Kasteleyn cokernels

Greg Kuperberg *

Department of Mathematics
University of California, Davis, CA 95616
greg@math.ucdavis.edu

Submitted: August 23, 2001; Accepted: June 24, 2002
MR Subject Classifications: 05A15, 11C20

Abstract

We consider Kasteleyn and Kasteleyn-Percus matrices, which arise in enumerating matchings of planar graphs, up to matrix operations on their rows and columns. If such a matrix is defined over a principal ideal domain, this is equivalent to considering its Smith normal form or its cokernel. Many variations of the enumeration methods result in equivalent matrices. In particular, Gessel-Viennot matrices are equivalent to Kasteleyn-Percus matrices.

We apply these ideas to plane partitions and related planar of tilings. We list a number of conjectures, supported by experiments in Maple, about the forms of matrices associated to enumerations of plane partitions and other lozenge tilings of planar regions and their symmetry classes. We focus on the case where the enumerations are round or q-round, and we conjecture that cokernels remain round or q-round for related "impossible enumerations" in which there are no tilings. Our conjectures provide a new view of the topic of enumerating symmetry classes of plane partitions and their generalizations. In particular we conjecture that a q-specialization of a Jacobi-Trudi matrix has a Smith normal form. If so it could be an interesting structure associated to the corresponding irreducible representation of $\mathrm{SL}(n,\mathbb{C})$. Finally we find, with proof, the normal form of the matrix that appears in the enumeration of domino tilings of an Aztec diamond.

1 Introduction

The permanent-determinant and Hafnian-Pfaffian methods of Kasteleyn and Percus give determinant and Pfaffian expressions for the number of perfect matchings of a planar graph (11; 21). Although the methods originated in mathematical physics, they have enjoyed new attention in enumerative combinatorics in the past ten years (10; 15; 16;

*Supported by NSF grants DMS #9704125 and DMS #0072342, and by a Sloan Foundation Research Fellowship

12; 34), in particular for enumerating lozenge and domino tilings of various regions in the plane. These successes suggest looking at further properties of the matrices that the methods produce beyond just their determinants or Pfaffians.

In this article we investigate the cokernel, or equivalently the Smith normal form, of a Kasteleyn or Kasteleyn-Percus matrix M arising from a planar graph G. One theme of our general results in Sections 3.3 and 4.1 is that the cokernel is a canonical object that can be defined in several different ways. More generally for weighted enumerations we consider M up to the equivalence relation of general row and column operations. If G has at least one matching, then the set of matchings is equinumerous with $\operatorname{coker} M$. (In Section 4.2, we conjecture an interpretation of this fact in the spirit of a bijection.) The cokernel of M is also interesting even when the graph G has no matchings, a situation which we call an *impossible enumeration*. Propp proposed another invariant of M that generalizes to impossible enumerations and that was studied by Saldanha (22; 25), namely the spectrum of M^*M.

The idea of computing cokernels as a refinement of enumeration also arose in the context of Kirchoff's determinant formula for the number of spanning trees of a connected graph. In this context the cokernels are called *tree groups* and they were proposed independently by Biggs, Lorenzini, and Merris (2; 18; 19). Indeed, Kenyon, Propp, and Wilson (13), generalizing an idea due to Fisher (7), found a bijection between spanning trees of a certain type of planar graph G and the perfect matchings of another planar graph G'. We conjecture that the tree group of G is isomorphic to the Kasteleyn-Percus cokernel of G'.

In Section 5.1 we study cokernels for the special case of enumeration of plane partitions in a box, as well as related lozenge tilings. We previously asked what is the cokernel of a Carlitz matrix, which is equivalent to the Kasteleyn-Percus matrix for plane partitions in a box with no symmetry imposed (22). We give two conjectures that together imply an answer. Finally in Section 5.3 we derive, with proof, the cokernel for the enumeration of domino tilings of an Aztec diamond.

Acknowledgments

The author would like to thank Torsten Ekedahl, Christian Krattenthaler, and Martin Loebl for helpful discussions. The author would especially like to thank Jim Propp for his diligent interest in this work.

2 Preliminaries

2.1 Graph conventions

In general by a *planar* graph we mean a graph embedded in the sphere S^2. We mark one point of S^2 outside of the graph as the infinite point; the face containing it is the infinite face. Our graphs may have both self-loops and multiple edges, although self-loops cannot participate in matchings.

2.2 Matrix algebra

Let R be a commutative ring with unit. We consider matrices M over R, not necessarily square, up to three kinds of equivalence: general row operations,

$$M \mapsto AM$$

with A invertible; general column operations,

$$M \mapsto MA$$

with A invertible; and stabilization,

$$M \mapsto \left(\begin{array}{c|c} 1 & 0 \\ \hline 0 & M \end{array}\right)$$

and its inverse. Any matrix M' which is equivalent to M under these operations is a *stably equivalent form* of M.

A matrix A over R is *alternating* if it is antisymmetric and has null diagonal. (Antisymmetric implies alternating unless 2 is a zero divisor in R.) We consider alternating matrices up to two kinds of equivalence: general symmetric operations,

$$A \mapsto BAB^T$$

with B invertible; and stabilization,

$$A \mapsto \left(\begin{array}{cc|c} 0 & 1 & \multirow{2}{*}{0} \\ -1 & 0 & \\ \hline \multicolumn{2}{c|}{0} & M \end{array}\right)$$

and its inverse. A matrix A' which is equivalent to A is also called a *stably equivalent form* of A.

As a special case of these notions, *elementary row operation* on a matrix M consists of either multiplying some row i by a unit in R, or adding some multiple of some row i to row $j \neq i$. *Elementary column operations* are defined likewise. We define a *pivot* on a matrix M at the (i,j) position as subtracting $M_{k,j}/M_{i,j}$ times row i from row k for all $k \neq i$, then subtracting $M_{i,k}/M_{i,j}$ times column j from column k for all $j \neq k$. This operation is possible when $M_{i,j}$ divides every entry in the same row and column. In matrix notation, if $M_{1,1} = 1$, the pivot at $(1,1)$ looks like this:

$$M = \left(\begin{array}{c|c} 1 & Y^T \\ \hline X & M' \end{array}\right) \mapsto \left(\begin{array}{c|c} 1 & 0 \\ \hline 0 & M' - XY^T \end{array}\right).$$

A *deleted pivot* consists of a pivot at (i,j) followed by deleting row i and column j from the matrix. The deleted pivot at $(1,1)$ on our example M looks like this:

$$M = \left(\begin{array}{c|c} 1 & Y^T \\ \hline X & M' \end{array}\right) \mapsto M' - XY^T.$$

If A is an alternating matrix, we define an *elementary symmetric operation* as an elementary row operation followed by the same operation in transpose on columns. We can similarly define a *symmetric pivot* and a *symmetric deleted pivot.* All of these operations are special cases of general symmetric matrix operations, and therefore preserve the alternating property.

If R is a principal ideal domain (PID), then an $n \times k$ matrix M is equivalent to one called its *Smith normal form* and denoted $\mathrm{Sm}(M)$. We define $\mathrm{Sm}(M)$ and prove its existence in Section 6. Note that if M' is a stabilization of M, then $\mathrm{Sm}(M')$ is a stabilization of $\mathrm{Sm}(M)$.

If R is arbitrary, then we can interpret M as a homomorphism from R^k to R^n. In this interpretation M has a kernel $\ker M$, an image $\operatorname{im} M$, and a cokernel

$$\operatorname{coker} M = R^n / \operatorname{im} M.$$

If R is a PID, the cokernel carries the same information as the Smith normal form. Over a general ring R, only very special matrices admit a Smith normal form. Determining equivalence of those that do not is much more complicated than for those that do. In particular inequivalent matrices may have the same cokernel. However, over any ring R the cokernel is invariant under stable equivalence and it does determine the determinant $\det M$ up to a unit factor. A special motivation for considering cokernels appears when $R = \mathbb{Z}$ and M is square. In this case the absolute determinant (*i.e.*, absolute value of the determinant) is the number of elements in the cokernel,

$$|\det M| = |\operatorname{coker} M|,$$

when the cokernel is finite, while

$$\det M = 0$$

if the cokernel is infinite.

An alternating matrix A over a PID is also equivalent to its antisymmetric Smith normal form $\mathrm{Sm}_a(A)$, which we also discuss in Section 6. Again $\operatorname{coker} A$ determines $\mathrm{Sm}_a(A)$.

Remark. If M is a matrix over the polynomial ring $\mathbb{F}[x]$ over an algebraically closed field $\mathbb{F}$, which is a PID, then the factor exhaustion method for computing $\det M$ (14) actually computes the Smith normal form (or cokernel) of M. Thus the Smith normal form plays a hidden role in a computational method which is widely used in enumerative combinatorics.

The basic version of the factor exhaustion method computes the rank of the reduction

$$M \otimes \mathbb{F}[x]/(x-r)$$

for all $r \in \mathbb{F}$. These ranks determine $\det M$ up to a constant factor if the Smith normal form of M is square free. It is tempting to conclude that the factor exhaustion method "fails" if the Smith normal form is not square free. But sometimes one can compute the cokernel of

$$M \otimes \mathbb{F}[x]/(x-r)^k,$$

for all r and k. This information determines $\operatorname{coker} M$, as well as $\det M$ up to a constant factor, regardless of its structure. Thus the factor exhaustion method always succeeds in principle.

3 Counting matchings

Most of this section is a review of Reference 16.

3.1 Kasteleyn and Percus

Let G be a connected finite graph. If we orient the edges of G, then we define the *alternating adjacency matrix* A of G by letting $A_{i,j}$ be the number of edges from vertex i to vertex j minus the number of edges from vertex j to vertex i. If G is simple, then the Pfaffian $\operatorname{Pf} A$ has one non-zero term for every perfect matching of G, but in general the terms may not have the same sign.

Theorem 1 (Kasteleyn). *If G is a simple, planar graph, then it admits an orientation such that all terms in* $\operatorname{Pf} A$ *have the same sign, where A is the alternating adjacency matrix of G (11).*

In general an orientation of G such that all terms in $\operatorname{Pf} A$ have the same sign is called a *Pfaffian orientation* of G. If an orientation of G is Pfaffian, then the absolute Pfaffian $|\operatorname{Pf} A|$ is the number of perfect matchings of G. Kasteleyn's rule for a Pfaffian orientation is that an odd number of edges of each (finite) face of G should point clockwise. We call such an orientation *Kasteleyn flat* and the resulting matrix A a *Kasteleyn matrix* for the graph G. Likewise an orientation may be Kasteleyn flat at a particular face if it satisfies Kasteleyn's rule at that face. Every planar graph has a Kasteleyn-flat orientation, although it is only flat at the infinite face of G if G has an even number of vertices. Forming a Kasteleyn matrix to count matchings of a planar graph is also called the *Hafnian-Pfaffian method* (15).

Percus (21) found a simplification of the Hafnian-Pfaffian method when G is bipartite. Suppose that G is a bipartite graph with the vertices colored black and white, and suppose that each edge has a sign $+$ or $-$, interpreted as the weight 1 or -1. Then we define the *bipartite adjacency matrix* M of G by letting $M_{i,j}$ be the total weight of all edges from the black vertex i to the white vertex j. If G is simple, then the determinant $\det M$ has a non-zero term for each perfect matching of G, but in general with both signs.

Theorem 2 (Percus). *If G is a simple, planar, bipartite graph, then it admits a sign decoration such that all terms in* $\det M$ *have the same sign, where M is the bipartite adjacency matrix of G.*

In the rule given by Percus, the edges of each face of G should have an odd number of $-$ signs if and only if the face has $4k$ sides. We call such a sign decoration of G *Kasteleyn flat* and the corresponding matrix M a *Kasteleyn-Percus matrix.* Every planar graph has a Kasteleyn-flat signing, although it is only flat at the infinite face if G has an even number of vertices. Forming a Kasteleyn-Percus matrix M is also called the *permanent-determinant method*. A Kasteleyn matrix A for a bipartite graph G can be viewed as two copies of a Kasteleyn-Percus matrix M:

$$A = \left(\begin{array}{c|c} 0 & M \\ \hline -M & 0 \end{array} \right).$$

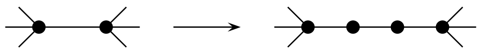

Figure 1: Tripling an edge in a graph.

If the graph G is not simple, then we may make it simple by tripling edges, as shown in Figure 1. The set of matchings of the new graph G' is naturally bijective with the set of matchings of G. A more economical approach is to define a Kasteleyn matrix A or a Kasteleyn-Percus matrix M for G directly. In this case A_{ij} is the number of edges from vertex i to vertex j minus the number from j to i, while M_{ij} is the number of positive edges minus the number of negative edges connecting i and j.

A variant of the Hafnian-Pfaffian method applies to a projectively planar graph G which is *locally but not globally bipartite*. This means that G is embedded in the projective plane and that all faces have an even number of sides, but that G is not bipartite. An equivalent condition is that all contractible cycles in G have even length and all non-contractible cycles have odd length.

Theorem 3. *If a projectively planar graph G is locally but not globally bipartite, then it admits a Pfaffian orientation. (16)*

The orientation constructed in the proof of Theorem 3 is one with the property that each face has an odd number of edges pointing in each direction. We call such an orientation Kasteleyn flat; it exists if G has an even number of vertices. (If G has an odd number of vertices, then every orientation is trivially Pfaffian.) We call the corresponding alternating adjacency matrix A the Kasteleyn matrix of G as usual.

The constructions of this section, in particular Theorems 1 and 2, generalize to weighted enumerations of the matchings of G, where each edge of G is assigned a weight and the weight of a matching is the product of the weights of its edges. We separately assign signs or orientations to G using the Kasteleyn rule (in the general case) or the Percus rule (in the bipartite case). The weighted alternating adjacency matrix A is called a Kasteleyn matrix of G. If G is bipartite, the weighted bipartite adjacency matrix M, with the weights multiplied by the signs, is a Kasteleyn-Percus matrix of G. Then $\operatorname{Pf} A$ or $\det M$ is the total weight of all matchings in G.

3.2 Polygamy and reflections

We can use the Hafnian-Pfaffian method to count certain generalized matchings among the vertices of a planar graph G using an idea originally due to Fisher (7). We arbitrarily divide the vertices of G into three types: Monogamous vertices, odd-polygamous vertices, and even-polygamous vertices. An *odd-polygamous vertex* is one that can be connected to any odd number of other vertices in a matching, while an *even-polygamous vertex* can be connected to any even number of other vertices (including none).

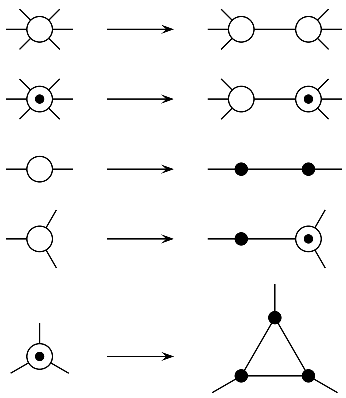

Figure 2: Resolving polygamy into monogamy.

If G is a graph with polygamous vertices, we can find a new graph G' such that the ordinary perfect matchings of G' are bijective with the generalized matchings of G. The graph G' defined from G using a series of local moves that are shown in Figure 2. (In this figure and later, an open circle is an even-polygamous vertex and a dotted circle is an odd-polygamous vertex.) We also describe the moves in words. First, if a polygamous vertex of G has valence greater than 3, we can split it into two polygamous vertices of lower valence with the same total parity. This leaves polygamous vertices of both parities of valence 1, 2, and 3. If a polygamous vertex v is even and has valence 1, we can delete it. If it is odd and has valence 1 or 2, it is the same as an ordinary monogamous vertex. If it is even and has valence 2, we can replace it with two monogamous vertices. If it is even and has valence 3, we can split it into an odd-polygamous divalent vertex and an odd-polygamous trivalent vertex. Finally, if it odd and has valence 3, we can replace it with a triangle. Each of these moves comes with an obvious bijection between the matchings before and after. Thus these moves establish the following:

Proposition 4 (Fisher). *Given a graph G with odd- and even-polygamous vertices, the polygamous vertices can be replaced by monogamous subgraphs so that the matchings of the new graph G' are bijective with those of G. If G is planar, then G' can be planar.*

We call the resulting graph G' a *monogamous resolution* of G. If G is planar, then G' admits Kasteleyn matrices, and we call any such matrix a Kasteleyn matrix of G as well.

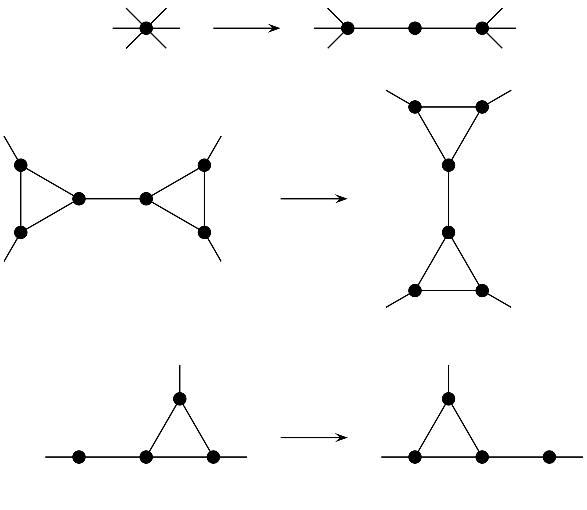

Figure 3: Moves on monogamous resolutions of a polygamous graph.

The monogamous resolution of a polygamous graph is far from unique. But we can consider moves that connect different monogamous resolutions of a polygamous graph. The moves are as shown in Figure 3: Doubly splitting a vertex, rotating a pair of triangles, and switching a triangle with an edge. Each of these moves comes with a bijection between the matchings of the two graphs that it connects.

Proposition 5. *Any two monogamous resolutions of a graph G are connected by the moves of vertex splitting and its inverse, switching triangles, and switching a triangle with an edge. The moves also connect any two planar resolutions of a planar graph G through intermediate planar resolutions.*

The proof of Proposition 5 is routine.

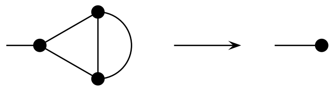

Figure 4: Removing a self-connected triangle.

Another interesting move is removing a self-connected triangle, as shown in Figure 4. This move induces a 2-to-1 map on the set of matchings before and after.

Polygamous matchings have two common applications. If a graph G is entirely polygamous, then we can denote the presence or absence of each edge by an element of $\mathbb{Z}/2$. Each vertex then imposes a linear constraint on the variables, so the number of matchings is therefore either 0 or 2^n for some n. The corresponding weighted enumerations are related to the Ising model (16; 33; 7). Another way to see that the number of matchings is a power of two is to use the moves in Figures 3 and 4 to reduce a monogamous resolution of G to a tree, which has at most one matching.

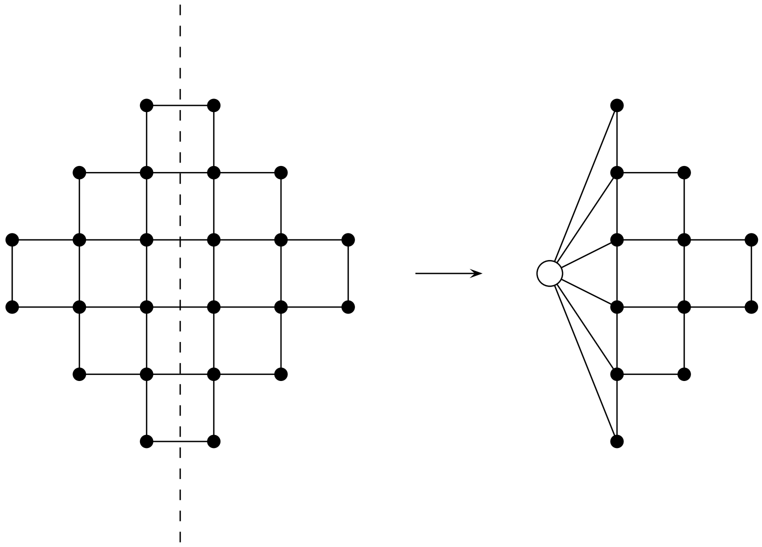

Figure 5: Using polygamy to count reflection-invariant matchings.

Another application is counting matchings invariant under reflections (15; 16). Suppose that a planar graph G has a reflection symmetry σ, and suppose that the line of reflection bisects some of the edges of G. Then the σ-invariant matchings of G are bijective with a modified quotient graph $G//\sigma$ in which the bisected edges are tied to a polygamous vertex, as in Figure 5. The parity of the polygamous vertex should be set so that the total parity (odd-polygamous plus monogamous vertices) is even. The same construction works if we divide G by any group acting on the sphere that includes reflections, since all of the reflective boundary can be reached by a single polygamous vertex.

3.3 Gessel-Viennot

The Gessel-Viennot method (9; 8) yields another determinant expression for a certain sum over the sets of disjoint paths in an acyclic, directed graph G. (Theorem 6 below, which is the basic result of the method, was independently found by Lindström (17). Gessel and Viennot were the first to use it for unweighted enumeration.) The graph G need not be planar. We label some of the vertices of G as left endpoints and some as right endpoints, and we separately order the left endpoints and the right endpoints. Let P be the set of collections of vertex-disjoint paths in G connecting the left endpoints to the right endpoints. If P is non-empty then there are the same number of left and right endpoints on left and right; if there are n of each we call the elements of P *disjoint n-paths*. The Gessel-Viennot matrix V is defined by setting $V_{i,j}$ to the number of paths in G from left endpoint i to right endpoint j.

Theorem 6 (Lindström, Gessel-Viennot). *Let G be a directed, acyclic, weighted graph with n ordered left endpoints and n ordered right endpoints. Let P be the set of disjoint n-paths in G connecting left to right. If V is the Gessel-Viennot matrix of G, then*

$$\det V = \sum_{\ell \in P} w(\ell)(-1)^{\ell}. \tag{1}$$

Here $(-1)^{\ell}$ is the sign of the bijection from the left to the right endpoints induced by the paths in the collection ℓ, and $w(\ell)$ is the product of the weights of the edges of G that appear in ℓ.

Proof. We outline a non-traditional proof that will be useful later. We first suppose that the left endpoints are the sources in G (the vertices with in-degree 0) and the right endpoints are the sinks (the vertices with out-degree 0). We argue by induction on the number of *transit vertices*, meaning vertices that are neither sources nor sinks.

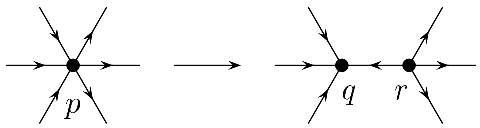

Figure 6: Splitting a transit vertex.

If G has no transit vertices, every path in G has length one. Consequently the n-paths in G are perfect matchings, and equation (1) is equivalent to the definition of the determinant. Suppose then that p is a transit vertex in G. We form a new graph G' by splitting p into two vertices q and r, with q a sink and r a source, as shown in Figure 6. We number q and r as the $n+1$st (last) source and sink in G'. We give the new edge between q and r a weight of -1. There is a natural bijection between disjoint n-paths ℓ in G and disjoint $n+1$-paths ℓ' in G': Every path in ℓ which avoids p is included in ℓ'. If some path in ℓ meets p, we break it into two paths ending at q and starting again at r. If ℓ is disjoint from p, we include the edge from r to q in ℓ'. In order to argue that the right side of equation (1) are the same for G and G', we check that

$$(-1)^{\ell} w(\ell) = (-1)^{\ell'} w(\ell').$$

If ℓ avoids p, the two sides are immediately the same. If ℓ meets p, then $(-1)^{\ell}$ and $(-1)^{\ell'}$ have opposite sign and so do $w(\ell)$ and $w(\ell')$. The left side of equation (1) is also the same: If V and V' are the Gessel-Viennot matrices of G and G', V is obtained from V' by a deleted pivot at $(n+1, n+1)$.

Now suppose that the left and right endpoints do not coincide with the sources and sinks. If G has a left endpoint q which is not a source, then there is an edge e from a vertex p to the vertex q. Let G' be G with e removed and let V' be its Gessel-Viennot matrix. Since the edge e is not in any n-path in G, the graph G' has the same n-paths with the same weights. If p is the ith left endpoint, we can obtain V' from V by subtracting $w(i,j)$

times row i from j, where $w(i,j)$ is the total weight in G of all paths from i to j. These row operations do not change the determinant. The same argument applies if G has a right endpoint which is not a sink.

Finally if G has a right endpoint source or a left endpoint sink, then it has no n-paths and some row or column of V is 0. If G has a source or a sink which is not an endpoint, we can delete it without changing the Gessel-Viennot matrix V or the set of n-paths. $\square$

We call the graph G' constructed in our proof of Theorem 6 the *transit-free resolution of* G. The transit-free resolution is a connection between the Gessel-Viennot method and the permanent-determinant method:

Corollary 7. *Let G be a connected, planar, directed, acyclic graph with n left and right endpoints on the outside face. Suppose that the left endpoints are segregated from the right endpoints on this face. If the left endpoints are the sources and the right endpoints are the sinks, then the Gessel-Viennot matrix V of G is obtained from a Kasteleyn-Percus matrix M of the transit-free resolution G' of G by deleted pivots. If not every left endpoint is a source or not every right endpoint is a sink, V is obtained by deleted pivots and other matrix operations.*

Note that by construction the matchings of the transit-free resolution G' of G are bijective with the n-paths in G. The planarity of G together with the position of its endpoints imply that all n-paths induce the same bijection and therefore have the same sign.

Proof. The Gessel-Viennot matrix V' of the graph G' is obtained from V by the stated operations. Thus it suffices to show that G' is planar and that V' is also a Kasteleyn-Percus matrix of G'.

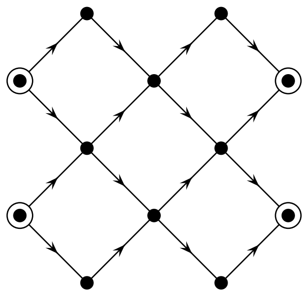

Figure 7: Left-to-right orientation implied by segregation of sources and sinks (circled).

We first establish that the orientation of G is qualitatively like that of the example in Figure 7: the orientations all point from left to right. More precisely, the edges incident to each transit vertex are segregated, in the sense that all incoming edges are adjacent and all outgoing edges are adjacent. The edges of each internal face are also segregated, in the sense that the clockwise edges are adjacent and the counterclockwise edges are adjacent. To prove that G has this structure, we reallocate the Euler characteristic of the sphere, 2, expressed as a sum over elements of G. In this sum, each vertex and face has Euler characteristic 1 and each edge has Euler characteristic -1. If a pair of edges shares both a vertex v and a face f, we deduct $\frac{1}{2}$ from the Euler characteristic of f if the edges both point to or both point from v, and otherwise we deduct $\frac{1}{2}$ from v. Since each edge participates in 4 such pairs, these deductions absorb the total Euler characteristic of all edges.

The reallocated characteristic of a vertex is 1 if it is a source or sink, 0 if it is a segregated transit vertex, and negative otherwise. The reallocated characteristic of a face is at most $2 - 2n$ if it is the outside face (since orientations must switch between clockwise and counterclockwise both at the sources and sinks and between them), 0 if it is a segregated internal face, and negative otherwise. (No face has positive reallocated characteristic since G is acyclic.) Thus the only way that the total can be 2 is if all internal faces and all transit vertices are segregated.

That the transit vertices of G are segregated implies that G' is planar. That each internal face f is segregated implies that if f has k sides, the corresponding face f' of G' has $2k-2$ sides. Moreover the $k-2$ new edges of f' have weight -1 in the proof of Theorem 6, which agrees with the Kasteleyn-Percus rule. Thus V' is a Kasteleyn-Percus matrix of G', as desired. $\square$

Finally we have not discussed a Pfaffian version of the Gessel-Viennot method defined by Stembridge (29). We believe that this method can be generalized further, and that it admits an analogue of Corollary 7.

4 Matchings and Smith normal form

4.1 Equivalences of Kasteleyn and Kasteleyn-Percus matrices

If M is a Kasteleyn-Percus matrix of a bipartite, planar graph G, then we can consider its cokernel, which by Section 2.2 is equinumerous with the number of matchings of M if it has any matchings. Furthermore, if $\operatorname{coker} M$ is infinite or if M isn't square, we can think of $\operatorname{coker} M$ as a way to "count" matchings in a graph that has none. We call such a computation an *impossible enumeration*. Both observations are reasons to study $\operatorname{coker} M$ as part of enumerative combinatorics.

If G is weighted by elements of some ring R, then we can consider M up to stable equivalence, whether or not it has a Smith normal form.

Suppose that M and M' are two Kasteleyn-Percus matrices for the same planar graph G. Then the signs on G given by M' and M differ by a 1-cocycle c with coefficients in the group $\{+,-\}$ (16). Since the sphere has no first homology, $c = \delta d$, where d is a 0-cochain. More explicitly d is a function from the vertices of G to $\{+,-\}$. We can use d to form two diagonal matrices A and B with diagonal entries ± 1 and such that $M' = AMB$. Evidently A and B are invertible over $\mathbb{Z}$, so M' and M have the same cokernel. In conclusion:

Proposition 8. *If G be a weighted bipartite planar graph, then all of its Kasteleyn-Percus matrices M are stably equivalent forms. In particular* $\operatorname{coker} M$ *is an invariant of G.*

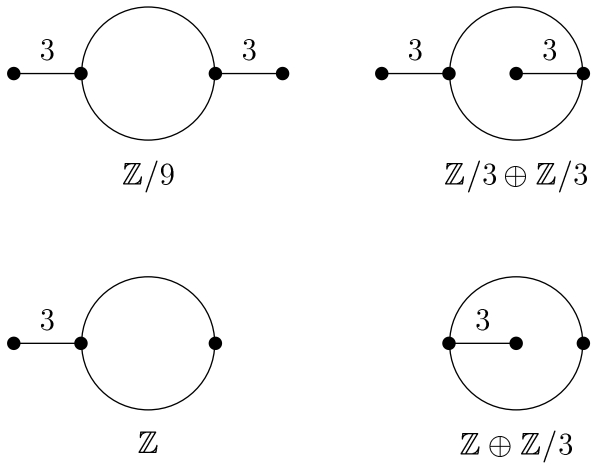

Figure 8: Embedding-dependent Kasteleyn-Percus cokernels.

Example 9. *The Smith normal form or cokernel of M can depend on the embedding of G in the plane. The top two graphs in Figure 8 have cokernels $\mathbb{Z}/9$ and $\mathbb{Z}/3 \oplus \mathbb{Z}/3$ in the two embeddings shown. If G has an odd number of vertices, then it cannot be Kasteleyn flat on its outside face. In this case changing which face is on the outside can change the cokernel as well. The bottom two graphs in Figure 8 are an example.*

Our analysis generalizes to the non-bipartite case. If G is a planar graph with a Kasteleyn matrix A, then we can consider A up to equivalence.

Again all Kasteleyn matrices we choose for the planar graph G are equivalent, because any two Kasteleyn-flat orientations of G differ by the coboundary of a 0-cochain on G with values in $\{+,-\}$. The matrices are consequently equivalent under the transformation

$$A \mapsto B^T A B$$

for some diagonal matrix B whose non-zero entries are ± 1. We can also pass from the usual clockwise-odd Kasteleyn rule to the counterclockwise-odd rule by negating A. We have no reason to believe that A and $-A$ are equivalent over a general ground ring R, but they do have the same cokernel and are therefore equivalent if R is a PID.

If G is projectively planar and locally but not globally bipartite, the argument is slightly different while the conclusion is the same. In this case the cohomology group

$$H^1(\mathbb{R}P^2, \mathbb{Z}/2) \cong \mathbb{Z}/2$$

is non-trivial. Suppose that we have two Kasteleyn-flat orientations of G whose matrices are A and A'. Their discrepancy is a 1-cocycle c which could represent either the trivial or the non-trivial class in $H^1(\mathbb{R}P^2, \mathbb{Z}/2)$. If c is trivial, then

$$A' = B^T A B$$

for some diagonal B. I.e., A and A' are equivalent. If c is non-trivial, then

$$A' = -B^T A B,$$

i.e., A' is equivalent to $-A$.

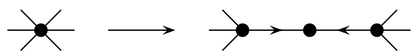

Figure 9: Preserving Kasteleyn flatness.

Kasteleyn and Kasteleyn-Percus matrices remain equivalent under more operations than just the choices of signs or orientations. In particular they remain equivalent under the moves in Section 3.2. In each move we make a graph G' from the graph G, and we need to choose related Kasteleyn-flat orientations of both graphs. For example consider a double vertex splitting. If G is Kasteleyn-flat, and if we orient the two new edges in the splitting in opposite directions as in Figure 9, then G' is also Kasteleyn flat. If the three vertices are numbered 1, 2, and 3, then the matrix A' of G' has a submatrix of the form

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

If we perform a deleted pivot at $(1,2)$ and $(2,1)$, we reduce A' to the matrix A of G.

4.2 Is it bijective?

Whenever two sets are known to have the same size, a traditional question in combinatorics is whether or not there is a bijection between them. In this section we conjecture a relationship between cokernels and matching sets which is similar to a bijection.

If M is a Kasteleyn-Percus matrix for an unweighted bipartite, planar graph G with at least one perfect matching, then two such sets to consider are $\operatorname{coker} M$ and P, the set of perfect matchings of G. To this pair we must add a third set, $\operatorname{coker} M^T$, since the choice between M and M^T is arbitrary. As explained in Section 6, $\operatorname{coker} M$ and $\operatorname{coker} M^T$ are isomorphic, but there is no canonical isomorphism. This is evidence against a natural bijection between $\operatorname{coker} M$ and $\operatorname{coker} M^T$, and therefore a natural bijection between either of them and P. On the other hand, the special planar structure of M might yield such bijections.

It may be better to consider quantum bijections or linearized bijections instead of traditional ones. If A and B are two finite sets, a *quantum bijection* is a unitary isomorphism

$$\mathbb{C}[A] \cong \mathbb{C}[B]$$

between the formal linear spans of A and B. A quantum bijection can be implemented by a quantum computer algorithm just as a traditional bijection can be implemented on a standard computer (28). A linear bijection is a linear isomorphism

$$\mathbb{F}[A] \cong \mathbb{F}[B],$$

not necessarily unitary, for some field $\mathbb{F}$. A linear bijection does not have the empirical computational interpretation that a traditional bijection or a quantum bijection does, but as a means of proving that A and B are equinumerous, it can be considered constructive.

If M is a non-singular $n \times n$ matrix over $\mathbb{Z}$, then there is a natural quantum bijection between $\operatorname{coker} M$ and $\operatorname{coker} M^T$, namely the discrete Fourier transform. (It is also a special case of Pontryagin duality (23)). We express it by defining a unitary matrix U whose rows are indexed by $x \in \operatorname{coker} M$ and whose columns are indexed by $y \in \operatorname{coker} M^T$. Given such x and y, we let X and Y be lifts in $\mathbb{Z}^n$. We then define

$$U_{x,y} = \frac{\exp(2\pi i Y^T M^{-1} X)}{\sqrt{|\det M|}}.$$

It is easy to check that $Y^T M^{-1} X$ changes by an integer if we change the lift X of x, because two such lifts differ by an element in $\operatorname{im} M$.

Given M, G, and P as above, it might be possible to factor the unitary map U into maps to and from $\mathbb{C}[P]$. However we may need to further relax the notion of a bijection. Sometimes when G is a finite group equinumerous with a finite set S, there is no natural bijection between them, but instead there is a natural, freely transitive group action of G on S. Having introduced quantum bijections, we can try to make $\mathbb{C}[P]$ a free unitary module over the group algebras $\mathbb{C}[\operatorname{coker} M]$ and $\mathbb{C}[\operatorname{coker} M^T]$. We can even ask that the two group actions be compatible with U by requiring the commutation relations

$$\alpha_y \alpha_x = \exp(2\pi i Y^T M^{-1} X) \alpha_x \alpha_y,$$

where $x \in \operatorname{coker} M$ and $y \in \operatorname{coker} M^T$ and α_x and α_y are their hypothetical actions on $\mathbb{C}[P]$. A standard theorem in representation theory says that for any M the algebra

$$D = \mathbb{C}[\operatorname{coker} M] \otimes \mathbb{C}[\operatorname{coker} M^T]$$

twisted by this commutation relation is isomorphic to a matrix algebra. This means that D has only one irreducible representation, and we conjecture that $\mathbb{C}[P]$ has the structure of this representation.

If A is a non-singular alternating matrix, then we can define a similar algebra D as a deformation of the group algebra $\mathbb{C}[\operatorname{coker} A]$. We let D be the formal complex span of elements α_x with $x \in \operatorname{coker} A$, and we arbitrarily order the elements in $\operatorname{coker} A$. For $x \leq y \in \operatorname{coker} A$, we define

$$\alpha_x \alpha_y = \alpha_{x+y},$$

and for arbitrary x and y, we impose the relation

$$\alpha_y \alpha_x = \exp(2\pi i Y^T A^{-1} X) \alpha_x \alpha_y.$$

The algebra D is again isomorphic to a matrix algebra, since Theorem 18 allows us to put A into

$$A = \left(\begin{array}{c|c} 0 & M \\ \hline -M & 0 \end{array} \right).$$

If we do so then algebra D then has the form previously described.

Conjecture 10. *If A is a Kasteleyn matrix of a planar graph G with non-empty matching set P, then there is a natural action of the algebra D on $\mathbb{C}[P]$, possibly depending on the way that G is embedded in the plane.*

5 Lozenge and domino tilings

5.1 Plane partitions and lozenge tilings

Plane partitions are an interesting source of enumerative planar matching problems. We consider the cokernels and Smith normal forms that arise in these problems, relying on the material in References 15 and 16.

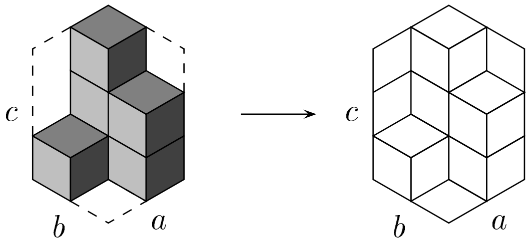

Figure 10: A plane partition and a tiling.

A plane partition in an $a \times b \times c$ box is equivalent to a lozenge tiling of an (a, b, c)-semiregular hexagon (Figure 10), which in turn is equivalent to a perfect matching of the dual hexagonal graph $Z(a, b, c)$ (Figure 11).

If a group G acts on the box and the plane partitions inside it, it also acts on the graph $Z(a, b, c)$. The G-invariant matchings correspond to matchings of a modified quotient graph $Z_G(a, b, c)$. To understand these graphs we recall the three generators of G:

- ρ, cyclic symmetry for plane partitions or rotation by 120 degrees for lozenge tilings, defined when $a = b = c$.

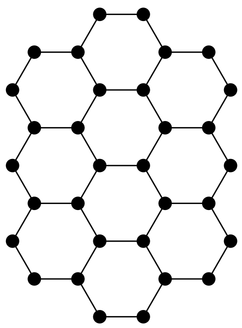

Figure 11: The graph $Z(2,2,3)$.

- τ, symmetry for plane partitions or diagonal reflection for lozenge tilings, defined when $b = c$.
- κ, complementation for plane partitions or rotation by 180 degrees for lozenge tilings.

We describe $Z_G(a,b,c)$ case by case:

- If $G = \langle\rho\rangle$ or $\langle\rho,\kappa\rangle$, or if a, b, and c are all even and G is $\langle\kappa\rangle$, then $Z_G(a,b,c)$ is the usual quotient graph $Z(a,b,c)/G$.
- If $G = \langle\kappa\tau\rangle$ or $G = \langle\kappa\tau,\rho\rangle$, then we delete the edges and vertices of $Z(a,b,c)$ along the lines of reflection and let $Z_G(a,b,c)$ be a connected component of the remainder.
- If $G = \langle\kappa\rangle$ and only one or two of a, b, and c is even, then $Z(a,b,c)$ has a central edge e invariant under κ. If one dimension is even, we define $Z_\kappa(a,b,c)$ by deleting e and its vertices and then quotienting by κ. If two dimensions are even, we define $Z_\kappa(a,b,c)$ by deleting e but not its vertices, and then quotienting by κ.
- If $G = \langle\tau\rangle$, $\langle\tau,\kappa\rangle$, or $\langle\tau,\rho,\kappa\rangle$, we start with $Z_H(a,b,c)$, where H is, respectively, $\langle 1\rangle$, $\langle\tau\kappa\rangle$, or $\langle\rho,\tau\kappa\rangle$. We cut $Z_H(a,b,c)$ by the line of reflection, and tie the cut edges of one region to a polygamous vertex to define $Z_G(a,b,c)$.
- If $G = \langle\tau,\rho\rangle$, we cut $Z(a,a,a)$ by three lines of reflection and tie the cut edges of one region (lying along two of the three lines) to a polygamous vertex to form $Z_G(a,a,a)$.

For every symmetry group G, the number of G-invariant plane partitions $N_G(a,b,c)$ is *round*, meaning a product of small factors. (This is not a completely rigorous notion. In this case $N_G(a,b,c)$ grows exponentially in $ab + ac + bc$ while its prime factors grow linearly in $a+b+c$.) In addition if G is a subgroup of $\langle\rho,\tau\rangle$, the G-invariant partitions have round q-enumerations, where the q-weight of a plane partition is sometimes the number of cubes and sometimes the number of G-orbits of cubes. (To say that a polynomial $P(q)$ is round means not only that it is a product of small factors, but also that the factors are cyclotomic. Equivalently $P(q)$ is a ratio of products of differences of monomials.)

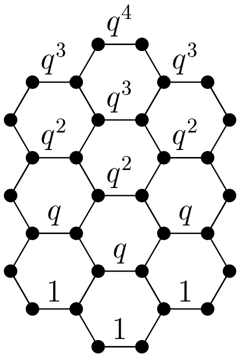

Figure 12: The graph $Z(2,2,3)$ weighted for q-enumeration.

All of these enumerations are proven (26; 15; 1; 31) except for the conjectured orbit q-enumeration of totally symmetric plane partitions. Both orbit and cube q-enumerations can be realized by weighted enumerations of matchings in $Z_G(a,b,c)$ (15; 16). Figure 12 shows an example where G is trivial (so that there is no distinction between cube and orbit q-enumeration in this case). Also the weight of a matching in the example only agrees with the q-weight of the corresponding plane partition up to a (matching-independent) factor of q. We will therefore consider q-enumerations over the ring $\mathbb{Z}[q,q^{-1}]$ and absorb powers of q in normalization.

We let $Z_G(a,b,c;q)$ be $Z_G(a,b,c)$ with weights chosen for cube q-enumeration, and we let $\tilde{Z}_G(a,b,c;q)$ be $Z_G(a,b,c)$ with weights chosen for orbit q-enumeration.

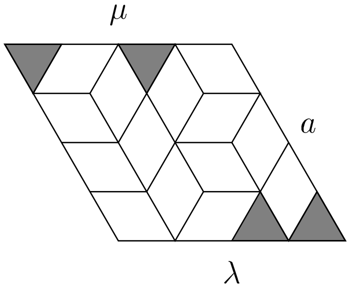

Figure 13: A lozenge tiling of a region with $a = 4$, $\lambda = (2,2)$, and $\mu = (1)$.

The problem of counting matchings in $Z(a,b,c)$ has two other interesting generalizations. The lozenge tilings of a parallelogram strip with notches on both sides (see Figure 13) are naturally bijective with the semi-standard skew tableaux of shape λ/μ and parts bounded by some a (27, §7.10), where λ and μ are two partitions with λ containing μ. If λ has b nonzero parts, the parallelogram then has dimensions a by $\lambda_1 + b$. On the bottom row it has notches at positions $\lambda_i + b + 1 - i$, counting from the left. On the top row it has notches at positions $\mu_i + b + 1 - i$, counting from the left and extending μ by 0 so that it also has b parts. We leave the proof of the bijection between tilings and

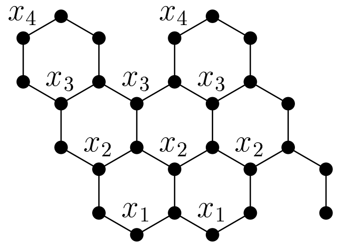

Figure 14: Weights of $Z(\lambda/\mu;\vec{x})$.

tableaux to the reader since it is similar to existing arguments in the literature. We let $Z(\lambda/\mu, a)$ be the graph dual to the tiling of this region by triangles, so that the lozenge tilings correspond to the matchings of $Z(\lambda/\mu, a)$. In an important weighted enumeration we assign the weight x_i to the northeast-pointing edges in the ith row, as in Figure 14. Call the weighted graph $Z(\lambda/\mu;\vec{x})$ The total weight of its matchings is the skew Schur function $s_{\lambda/\mu}(\vec{x})$ (in finitely many variables). If we let

$$\vec{q_a} = (1, q, \ldots, q^{a-1}),$$

then the specialization $s_{\lambda/\mu}(\vec{q_a})$ is the standard q-enumeration of skew tableaux. In particular, if we omit μ (by setting it to the empty partition), then $s_\lambda(\vec{x})$ is the character of an irreducible representation $V(\lambda)$ of $\mathrm{GL}(a, \mathbb{C})$, while $s_\lambda(\vec{q_a})$ is round by the q-Weyl dimension formula.

A second generalization is to count tilings of a semiregular hexagon with side lengths $a, b+d, c, a+d, b, c+d$ and with a triangle of size d removed (3). We do not know of a q-enumeration of these, although the symmetry $\kappa\tau$ (corresponding to TCPPs) appears when $b = c$ and the symmetry ρ (corresponding to CSPPs) appears when $a = b = c$. We let $Z(a, b, c, d, d)$ be the dual graph and we analogously define $Z_G(a, b, c, d, d)$ if G is a subgroup of $\langle \kappa\tau, \rho\rangle$.

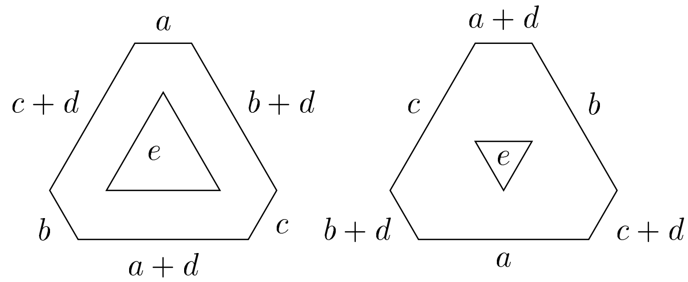

Figure 15: Hexagonal regions without central triangles.

As we mentioned in Section 4, we can consider the Kasteleyn or Kasteleyn-Percus cokernel for impossible enumerations where there are no matchings. This allows us to vary the above graphs in several ways, which we describe case by case:

- We can consider $Z_\kappa(a,b,c)$ when all three of a, b, and c are odd, as well as $Z_{\langle\kappa,\rho\rangle}(a,a,a)$ when a is odd.
- If one of a, b, and c is odd, we define the graph $Z'_\kappa(a,b,c)$ by removing the central edge e but not its vertices before quotienting by κ. If two are odd we define $Z'_\kappa(a,b,c)$ by removing e and its vertices before quotienting by κ.
- If $Z_G(a,b,c)$ has a polygamous vertex, we can give it the wrong parity to make $Z'_G(a,b,c)$. If $G=\langle\tau\rangle$ or $G=\langle\rho,\tau\rangle$, then we also define the q-weighted forms $Z'_G(a,b,c;q)$ and $\tilde{Z}'_G(a,b,c;q)$.
- We define the graph $Z(a,b,c,d,e)$ by removing a triangle of size e from a semiregular hexagon of side lengths $a, b+d, c, a+d, b, c+d$, with $d\neq e$, as shown in Figure 15. The parameters d and e might even have opposite sign, which we interpret by turning the triangle upside-down, as also indicated in Figure 15.

For each of the graphs defined in this section, we denote a corresponding Kasteleyn-Percus matrix by replacing Z by M if it is bipartite and monogamous, and the corresponding Kasteleyn matrix by replacing Z by A otherwise.

Conjecture 11. *Each of the matrices $M(a,b,c;q)$, $M_\rho(a,a,a;q)$, $A_\tau(a,b,b;q)$, $\tilde{A}_\tau(a,b,b;q)$, $A'_\tau(a,a,b;q)$, $\tilde{A}'_\tau(a,a,b;q)$, $\tilde{A}_{\langle\rho,\tau\rangle}(a,a,a;q)$, $\tilde{A}'_{\langle\rho,\tau\rangle}(a,a,a;q)$, and $M(\mu;\vec{q}_a)$ admits a Smith normal form over $\mathbb{Z}[q,q^{-1}]$, and the entries are q-round. The Smith normal form over $\mathbb{Z}$ of each of the matrices $M(a,b,c,d,e)$, $M_{\kappa\tau}(2a,b,b)$, $M_{\langle\rho,\kappa\tau\rangle}(2a,2a,2a)$, $A_\kappa(a,b,c)$, $A'_\kappa(a,b,c)$, $A_{\langle\kappa,\tau\rangle}(2a,b,b)$, $A'_{\langle\kappa,\tau\rangle}(2a,b,b)$, $A_{\langle\rho,\kappa\rangle}(a,a,a)$, $A_{\langle\rho,\kappa,\tau\rangle}(2a,2a,2a)$, and $A'_{\langle\rho,\kappa,\tau\rangle}(2a,2a,2a)$ has round entries.*

We could have stated conjecture 11 in greater generality by combining more of the variations above. For example the graph $Z(a,a,a,d,e)$ has symmetries and we could consider the matrix of a suitably modified quotient graph. Indeed we do not know how to state Conjecture 11 in full generality, since there are yet other variations of counting lozenge tilings in a hexagon with round enumerations (4). Presumably many or all of these variations also have impossible counterparts, and some may have q-enumerations as well. Also the spirit of the conjecture is to find the Smith normal forms or the cokernels explicitly.

Conjecture 12. *The Smith normal forms over $\mathbb{Z}[q,q^{-1}]$ of each of the matrices $M(a,b,c;q)$, $M_\rho(a,a,a;q)$, $A_\tau(a,b,b;q)$, $\tilde{A}_\tau(a,b,b;q)$, and $\tilde{A}_{\langle\rho,\tau\rangle}(a,a,a;q)$ are square free.*

Conjecture 12 may also not be fully general, although we note that the Smith normal form of $M(\lambda;\vec{q}_a)$ is not always square free. Note that Conjectures 12 and 11 would together solve problem 5 in Propp's problem list (22), which asks for the Kasteleyn cokernel of the graph $Z(a,b,c)$. The reason is that the Smith normal form of any non-singular matrix

M is uniquely determined by $\det M$, provided that the normal form exists and is square free. For example the q-enumeration of plane partitions tells us that

$$\det M(2,2,2;q) = (2)_q^2(5)_q,$$

where

$$(n)_q = \frac{q^n - 1}{q - 1}.$$

So Conjectures 11 and 12 assert that $\mathrm{Sm}(M(2,2,2;q))$ exists and its nontrivial entries are $(2)_q$ and $(2)_q(5)_q$. If we specialize at $q = 1$, we obtain the correct prediction that

$$\operatorname{coker} M(2,2,2) \cong \mathbb{Z}/2 \oplus \mathbb{Z}/10.$$

Stembridge (30) noticed a family of relations, called the $q = -1$ phenomenon, between -1-enumeration of some symmetry classes of plane partitions and ordinary enumeration of other symmetry classes. We conjecture an extension of the phenomenon that is easier to state in terms of cokernels in some cases and Smith normal forms in others.

Conjecture 13. *If G is $\langle 1 \rangle$ or $\langle \rho \rangle$ and if $G' = \langle G, \kappa \rangle$, then*

$$\mathrm{Sm}(A_{G'}(a,b,c)) = \mathrm{Sm}(M_G(a,b,c)_{-1})^2.$$

If G is $\langle \tau \rangle$ or $\langle \rho, \tau \rangle$, and if G' is, respectively, $\langle \kappa\tau \rangle$ or $\langle \rho, \kappa\tau \rangle$, then

$$\operatorname{coker} A_G(a,b,c)_{-1} \cong \operatorname{coker} M_{G'}(a,b,c)^{\oplus 2}.$$

If G is $\langle \tau \rangle$ or $\langle \rho, \tau \rangle$ and if $G' = \langle G, \kappa \rangle$, then

$$\operatorname{coker} A_{G'}(a,b,c) \cong \operatorname{coker} A'_G(a,b,c)_{-1}.$$

5.2 Jacobi-Trudi matrices

One conclusion of our construction for skew tableaux is a novel determinant formula for skew Schur functions:

$$s_{\lambda/\mu}(\vec{x}) = \det M(\lambda/\mu;\vec{x}).$$

We can compare this to two other determinant formulas for skew Schur functions (27, §7.16). Let

$$J(\lambda/\mu;\vec{x})_{i,j} = h_{\lambda_i - \mu_j - i + j}(\vec{x}),$$

where $h_n(\vec{x})$ is the nth complete symmetric function of $\vec{x}$, and let

$$D(\lambda/\mu;\vec{x})_{i,j} = e_{\lambda'_i - \mu'_j - i + j}(\vec{x}),$$

where $e_n(\vec{x})$ is the nth elementary symmetric function of $\vec{x}$ and λ' is the partition conjugate to λ. Then the Jacobi-Trudi identity states that

$$s_{\lambda/\mu}(\vec{x}) = \det J(\lambda/\mu;\vec{x})$$

while the dual Jacobi-Trudi identity states that

$$s_{\lambda/\mu}(\vec{x}) = \det D(\lambda/\mu;\vec{x}).$$

Theorem 14. *The Jacobi-Trudi matrices $J(\lambda/\mu;\vec{x})$ and $D(\lambda/\mu;\vec{x})$ are stably equivalent over the ring $\mathbb{Z}[\vec{x}]$ to the matrix $M(\lambda/\mu;\vec{x})$.*

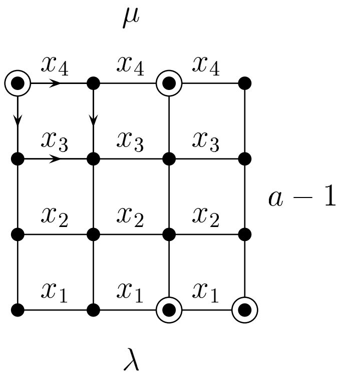

Figure 16: First Gessel-Viennot model of lozenge tilings. (Not all orientations are shown.)

Proof. The theorem is a special case of Corollary 7. Let $X(\lambda/\mu;\vec{x})$ be the graph exemplified in Figure 16: a square grid of height $a-1$ and width λ_1+b. On the top there is a left endpoint at each position $\mu_i+b+1-i$ and on the bottom a right endpoint at each position $\lambda_i+b+1-i$. Each horizontal edge points to the right and on the ith row has weight x_i. Each vertical edge points down. By a standard argument the Gessel-Viennot matrix of $X(\lambda/\mu;\vec{x})$ is $J(\lambda/\mu;\vec{x})$. At the same time, one can check that the transit-free resolution of $X(\lambda/\mu;\vec{x})$ is $Z(\lambda/\mu;\vec{x})$. Thus Corollary 7 says that $J(\lambda/\mu;\vec{x})$ is equivalent to $M(\lambda/\mu;\vec{x})$.

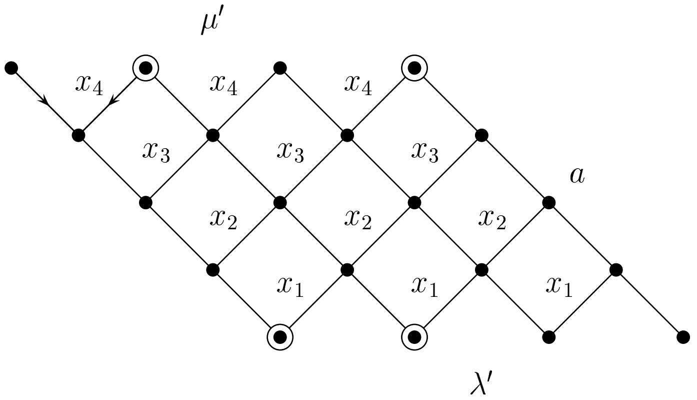

Figure 17: Second Gessel-Viennot model of lozenge tilings. (Not all orientations are shown.)

Let $Y(\lambda/\mu;\vec{x})$ be the graph exemplified in Figure 17. It is a diamond grid of height a and width λ_1+b. All edges point down, and those in the ith row that point southwest have weight x_i. On the top row there is a left endpoint at each position $b+i-\mu_i'$ and on the bottom a right endpoint at each position $b+i-\lambda_i'$. Note that the left endpoints are at all of the positions *not* of the form $\mu_i+b+1-i$ and likewise the right endpoints are at the positions *not* of the form One can check that the Gessel-Viennot matrix of $X(\lambda/\mu;\vec{x})$ is $D(\lambda/\mu;\vec{x})$. At the same time, noting the gaps between the marked endpoints, one can check that the transit-free resolution of $Y(\lambda/\mu;\vec{x})$ is $Z(\lambda/\mu;\vec{x})$. Thus Corollary 7 says that $J(\lambda/\mu;\vec{x})$ is equivalent to $D(\lambda/\mu;\vec{x})$. □

Question 15. *Are $J(\lambda/\mu;\vec{x})$ and $D(\lambda/\mu;\vec{x})$ stably equivalent over the ring of symmetric functions?*

Theorem 14 implies that there are several ways that equivalent forms of the matrix $M(\lambda/\mu;\vec{x})$ arise in several common guises. If we set $\vec{x}=\vec{q}_a$ and $\mu=0$, then Conjecture 11 asserts that $M(\lambda;\vec{q}_a)$ admits a Smith normal form a Smith normal form over the ring $\mathbb{Z}[q,q^{-1}]$. This suggests that coker $M(\lambda;\vec{q}_a)$ is an important extra structure that one can associate to the representation $V(\lambda)$ mentioned previously. Moreover the relationship between $M(\lambda)$ and $V(\lambda)$ is an interesting special case of Conjecture 10.

5.3 Domino tilings

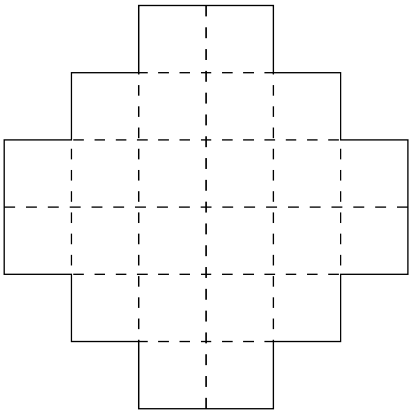

Figure 18: An Aztec diamond of order 3

Domino tilings of an Aztec diamond are a well-known analogue of lozenge tilings of a hexagon. Recall that an Aztec diamond of order n is the polyomino consisting of those unit squares lying entirely inside the region $|x|+|y|\le n+1$, as shown in Figure 18. A domino tiling of an Aztec diamond corresponds to a matching of the graph dual to the tiling by squares; an example of such a graph is the one on the left in Figure 5. Denote this graph by $Z_A(n)$ and let $M_A(n)$ be a Kasteleyn-Percus matrix for it.

Theorem 16. *If $M_A(n)$ is a Kasteleyn-Percus matrix for domino tilings of an Aztec diamond of order n, then*

$$\operatorname{coker} M_A(n) \cong \mathbb{Z}/2 \oplus \mathbb{Z}/4 \oplus \cdots \oplus \mathbb{Z}/2^n.$$

Theorem 16 extends the result that the number of domino tilings of an Aztec diamond of order n is $2^{n(n+1)/2}$ (20; 5; 6).

Proof. We give two arguments, one using Kasteleyn-Percus matrices and the other using Gessel-Viennot matrices. First, define a binomial coefficient matrix $B(n)$ by

$$B(n)_{0 \leq i,j < n} = \binom{i}{j}.$$

(We will assume that rows and columns of other matrices are numbered from 0 as well.) Define an $n \times n+1$ matrix $L(n)$ by putting the $n \times n$ identity matrix to the left of a null column, and define $R(n)$ by putting the $n \times n$ identity matrix to the right of a null column. For example here are $B(4)$, $L(3)$, and $R(3)$:

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 2 & 1 & 0 \\ 1 & 3 & 3 & 1 \end{pmatrix} \quad \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

One can check that the matrix

$$\begin{aligned} M_A(n) = {} & R(n) \otimes R(n)^T + L(n) \otimes R(n)^T \\ & + R(n) \otimes L(n)^T - L(n) \otimes L(n)^T \end{aligned}$$

is a Kasteleyn-Percus matrix for the graph $Z_A(n)$. At the same time,

$$B(n)L(n)B(n+1)^{-1} = L(n)$$

and

$$B(n)R(n)B(n+1)^{-1} = R(n) - L(n).$$

It follows that

$$\begin{aligned} M'(n) &\stackrel{\text{def}}{=} (B(n) \otimes B(n+1)^{-T})M(n)(B(n+1)^{-1} \otimes B(n)^T) \\ &= R(n) \otimes R(n)^T - 2L(n) \otimes L(n)^T. \end{aligned}$$

Since $B(n)$ is a triangular matrix, $M'(n)$ is equivalent to $M(n)$. On the other hand, one can check that, modulo permuting rows and columns,

$$\begin{aligned} M'(n) = {} & X(1) \oplus X(2) \oplus \cdots \oplus X(n) \\ & \oplus Y(1) \oplus Y(2) \oplus \cdots \oplus Y(n), \end{aligned}$$

where

$$X(k) = \begin{pmatrix} 1 & -2 & 0 & & 0 \\ 0 & 1 & -2 & \cdots & 0 \\ 0 & 0 & 1 & & 0 \\ & \vdots & & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}$$

and

$$Y(k) = \begin{pmatrix} -2 & 1 & 0 & & 0 \\ 0 & -2 & 1 & \cdots & 0 \\ 0 & 0 & -2 & & 0 \\ & \vdots & & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & -2 \end{pmatrix}.$$

Since

$$\operatorname{coker} X(k) = 0 \qquad \operatorname{coker} Y(k) = \mathbb{Z}/2^k,$$

the theorem follows.

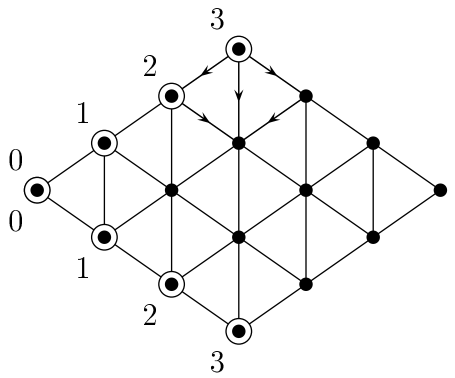

Figure 19: Gessel-Viennot model of an Aztec diamond. (Not all orientations are shown.)

In the Gessel-Viennot approach we first observe that if $G(n)$ is a graph of the type in Figure 19 then its transit-free resolution $G'(n)$ as described in Theorem 6 is the Aztec diamond graph $Z_A(n)$. (The left and right endpoints of $G(n)$ are numbered in the figure, with the left endpoints on top. The first left and right endpoints coincide, which is degenerate but allowed.) Let $V(n)$ be the Gessel-Viennot matrix of $G(n)$. The entry $V(n)_{i,j}$ is the Delannoy number $D(i,j)$ (27, §6.3) (see also Sachs and Zernitz (24)), because $G(n)$ matches the defining recurrence

$$\begin{aligned} D(i,j) = D(i,j-1) + D(i-1,j) + D(i-1,j-1) \\ D(0,i) = D(i,0) = 1. \end{aligned}$$

A standard formula for Delannoy numbers is

$$D(i,j) = \sum_k \binom{i}{k}\binom{j}{k} 2^k.$$

In matrix form this identity can be expressed

$$V(n) = B(n)V'(n)B(n)^T,$$

where $V'(n)$ is the Smith normal form of $V(n)$ with

$$V'(n)_{k,k} = 2^k.$$

Thus the Smith normal form of $V(n)$ also establishes the theorem. □

Finally, we could at least conjecturally extend Theorem 16 with the same variations as those we considered for lozenge tilings: Domino tilings with symmetry, q-enumerations, impossible enumerations, Aztec diamonds with teeth missing, etc. For example, Tokuyama (32) established a relation between generating functions of lozenge-type and Aztec-type Gelfand triangles. We conjecture that this identity can be extended to an equivalence of Kasteleyn-Percus (or Gessel-Viennot) matrices. We leave this and other possibilities to future work.

6 Appendix: Smith normal form

Theorem 17. *If M is a $k \times n$ matrix over a principal ideal domain R, then there exist invertible matrices A and B such that*

$$\mathrm{Sm}(M) = AMB$$

is diagonal and $\mathrm{Sm}(M)_{i,i}$ *divides* $\mathrm{Sm}(M)_{i+1,i+1}$. *The matrix* $\mathrm{Sm}(M)$ *is a* Smith normal form *of M.*

Evidently

$$\begin{aligned}\operatorname{coker} M &\cong \operatorname{coker} \mathrm{Sm}(M) \\ &= R^{k-n} \oplus R/\mathrm{Sm}(M)_{1,1} \oplus \cdots \oplus R/\mathrm{Sm}(M)_{n,n}\end{aligned}$$

if $k \geq n$ and

$$\begin{aligned}\operatorname{coker} M &\cong \operatorname{coker} \mathrm{Sm}(M) \\ &= R/\mathrm{Sm}(M)_{1,1} \oplus \cdots \oplus R/\mathrm{Sm}(M)_{k,k}\end{aligned}$$

otherwise. (If $R = \mathbb{Z}$ and we allow n, but not k, to be infinite, then Theorem 17 is equivalent to the classification of finitely generated abelian groups.) It is easy to show that this decomposition of $\operatorname{coker} M$ is unique and that $\mathrm{Sm}(M)$ is unique up to multiplying the entries by units.

Proof. We argue by induction on n or k. Since R is a PID, any two elements a and b have a greatest common divisor

$$\gcd(a,b) = ax + by$$

which is unique up to a unit factor. Moreover, R is Noetherian, which means that any chain of divisibilities $a_{i+1}|a_i$ must eventually be constant up to unit factors. Thus we can argue by induction with respect to divisibility of non-zero elements of R. Since we can replace M by any equivalent form, we assume that the entry $M_{1,1} = a$ is not divisible by any entry in any form of M, except for those that are a times a unit. We claim that a divides every entry of M. Otherwise there is an entry $M_{i,j} = b$ such that $a \nmid b$ and $b \nmid a$. If b is in the first row or column, then after permuting rows and columns and possibly transposing M, $M_{1,2} = b$:

$$M = \left(\begin{array}{cc|c} a & b & \cdots \\ \hline & \vdots & \ddots \end{array}\right).$$

Let $c = ax + by$ be a greatest common divisor of a and b. If we post-multiply M by the matrix

$$B = \left(\begin{array}{cc|c} x & a/c & \multirow{2}{*}{0} \\ y & -b/c & \\ \hline \multicolumn{2}{c|}{0} & I \end{array}\right),$$

the result is a form of M with the entry c, which contradicts the choice of a. If a divides every entry in the first row and column but not some entry b elsewhere, then we first pivot at the position $(1,1)$ to obtain:

$$M = \left(\begin{array}{cc|c} a & 0 & \cdots \\ 0 & b & \\ \hline & \vdots & \ddots \end{array}\right).$$

If we pre-multiply by

$$A = \left(\begin{array}{cc|c} 1 & 1 & \multirow{2}{*}{0} \\ 0 & 1 & \\ \hline \multicolumn{2}{c|}{0} & I \end{array}\right)$$

and post-multiply by B above, we again produce the entry c.

Given that $M_{1,1} = a$ does divide the rest of M we perform a deleted pivot at $(1,1)$ and assume normal form for the remaining submatrix by induction. $\square$

Theorem 18. *If A is an alternating $n \times n$ matrix over a principal ideal domain R, then there exists an invertible matrix B such that*

$$\mathrm{Sm}_a(A) = B^T A B$$

is block-diagonal with 2×2 blocks, and such that $\mathrm{Sm}_a(A)_{2i,2i+1}$ divides $\mathrm{Sm}_a(A)_{2i+2,2i+3}$. The matrix $\mathrm{Sm}_a(A)$ is the alternating Smith normal form *of A.*

Proof. The argument is the same as the one for Theorem 17. We assume that $A_{1,2} = a$ is minimal among all entries of all forms of A. If it does not divide some entry in the first two rows and columns, then after permuting rows and columns, that entry is $A_{1,3} = b$:

$$A = \left(\begin{array}{ccc|c} 0 & a & b & \\ -a & 0 & \cdot & \cdots \\ -b & \cdot & 0 & \\ \hline & \vdots & & \ddots \end{array}\right).$$

Let $c = ax + by$ be a common divisor and let

$$B = \left(\begin{array}{ccc|c} 1 & 0 & 0 & \\ 0 & x & a/c & 0 \\ 0 & y & -b/c & \\ \hline & 0 & & I \end{array}\right).$$

Then $B^T AB$ has the entry c, a contradiction. If a divides every entry in the first two rows and columns but not some other entry b, then we can perform a symmetric pivot at $(1, 2)$. After permuting rows and columns A then has the form:

$$A' = \left(\begin{array}{cccc|c} 0 & a & 0 & 0 & \\ -a & 0 & 0 & 0 & \cdots \\ 0 & 0 & 0 & b & \\ 0 & 0 & -b & 0 & \\ \hline & & \vdots & & \ddots \end{array}\right).$$

Let $c = ax + by$ be a common divisor and let

$$B = \left(\begin{array}{cccc|c} 1 & 0 & 0 & 0 & \\ 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & \\ 0 & 0 & 0 & 1 & \\ \hline & & 0 & & I \end{array}\right).$$

Then $B^T AB$ has the form of the previous case.

Finally if a divides every entry of A, we perform a deleted symmetric pivot at $(1, 2)$ and inductively assume the normal form for the remaining submatrix. □

References

[1] George E. Andrews, *Plane partitions V: The T. S. S. C. P. P. conjecture*, J. Combin. Theory Ser. A **66** (1994), no. 1, 28–39.

[2] Norman Biggs, *Algebraic potential theory on graphs*, Bull. London Math. Soc. **29** (1997), no. 6, 641–682.

[3] Mihai Ciucu, Theresia Eisenkölbl, Christian Krattenthaler, and Douglas Zare, *Enumeration of lozenge tilings of hexagons with a central triangular hole*, arXiv:math.CO/9912053.

[4] Mihai Ciucu and Christian Krattenthaler, *The number of centered lozenge tilings of a symmetric hexagon*, J. Combin. Theory Ser. A **86** (1999), no. 1, 103–126, arXiv:math.CO/9712209.

[5] Noam Elkies, Greg Kuperberg, Michael Larsen, and James Propp, *Alternating-sign matrices and domino tilings, I*, J. Algebraic Combin. **1** (1992), no. 2, 111–132.

[6] ———, *Alternating-sign matrices and domino tilings, II*, J. Algebraic Combin. **1** (1992), no. 3, 219–234.

[7] Michael E. Fisher, *On the dimer solution of planar Ising models*, J. Math. Phys. **7** (1966), 1776–1781.

[8] Ira Gessel and Xavier G. Viennot, *Determinants, paths, and plane partitions*, 1989, Preprint.

[9] Ira M. Gessel and Xavier G. Viennot, *Binomial determinants, paths, and hook length formulas*, Adv. Math. **58** (1985), no. 3, 300–321.

[10] William Jockusch, *Perfect matchings and perfect squares*, J. Combin. Theory Ser. A **67** (1994), 100–115.

[11] P. W. Kasteleyn, *Graph theory and crystal physics*, Graph theory and theoretical physics (F. Harary, ed.), Academic Press, 1967.

[12] Richard Kenyon, *Local statistics of lattice dimers*, Ann. Inst. H. Poincaré Probab. Statist. **33** (1997), 591–618, arXiv:math.CO/0105054.

[13] Richard Kenyon, James Propp, and David Wilson, *Trees and matchings*, Electron. J. Combin. **7** (2000), #R25, arXiv:math.CO/9903025.

[14] Christian Krattenthaler, *Advanced determinant calculus*, Séminaire Lotharingien Combin. **42** (1999), B42q, arXiv:math.CO/9902004.

[15] Greg Kuperberg, *Symmetries of plane partitions and the permanent-determinant method*, J. Combin. Theory Ser. A **68** (1994), no. 1, 115–151.

[16] ———, *An exploration of the permanent-determinant method*, Electron. J. Combin. **5** (1998), no. #R46, 16pp, arXiv:math.CO/9810091.

[17] B. Lindström, *On the vector representations of induced matroids*, Bull. London Math. Soc. **5** (1973), 85–90.

[18] Dino J. Lorenzini, *A finite group attached to the Laplacian of a graph*, Discrete Math. **91** (1991), no. 3, 277–282.

[19] Russell Merris, *Unimodular equivalence of graphs*, Linear Algebra Appl. **173** (1992), 181–189.

[20] William H. Mills, David P. Robbins, and Howard Rumsey, *Alternating-sign matrices and descending plane partitions*, J. Combin. Theory Ser. A **34** (1983), no. 3, 340–359.

[21] Jerome K. Percus, *One more technique for the dimer problem*, J. Math. Phys. **10** (1969), 1881–1888.

[22] James Propp, *Enumeration of matchings: problems and progress*, New perspectives in algebraic combinatorics (Berkeley, CA, 1996–97), Cambridge Univ. Press, Cambridge, 1999, arXiv:math.CO/9904150, pp. 255–291.

[23] Walter Rudin, *Fourier analysis on groups*, Wiley Interscience, 1962.

[24] Horst Sachs and Holger Zernitz, *Remark on the dimer problem*, Discrete Appl. Math. **51** (1994), no. 1-2, 171–179.

[25] Nicolau C. Saldanha, *Singular polynomials of generalized Kasteleyn matrices*, arXiv:math.CO/0101002.

[26] Richard P. Stanley, *Symmetries of plane partitions*, J. Combin. Theory Ser. A **43** (1986), no. 1, 103–113.

[27] ______, *Enumerative combinatorics*, vol. 2, Cambridge University Press, Cambridge, England, 1999.

[28] Andrew Steane, *Quantum computing*, Rep. Progr. Phys. **61** (1998), no. 2, 117–173, arXiv:quant-ph/9708022.

[29] John R. Stembridge, *Non-intersecting paths, Pfaffians, and plane partitions*, Adv. Math. **83** (1990), no. 1, 96–131.

[30] ______, *Some hidden relations involving the ten symmetry classes of plane partitions*, J. Combin. Theory Ser. A **68** (1994), 372–409.

[31] ______, *The enumeration of totally symmetric plane partitions*, Adv. Math. **111** (1995), no. 2, 227–243.

[32] Takeshi Tokuyama, *A generating function of strict Gel'fand patterns and some formulas on characters of general linear groups*, J. Math. Soc. Japan **40** (1988), 671–685.

[33] B. L. van der Waerden, *Die lange Reichweite der regelmaessigen Atomanordnung in Mischkristallen*, Zeitschrift für Physik **118** (1941), 473–488.

[34] Bo-Yin Yang, *Two enumeration problems about the Aztec diamonds*, Ph.D. thesis, Massachusetts Institute of Technology, 1991.