample from a generic Markov chain and generate a random spanning tree of a directed graph. *J. Algorithms*, 27(2):170–217, 1998. 7th Annual ACM-SIAM Symposium on Discrete Algorithms (Atlanta, GA, 1996).

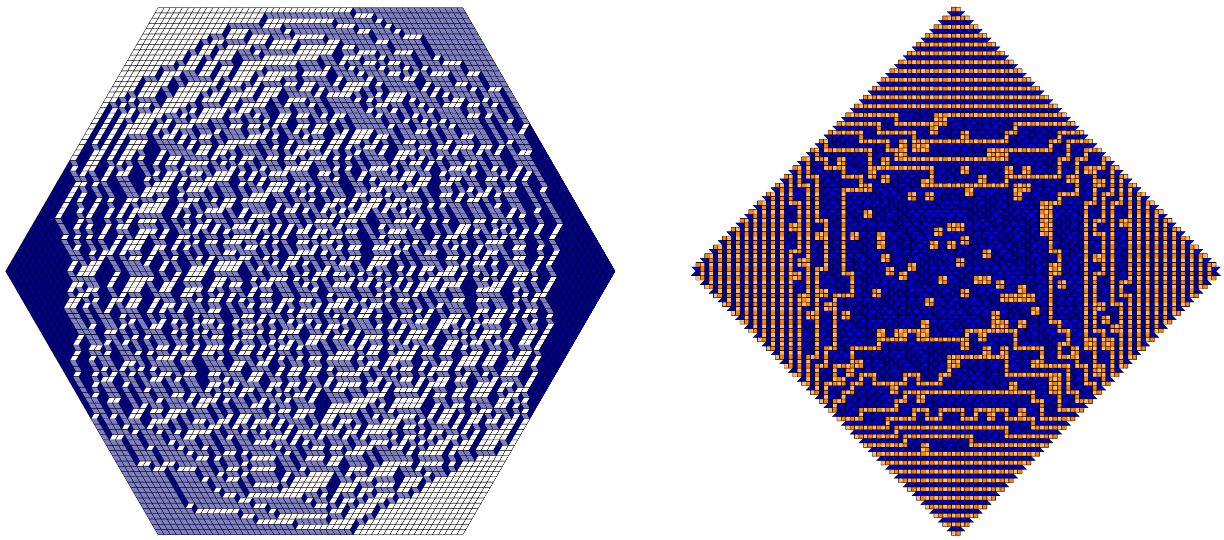

Figure 20: Examples of tiling of the duals of the 6-6-6 and 4-8-8 (semi-)regular tesselations generated by the algorithm.

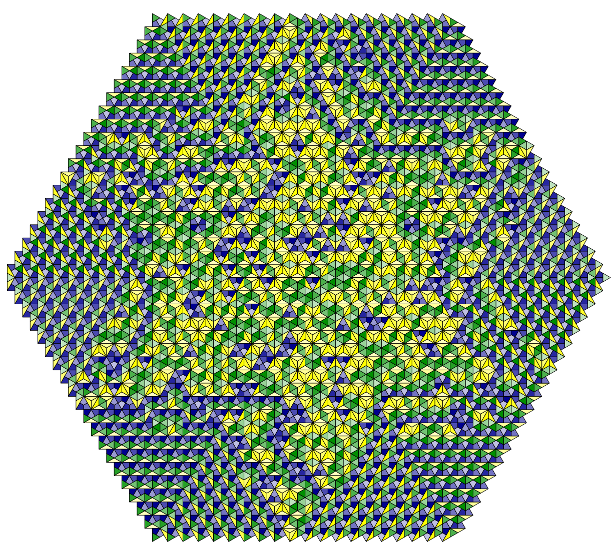

Figure 21: Example of tiling of the dual of the 4-6-12 semi-regular tesselations generated by the algorithm.