

Pfaffian Algorithms for Sampling Routings on Regions with Free Boundary Conditions^{*}

Russell A. Martin¹ and Dana Randall²

¹ School of Mathematics

² College of Computing and School of Mathematics,
Georgia Institute of Technology, Atlanta GA 30332, USA
{martin, randall}@math.gatech.edu

Abstract. Sets of non-intersecting, monotonic lattice paths, or *fixed routings*, provide a common representation for several combinatorial problems and have been the key element for designing sampling algorithms. Markov chain algorithms based on routings have led to efficient samplers for tilings, Eulerian orientations [8] and triangulations [9], while an algorithm which successively calculates ratios of determinants has led to a very fast method for sampling fixed routings [12]. We extend Wilson's determinant algorithm [12] to sample *free routings* where the number of paths, as well as the endpoints, are allowed to vary. The algorithm is based on a technique due to Stembridge for counting free routings by calculating the Pfaffian of a suitable matrix [11] and a method of Colbourn, Myrvold and Neufeld [1] for efficiently calculating ratios of determinants. As an application, we show how to sample tilings on planar lattice regions with free boundary conditions.

1 Introduction

Physicists study combinatorial structures on lattices in order to understand various physical systems. For example, tilings on planar lattice regions model systems of diatomic molecules, or dimers. By studying statistics of random configurations on families of regions of finite size (such as the $n \times n$ square or the Aztec diamond), physicists gain insight into the behavior of these systems on the infinite lattice, the so-called thermodynamic limit.

It is well known that the boundary of the region plays a crucial role. There are two relevant boundary effects. The first is the *shape* of the family of finite regions; the second is the *type* of boundary conditions defined for the regions. So far sampling has primarily been done for *fixed* boundary conditions, where the configurations are forced to precisely agree with the boundary. In the case of domino tilings this means that tiles are forced to cover all of the squares inside, and only inside, the region. Another important type of boundary condition considered permits all configurations that can be seen within a *window* in the shape of the region. Returning to tilings, this means that tiles can overlap the boundary (as long as the configuration can be extended to a tiling of the plane). In the context of tilings, these are commonly referred to as *free* boundary conditions.

^{*} Research supported in part by NSF Career Grant No. CCR-9703206




Fig. 1. Domino tilings on regions with fixed and free boundary conditions

One reason for studying free boundary conditions is that these eliminate the boundary effect due to the shape of the region (in the limit). This is not true for families of regions with fixed boundary conditions, where properties of random configurations can vary drastically according to their shape. Consider, for example, the entropy of the system, defined as $h(\Lambda) = \lim_{n \rightarrow \infty} \frac{\log \#(\lambda_n)}{\text{Area}(\lambda_n)}$, where $\Lambda = \{\lambda_n\}$ is a nested family of regions tending towards the infinite lattice and $\#(\lambda_n)$ is the number of tilings of λ_n . With fixed boundary conditions, the family of square regions has been proven to have *maximal* entropy over all finite families of regions Λ . In contrast, the family of Aztec diamonds is known to have lower entropy, which is related to the arctic circle phenomenon whereby frozen regions of the Aztec diamond emerge having a completely predictable local tiling [5]. On the other hand, with free boundary conditions, for any family of regions where the ratio of the length of the perimeter to the area of the region tends to zero, the entropy will converge to the same (maximal) value. In other words, statistics of tilings of square regions with free boundary conditions will agree with statistics of tilings of Aztec diamonds with free boundary conditions.

Several algorithms for sampling tilings on regions with fixed boundary conditions rely on a bijection between tilings and *fixed routings*, or sets of non-intersecting lattice paths where the number of paths and the position of their endpoints are fixed. The first is a Markov chain approach of Luby, Randall and Sinclair [8] which samples routings uniformly (and can be extended to the case where the paths are edge disjoint, but not necessarily vertex disjoint). A second approach, due to Wilson [12], uses the Gessel-Viennot method for enumerating routings by calculating a determinant [2] (and the close relationship between counting and sampling formalized by Jerrum, Valiant and Vazirani [4]). Wilson utilizes a technique introduced by Colbourn, Myrvold and Neufeld [1] which allows ratios of determinants of closely related matrices to be computed quickly without having to evaluate both determinants.

In this paper we sample *free routings*, or sets of non-intersecting lattice paths where the positions of the endpoints of the paths, as well as the number of paths, are allowed to vary. Our result relies on Stembridge's algorithm counting the number of free routings of a region by evaluating a Pfaffian [11]. We adapt the method of Colbourn, Myrvold and Neufeld to allow ratios of Pfaffians to be evaluated quickly, a special case of a technique of Kenyon for calculating statistics of random tilings [6]. The running time of our algorithm is $O(l^2 n)$, where n is the size of the region and l is the maximal number of paths in

a routing. Typically $l = O(\sqrt{n})$, yielding an $O(n^2)$ algorithm. We apply this sampling method to generate random domino and lozenge tilings of hexagonal regions with free boundary conditions.

The remainder of the paper is organized as follows. In section 2 we review the counting techniques of Gessel-Viennot and Stembridge for fixed and free routings, respectively. In section 3 we present our algorithm for uniformly sampling free routings. Finally, in section 4 we show the bijections between free routings and tilings on regions with free boundary conditions which allow us to sample these tilings efficiently.

2 Background: Counting routings

First we begin with an overview of the method of Gessel and Viennot for counting fixed routings and that of Stembridge for counting free routings. Wilson shows how to sample fixed routings using self-reducibility and iterative applications of the Gessel-Viennot method. We give a similar method to sample free routings, utilizing Stembridge's method for counting free routings.

2.1 The Gessel-Viennot method

Gessel and Viennot[2, 3], and Lindström[7] introduce a method for finding the number of non-intersecting paths, with specified sources and sinks, in certain directed graphs by computing a determinant of a matrix. For their technique to work, the graph must be directed and acyclic. Furthermore, the sources and sinks must satisfy a condition known as *compatibility*. In this definition, we require that both the set of sources \mathcal{S} and the set of sinks \mathcal{T} be ordered.

Definition 1. *Let D be a directed acyclic graph. The ordered sets \mathcal{S} and \mathcal{T} are said to be compatible if $s < s'$ in \mathcal{S} and $t < t'$ in \mathcal{T} implies that every $s - t'$ path intersects every $s' - t$ path.*

Thus, if there is a set of l non-intersecting paths using sources $s_1 < s_2 < \dots < s_l$ and sinks $t_1 < t_2 < \dots < t_l$, then it must be the case that s_i is joined to t_i for all i . We call such a set of l non-intersecting paths a *fixed routing* of D .

Let D denote an acyclic directed graph with compatible sources $\mathcal{S} = \{s_1, \dots, s_l\}$ and sinks $\mathcal{T} = \{t_1, \dots, t_l\}$. Let p_{ij} denote the number of directed paths in D with source s_i and sink t_j . Since the graph is assumed to be acyclic, this number is finite for all i and j . Let P be the matrix with entries p_{ij} .

We have the following theorem [2, 11]:

Theorem 1. *With D , \mathcal{S} , \mathcal{T} , (\mathcal{S} and \mathcal{T} compatible) and P as above, the number of non-intersecting sets of l paths in D is equal to $\det(P)$.*

If D is not acyclic, or if \mathcal{S} and \mathcal{T} are not compatible, then the preceding theorem fails. (See [11] for an example for which the theorem fails.)

Theorem 2 (Wilson [12]). *Let D be a planar, acyclic digraph with n vertices, having compatible sources and sinks. Fixed routings of D can be uniformly sampled in $O(l^{1.688}n)$ time.*

2.2 Stembridge's extension

Stembridge[11] extends the Gessel-Viennot method to count free routings of a directed acyclic graph, D , with sources \mathcal{S} and sinks \mathcal{T} . In the case of free routings, the number of paths is no longer fixed, so S is really the set of *potential* sources and T is the set of *potential* sinks. Also, it is no longer always true that s_i will be joined to t_i , as was true in the case of fixed routings.

If $s_i \in \mathcal{S}$ is a source in a free routing, we say that s_i is *used* in the routing; otherwise s_i is *unused*. Here we assume there are l sources and l sinks. First, we need a bit of linear algebra.

Definition 2. Let B be a $2n \times 2n$ skew-symmetric matrix (i.e. $B^T = -B$), and let

$$\pi = \{\{i_1, j_1\}, \{i_2, j_2\}, \dots, \{i_n, j_n\}\}$$

be a partition of the set $\{1, \dots, 2n\}$ into pairs. Let

$$b_\pi = \text{sgn} \begin{pmatrix} 1 & 2 & 3 & 4 & \dots & 2n-1 & 2n \\ i_1 & j_1 & i_2 & j_2 & \dots & i_n & j_n \end{pmatrix} b_{i_1, j_1} b_{i_2, j_2} \cdots b_{i_n, j_n}.$$

The Pfaffian of B , denoted $\text{Pf}(B)$, is defined by

$$\text{Pf}(B) = \sum_{\pi} b_\pi.$$

Theorem 3. If B is a skew-symmetric matrix of even size, then $\det(B) = \text{Pf}(B)^2$.

A skew-symmetric matrix, Q , will take the role of the matrix P in theorem 1, but instead of the determinant of Q , we look at its Pfaffian. For $1 \leq i < j \leq l$ and $1 \leq h < k \leq l$, let $\alpha_{ij}(h, k)$ denote the number of non-intersecting paths in D with sources s_i, s_j , and sinks t_h, t_k . We find $\alpha_{ij}(h, k) = \det \begin{pmatrix} p_{ih} & p_{ik} \\ p_{jh} & p_{jk} \end{pmatrix}$ using theorem 1, where, recall, p_{ih} is the number of paths from s_i to t_h .

Let $q_{ij} = \sum_{h < k} \alpha_{ij}(h, k)$. Then q_{ij} is the number of pairs of non-intersecting paths with sources s_i and s_j , where the sinks range over all pairs where t_h precedes t_k in the ordering of \mathcal{T} . Finally, let q_i denote the number of paths with source s_i to any sink in \mathcal{T} .

We assume that l is odd; if not, we can add an additional isolated vertex s_{l+1} to \mathcal{S} . The following is due to Stembridge[11]:

Theorem 4. Let $\mathcal{S} = (s_1, \dots, s_l)$ be an l -tuple of vertices in an acyclic digraph D , with l odd. Let \mathcal{T} be an ordered subset of vertices that is compatible with \mathcal{S} . Let Q be the skew-symmetric matrix where the upper triangular entries are given by

$$[Q]_{ij} = (-1)^{i+j-1} + q_{ij}^*$$

for $1 \leq i < j \leq l+1$, where $q_{ij}^* = q_{ij}$ for $j \leq l$ and $q_{i, l+1}^* = q_i$. Then $\Phi = \text{Pf}(Q)$ is the number of free routings of D .

The matrix Q looks like:

$$Q = \begin{pmatrix} 0 & 1 + q_{12} & \dots & -1 + q_{1l} & 1 + q_1 \\ -1 - q_{12} & 0 & \dots & 1 + q_{2l} & -1 + q_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ -1 - q_1 & 1 - q_2 & \dots & -1 - q_l & 0 \end{pmatrix}.$$

Stembridge uses this theorem to study shifted tableaux, plane partitions, and Schur's Q -functions. As we will see in section 4 it can be used to count and generate tilings with free boundary conditions. We first give an extension of Stembridge's result.

2.3 Fixing sources

We can extend theorem 4 to count the number of free routings where we specify that certain sources must be used (or unused) in a routing.

Informally, if s_i is used as a source then we replace terms $\pm 1 + q_{ij}$ in Q by q_{ij} , and if not used, by ± 1 . The following theorem formalizes this:

Theorem 5. *Let $\mathcal{S} = (s_1, \dots, s_l)$ be an l -tuple of vertices in an acyclic digraph D , with l odd. Let \mathcal{T} be an ordered subset of vertices that is compatible with \mathcal{S} . Suppose $\mathcal{S}_{in}, \mathcal{S}_{out} \subseteq \mathcal{S}$ with $\mathcal{S}_{in} \cap \mathcal{S}_{out} = \emptyset$. Let Q be the skew-symmetric matrix where the upper triangular entries are given by*

$$[Q]_{ij} = \begin{cases} 0 & \text{if } s_i \in \mathcal{S}_{in} \text{ and } s_j \in \mathcal{S}_{out}, \text{ (or vice-versa)} \\ q_{ij} & \text{if } j \leq l \text{ and } (s_i \text{ and/or } s_j \text{ in } \mathcal{S}_{in}, \text{ neither in } \mathcal{S}_{out}) \\ (-1)^{i+j-1} & \text{if } j \leq l \text{ and } (s_i \text{ and/or } s_j \text{ in } \mathcal{S}_{out}, \text{ neither in } \mathcal{S}_{in}) \\ q_i & \text{if } j = l + 1 \text{ and } s_i \in \mathcal{S}_{in} \\ (-1)^{i+l} & \text{if } j = l + 1 \text{ and } s_i \in \mathcal{S}_{out} \\ (-1)^{i+j-1} + q_{ij}^* & \text{otherwise, where } q_{ij}^* \text{ is as in theorem 4} \end{cases}$$

for $1 \leq i < j \leq l + 1$. Then $\Phi = \text{Pf}(Q)$ is the number of free routings of D with \mathcal{S}_{in} included in the set of used sources, and \mathcal{S}_{out} in the set of unused sources.

Proof. For $J \subseteq \{1, 2, \dots, n\}$, let A_J denote the square submatrix of A obtained by selecting the rows and columns indexed by J . We use the result (from [11, lemma 4.2]) that (for n even, A and B $n \times n$ in size) we can write

$$\text{Pf}(A + B) = \sum_J (-1)^{\sigma(J) - \frac{|J|}{2}} \text{Pf}[A_J] \text{Pf}[B_{J^c}] \quad (1)$$

where $\sigma(J) = \sum_{j \in J} j$, and the sum is taken over all partitions J, J^c of $\{1, \dots, n\}$ with $|J|$ even.

Decompose Q into a sum of two matrices, A and B , where $[A]_{ij} \in \{0, 1, -1\}$ and $[B]_{ij} \in \{0, q_{ij}^*, -q_{ij}^*\}$. Now apply the above result for the Pfaffian of a sum. We have $\text{Pf}[A_J] = 0$ if $J \cap \mathcal{S}_{in} \neq \emptyset$ since A will contain a row of zeros. Similarly, $\text{Pf}[B_{J^c}] = 0$ if $J^c \cap \mathcal{S}_{out} \neq \emptyset$. So the only terms that survive in the sum (1)

are those with $\mathcal{S}_{in} \subseteq J^C$ and $\mathcal{S}_{out} \subseteq J$. Note that if q_{ij} or q_i appears in one of the terms of Φ , that term corresponds to a set of paths (not necessarily non-intersecting) that uses s_i . Thus, if $s_i \in \mathcal{S}_{in}$, by the choice of the entries of Q we ensure that s_i is used in every routing of D , as one of the q_{ij} 's or q_i will appear in each term of Φ . Similarly, if $s_i \in \mathcal{S}_{out}$ then none of the q_{ij} 's or q_i appears in Φ , so that s_i is unused as a source in every routing of D . \square

3 Generating random routings

We present an algorithm to uniformly generate a free routing of a planar acyclic digraph D with compatible sources and sinks. This algorithm is similar to the determinant algorithm of [12] for generating fixed routings. Once again, we assume that $|\mathcal{S}| = |\mathcal{T}| = l$, with l odd.

We use self-reducibility to find the routing by building paths one edge at a time. We move through the graph, deciding (probabilistically) if a source s_i is used in the routing. Then, if used, we select one of its out-going edges, with appropriate probabilities, and add it to the routing, thereby starting a path using s_i . We push the source forward to w , the other end of the selected edge, and will eventually complete a path from s_i into \mathcal{T} . We use the fact that ratios of Pfaffians can be computed efficiently to determine the probability of using a particular source or edge. The following theorem is analogous to the result in [1] for ratios of determinants of matrices that differ by a single row.

Theorem 6. *Let A be an invertible, skew-symmetric matrix and let B be a skew-symmetric matrix which differs from A by only the i th row and column. Then*

$$\frac{\text{Pf}(B)}{\text{Pf}(A)} = [BA^{-1}]_{ii}.$$

Proof. The proof relies on a closely related fact that if A is an invertible matrix and C differs from A by only the i th row, then

$$\frac{\det(C)}{\det(A)} = [CA^{-1}]_{ii}$$

(which follows from Cramer's rule). Given A , an invertible, skew-symmetric matrix and B , a skew-symmetric matrix which differs from A by only the i th row and column, let C be the matrix formed by replacing the i th row of A with the i th row of B (so C differs from A by the i th row and from B by the i th column). Assume first that C is invertible. Then we have that $\det(C)/\det(A) = [CA^{-1}]_{ii}$ and $\det(B)/\det(C) = [B^T C^{-1}]_{ii} = [C^{-1}B]_{ii}$. Finally, we use the fact that since A and C differ in only the i th row and $A_{ii} = C_{ii} = 0$ since A and B are skew-symmetric, then the i th rows of A^{-1} and C^{-1} must agree. Hence,

$$\begin{aligned} \left(\frac{\text{Pf}(B)}{\text{Pf}(A)}\right)^2 &= \frac{\det(B)}{\det(A)} = \frac{\det(B)}{\det(C)} \cdot \frac{\det(C)}{\det(A)} = [C^{-1}B]_{ii}[CA^{-1}]_{ii} \\ &= [A^{-1}B]_{ii}[BA^{-1}]_{ii} = ([BA^{-1}]_{ii})^2. \end{aligned} \tag{2}$$

If C is not invertible, let B' be obtained from B by perturbing the i th row of B by $\varepsilon \cdot$ (random vector) and the i th column so that B' is skew-symmetric (and differs from A only in the i th row and i th column). With C' as the matrix formed by replacing the i th row of A by the i th row of B' we proceed as before (C' is invertible), then let $\varepsilon \rightarrow 0$ (so $B' \rightarrow B$) to get the same result as in (2).

Taking square roots, and recalling that $A^{-1} = \text{adj}(A)/\det(A)$, where $\text{adj}(A)$ is the (classical) adjoint of A , we can write (2) as

$$\text{Pf}(B) \det(A) = \pm \text{Pf}(A)[B \text{adj}(A)]_{ii}. \quad (3)$$

Taking an invertible skew-symmetric matrix A and letting $B = A$, the sign in that case is $+$, and by continuity the sign for a whole neighborhood of the parameter values is also $+$. By taking partial derivatives and evaluating at 0, we see that the coefficients of the polynomials must be equal, so that the sign in (3) is everywhere $+$. \square

We use the Sherman-Morrison formula for updating A^{-1} after changing a single row or column of the matrix A . In our case we will be changing both a row and a column, but we can update the inverse by applying the Sherman-Morrison formula twice. Updating A^{-1} can be done in $\Theta(l^2)$ time using this method (for details see, e.g., [1, 12]). (The Sherman-Morrison formula for updating an inverse has shown some numerical instability in practice; we may achieve greater numerical stability by using other schemes for updating A^{-1} at a small cost in the running time of the algorithm.) Now we are ready to describe the algorithm.

The input to the algorithm is the planar digraph, D , having n vertices and m edges, and sets \mathcal{S} and \mathcal{T} , the sources and sinks, respectively. The variable x_i records the current position of source i , and the array R records the routing as it is constructed. We maintain a matrix Q , initially equal to the matrix Q of theorem 4, and a matrix U , initially the inverse of Q , which we use to compute probabilities of using sources or edges in the routing. Q and U are updated as we move through the digraph. We compute $P[v, i]$, which will be the number of paths from vertex v to sink t_i , and $\hat{P}[v, i]$, the number of paths from v to any of the sinks t_i, t_{i+1}, \dots, t_l . (We use the $\hat{P}[v, i]$'s to help initialize the matrix Q in time $O(l^3)$ instead of $O(l^4)$, and later for updating entries of Q as we move through D .) With " $v \rightarrow w$ " denoting that there is a directed edge from v to w the algorithm is:

FreeRoute($D, \mathcal{S}, \mathcal{T}$)

1. Do a topological sort on D , numbering the vertices 1 through n , so that $v \rightarrow w$ implies $v < w$. ($O(n)$ time.)
2. For $v = 1$ to n , set $q_v = 0$. ($O(n)$ time.)
3. For $i = l$ down to 1 (Dynamic programming step) ($O(ln)$ time.)
 - (a) Set $x_i = s_i$.
 - (b) For $v = n$ down to 1
 - i. If $v = t_i$, set $P[v, i] = 1$, else set $P[v, i] = \sum_{w:w \rightarrow v} P[w, i]$.
($P[v, i]$ now contains the number of paths from v to t_i .)
 - ii. Set $q_v = q_v + P[v, i]$.

- iii. If $i = l$, set $\hat{P}[v, i] = P[v, i]$, else set $\hat{P}[v, i] = \hat{P}[v, i + 1] + P[v, i]$.
 $(O(l^3)$ time.)
4. For $i = 1$ to l
- (a) Find v such that $v = s_i$.
 - (b) Set $q_i = q_v$. (Initialize q_i 's.)
 - (c) For $j = i + 1$ to l
 - i. Find w such that $w = s_j$.
 - ii. Set $q_{ij} = \sum_{k=1}^{l-1} \det \begin{pmatrix} P[v, k] & \hat{P}[v, k + 1] \\ P[w, k] & \hat{P}[w, k + 1] \end{pmatrix}$. (Initialize q_{ij} 's.)
5. Initialize the matrix Q as in theorem 4 using the q_i 's and q_{ij} 's, find $U = Q^{-1}$, and set $\mathcal{S}_{in} = \emptyset$ and $\mathcal{S}_{out} = \emptyset$.
 $(O(l^3)$ time.)
6. For $v = 1$ to n , if $v = x_i$ for some i then
- (a) If $v \in \mathcal{S} \setminus (\mathcal{S}_{in} \cup \mathcal{S}_{out})$ then decide if v is used as a source (see details below). If it is, add v to \mathcal{S}_{in} . If not, add v to \mathcal{S}_{out} and set $R[v] = 0$.
In either case, update row and column i of Q and U . $(O(n + l^3)$ time.)
 - (b) If $v \in \mathcal{S}_{in}$ then decide which edge leaving v to include in the path of the routing (see details and remarks). Let w be the other endpoint of this edge.
If $w = s_k$ for some k , see remark 1 below.
Set $R[v] = w$, $x_i = w$, and add w to \mathcal{S}_{in} .
Set $q_i = q_w$. Update the i th row and column of Q and update U . $(O(l^2 n)$ time.)

Remark 1. In step 6(b), we may try to push $v = x_i$ forward to an (as yet) unused source $w = s_k \in \mathcal{S}$. In this case, we want to add w to \mathcal{S}_{out} so that it is not used during some later step to begin a different path. However, we also want to add w to \mathcal{S}_{in} so that in later iterations of step 6(b) we push w forward to complete a full path into \mathcal{T} that started from s_i . This conflicts with the condition of theorem 5 that $\mathcal{S}_{in} \cap \mathcal{S}_{out} = \emptyset$. We get around this difficulty as follows: Remove x_i from \mathcal{S}_{in} and add it to \mathcal{S}_{out} , then add $w = s_k$ to \mathcal{S}_{in} so that it is pushed forward in later steps of the algorithm. Update row and column i of Q to reflect that x_i is unused, then row and column k so that s_k is used, and update U accordingly with successive applications of the Sherman-Morrison formula. Finally, set $R[v] = w$ to join the path between s_i and w to the path from w into \mathcal{T} . We will see examples of digraphs in which this situation might arise in section 4 where we consider tilings of reduced Aztec diamonds.

Remark 2. During step 6(b) it is possible that $x_i \in \mathcal{T}$ but we might still push x_i forward. This could occur if x_i has out-neighbors that are also in \mathcal{T} . Informally, in this situation we may consider that x_i is joined to a phantom sink by a single (phantom) edge. Pushing x_i forward to this phantom sink corresponds to terminating the path at x_i and not continuing to any of x_i 's neighbors. In practice, we need not handle this situation as a special case, since we can examine all of the out-neighbors of x_i in turn and if we reject using any of them then terminate the path, i.e., x_i is not pushed.

Details for step 6(a): In this step, we determine if the source s_i is used in a routing. The probability that s_i is used is given by $\frac{\text{Pf}(Q')}{\text{Pf}(Q)}$, where Q' is a skew-symmetric matrix differing from Q in the i th row and i th column. In particular, the i th row of Q' can be found using theorem 5, where we apply the theorem

with s_i used in the set of current potential sources (the x_j 's, restricting \mathcal{S}_{in} and \mathcal{S}_{out} to that set). We use theorem 6 to compute this probability as the dot product of the new i th row of Q' with column i of U . If v is used, we replace the i th row of Q by the i th row of Q' to reflect this (and then update the i th column of Q so that it remains skew-symmetric), and add v to \mathcal{S}_{in} . If v is not used, we update row and column i of Q as appropriate in theorem 5, where v is now in \mathcal{S}_{out} . In either case, we update U (so that it is still equal to Q^{-1}), using two successive applications of the Sherman-Morrison formula, once for changing row i of Q , and again for changing column i . Updating U takes time $\Theta(l^2)$, and hence the total time spent in step 6(a) is $O(n + l^3)$.

Details for step 6(b): Moving the source x_i forward in step 6(b) changes the i th row and column of Q . As before, the probability that the edge $v \rightarrow w$ is used is $\frac{\text{Pf}(Q')}{\text{Pf}(Q)}$, where Q' is the matrix with w used as a source in place of x_i . If this edge is taken, we update Q (and U) by replacing the i th row and column of Q with those of Q' . In the special case that $w \in \mathcal{S}$, we proceed as outlined in remark 1. The time to update U (at any instance when Q is updated) is $\Theta(l^2)$, so the total time spent in step 6(b) is $O(l^2n)$.

We have demonstrated the following theorem:

Theorem 7. *Let D be a planar acyclic digraph with n vertices, having compatible sources and sinks. `FreeRoute` uniformly samples a free routing of D in time $O(l^2n)$.*

4 Lattice paths

In this section, we demonstrate applications of the techniques from the previous section. We show how to generate random domino tilings of the reduced Aztec diamond with free boundary conditions and lozenge tilings of the hexagon with free boundary conditions. The key idea is the existence of a bijection between the set of tilings of this region and the set of free routings in a related digraph. For details of the analogous bijections in the case of fixed boundary conditions, we refer the reader to [8].

4.1 Domino tilings of the reduced Aztec diamond

The *reduced Aztec diamond of order n* , denoted Γ_n , is a region composed of $2n^2$ unit squares arranged as $2n$ centered rows of squares, where the k th row has $\min\{2k - 1, 4n - 2k + 1\}$ squares in it. A domino tiling is a cover of Γ_n using non-overlapping dominoes, where a domino covers two adjacent squares. A domino tiling with *free boundary conditions* is a tiling in which all the squares of Γ_n are covered, but the dominoes are allowed to “stick out” of (or overlap) the boundary of the region. We assume that we know the orientation of a domino that overlaps the boundary, i.e., a single square (or half-domino) is designated as the bottom, top, left or right half of a domino.

Given a tiling of Γ_n with free boundary conditions (or simply, a free tiling), we define a routing of a digraph, D_n . To get D_n , first color the left square of row n of Γ_n black, then extend the coloring to Γ_n using alternating black and white squares (as on the underlying infinite chessboard). Mark the midpoint of each vertical edge that has a black square to its right. Fix $(0,0)$ as the coordinates of the point on the left edge of row n . Add $n + 1$ additional points at coordinates $(-1, -1), (0, -2), (1, -3), \dots, (n - 1, -n - 1)$, and another $n + 1$ points at $(n, n), (n + 1, n - 1), (n + 2, n - 2), \dots, (2n, 0)$. Join a point with coordinates (x, y) to the points $(x + 1, y + 1), (x + 1, y - 1)$, and $(x + 2, y)$. Finally, delete edges that lie completely outside the boundary of Γ_n .




Fig. 2. The reduced Aztec diamond

The marked points form the vertex set of D_n , and the edges of D_n are those that remain between points after the deletion step. Direct edges from left to right. Starting at the source in the top square, label the sources $s_1, s_2, \dots, s_{2n+1}$ in the counterclockwise direction, and then label each sink t_i where s_i is the last unmatched source. The left picture of figure 2 is Γ_4 , the reduced Aztec diamond of order 4, along with the sources and sinks of D_4 . The right picture is the digraph D_4 .

Theorem 8. *There is a bijection between free boundary tilings of Γ_n and free routings of D_n .*

Proof. Given a free tiling of Γ_n we map it to a free routing of D_n as follows: Examine the sources in this order: $s_n, s_{n-1}, \dots, s_1, s_{n+1}, s_{n+2}, \dots, s_{2n+1}$. It's possible that no source lies on the edge of a domino, in which case the routing is empty. Otherwise, the routing consists of the paths constructed as follows: If s_n lies on the edge of a domino, this determines the first edge in a path starting at s_n (otherwise move onto s_{n-1}). Connect s_n to the unique vertex in D_n that lies on the right side of the domino. This new vertex must lie on the left side of another domino, so repeat this process. Stop when we reach a vertex in \mathcal{T} that does not have a domino to its right. Choose the next source, in the prescribed order, that is not on a path already constructed, and repeat this procedure. The paths are non-intersecting since dominoes cannot overlap and because of the order in which the sources were examined. See figure 3 for an example of a free boundary tiling of Γ_4 and the corresponding routing. (An arrow in the tiling

points to the location of the other half of a domino that overlaps the boundary.) The proof that this map forms a bijection is analogous to the proof given in [8] which establishes a similar bijection between domino tilings of regions with fixed boundary conditions and fixed routings of related regions. \square




Fig. 3. A domino tiling with free boundary conditions and its free routing

It follows from this connection between tilings of Γ_n and routings of D_n that we can generate free boundary tilings of Γ_n by using the algorithm given in section 3 for generating free routings of D_n .

4.2 Lozenge tilings of the hexagon

We use a similar approach as in the previous section to generate lozenge tilings of a hexagonal region of the triangular lattice with free boundary conditions. There is a bijection between the collection of free boundary tilings, and the set of free routings of a related digraph.

Let H_n denote a hexagonal region on the triangular lattice with n edges on each side. A *lozenge tiling* of H_n is a covering of the region with lozenges, where a lozenge covers two adjacent triangles, and lozenges do not overlap. As in the previous section, a lozenge tiling of H_n with *free boundary conditions* is a tiling in which lozenges may overlap the boundary of the region. We describe a digraph, G_n , associated with H_n , in which free routings correspond to free boundary tilings of H_n . First, augment H_n to get a region \hat{H}_n by adding the triangles in the underlying lattice that share an edge with the boundary of H_n . Mark the midpoint of each vertical edge in \hat{H}_n . These marked points form the vertex set of G_n . Join two points if they lie on adjacent triangles. These are the edges of G_n . Direct these edges from left to right. A free boundary lozenge tiling of H_n corresponds to a free routing of G_n . Again, the proof of this bijection follows analogously to the proof given in [8] for establishing a bijection between fixed lozenge tilings and fixed routings. Figure 4 provides a pictorial illustration of this correspondence.

Applying the *FreeRoute* algorithm of section 3 allows us to uniformly generate free routings of G_n , which we may then map to their corresponding free boundary tilings of H_n .




Fig. 4. A lozenge tiling with free boundary conditions and its free routing

Acknowledgements: We would like to thank Dan Klain, Tom Morley and the referees for many useful comments.

References

1. C. J. Colbourn, W. J. Myrvold, and E. Neufeld. Two algorithms for unranking arborescences. *J. Algorithms*, 20:268–281, 1996.
2. I. M. Gessel and X. G. Viennot. Binomial determinants, paths, and hook length formulae. *Advances in Mathematics*, 58:300–321, 1985.
3. I. M. Gessel and X. G. Viennot. Determinants, paths, and plane partitions. Preprint, 1989.
4. M. R. Jerrum, L. G. Valiant, and V. V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986.
5. W. Jockusch, J. Propp, and P. Shor. Random domino tilings and the arctic circle theorem. Preprint, 1995.
6. R. Kenyon. Local statistics of dimers on a lattice. *Annales de l’Institut Henri Poincaré – Probabilités et Statistiques*, 33:591–618, 1997.
7. B. Lindström. On the vector representations of induced matroids. *Bulletin of the London Mathematical Society*, 5:85–90, 1973.
8. M. Luby, D. Randall, and A. Sinclair. Markov chain algorithms for planar lattice structures. *Proc. 36th IEEE Symposium on Foundations of Computer Science*, 150–159, 1995.
9. L. McShine and P. Tetali. On the mixing time of the triangulation walk and other Catalan structures. *Randomization Methods in Algorithm Design, DIMACS-AMS* 43, 1998.
10. J. Sherman and W. J. Morrison. Adjustment of an inverse matrix corresponding to changes in the elements of a given column or a given row of the original matrix. *The Annals of Mathematical Statistics* 20:621, 1949.
11. J. R. Stembridge. Non-intersecting paths, Pfaffians, and plane partitions. *Advances in Mathematics*, 83:96–131, 1990.
12. D. B. Wilson. Determinant algorithms for random planar structures. *Proc. 8th Annual ACM-SIAM Symposium on Discrete Algorithms*, 258–267, 1997.