

Channels, Billiards, and Perfect Matching 2-Divisibility

Grant T. Barkley and Ricky Ini Liu

Abstract

Let m_G denote the number of perfect matchings of the graph G . We introduce a number of combinatorial tools for determining the parity of m_G and giving a lower bound on the power of 2 dividing m_G . In particular, we introduce certain vertex sets called channels, which correspond to elements in the kernel of the adjacency matrix of G modulo 2. A result of Lovász states that the existence of a nontrivial channel is equivalent to m_G being even. We give a new combinatorial proof of this result and strengthen it by showing that the number of channels gives a lower bound on the power of 2 dividing m_G when G is planar. We describe a number of local graph operations which preserve the number of channels. We also establish a surprising connection between 2-divisibility of m_G and dynamical systems by showing an equivalency between channels and billiard paths. We exploit this relationship to show that $2^{\frac{\gcd(m+1, n+1)-1}{2}}$ divides the number of domino tilings of the $m \times n$ rectangle. We also use billiard paths to give a fast algorithm for counting channels (and hence determining the parity of the number of domino tilings) in simply connected regions of the square grid.

1 Introduction

Given a graph G , a *perfect matching* of G is a set of edges μ such that each vertex of G is contained in a unique edge in μ . We let m_G denote the number of distinct perfect matchings of G . The problem of determining m_G arises in various mathematical contexts, particularly in tiling problems, but also in statistical mechanics [6], spectral graph theory [8], network analysis [4], total positivity [11], and representation theory [5]. Exact formulas for m_G over an infinite family of graphs are quite rare. One notable exact formula is for $G = \mathcal{R}_{m \times n}$, the rectangular subgraph of the square lattice with m rows of n vertices. In this case, a famous result of Kasteleyn [6] gives

$$m_G^4 = \prod_{j=1}^n \prod_{k=1}^m \left(4 \cos^2 \frac{j\pi}{n+1} + 4 \cos^2 \frac{k\pi}{m+1} \right).$$

From this product we may extract certain number-theoretic information: for example, m_G is always divisible by $2^{\frac{\gcd(n+1, m+1)-1}{2}}$ [13]. Studying similar 2-divisibility patterns is a common theme in the literature on domino tilings, which are equivalent to perfect matchings of subgraphs of the square lattice (see, e.g., [1, 3, 13, 15, 17, 18]). It is often the case that the 2-component of the prime factorization of m_G follows a predictable pattern, even when an exact formula for m_G is elusive or unwieldy. In Propp's perfect matching problem anthology [15], he gives a number of conjectured and known power of 2 patterns for various graphs. For example, the following is a refinement by Pachter [13] of one of these conjectures.

Conjecture 1.1 (Deleting from step-diagonals). *Let $\mathcal{R}_{2r \times 2r}$ be the $2r \times 2r$ grid graph shown in Figure 1. If G is a subgraph with the same vertex set, constructed by deleting any k of the edges highlighted in the figure, then*

$$m_G = 2^{r-k}b$$

for some odd integer b .




Figure 1: The step-diagonal for $r = 4$.

The following special case of a theorem of Ciucu is perhaps the most widely used 2-divisibility result in the literature.

Proposition 1.2 (Ciucu’s Factorization Theorem [1]). *If a bipartite planar graph has a line of symmetry containing $2r$ vertices, and no edges lie on the line or connect two vertices on opposite sides of the line, then the number of perfect matchings of the graph is divisible by 2^r .*

However, the symmetry requirement in Ciucu’s theorem means that it does not apply to graphs such as those described in Conjecture 1.1 or to $\mathcal{R}_{m \times n}$ for $m \neq n$ (though it is quite important for studying the square $\mathcal{R}_{m \times m}$). The results we introduce here provide a uniform (partial) explanation of power of 2 patterns in terms of the geometry of the graph. Our foundational construction is based on the following result known to Lovász.

Proposition 1.3 ([12], Problem 5.18). *Let G be a graph. Then m_G is even if and only if there is a nonempty vertex set $C \subseteq V$ such that every vertex in G is adjacent to an even number of vertices in C .*

We will call a vertex set C satisfying the hypothesis of Proposition 1.3 a *channel*. We also count the empty set as a trivial channel. Figure 2 shows examples of channels in different graphs. (The name “channel” is intended to evoke the image of a river winding its way around G , which may be redirected by “digging” edges and vertices out of G .) The symmetric difference of two channels is again a channel, so they form a vector space over $\mathbf{Z}/2\mathbf{Z}$. This implies that the number of distinct channels in a graph is always a power of 2. The space of channels can be identified with the kernel of the adjacency matrix of G modulo 2 (see Lemma 3.2).




Figure 2: Three graphs, each with a channel C depicted by colored rings around the vertices contained in C .

Lovász's result already shows the importance of channels for determining the parity of m_G . The main theorem of this paper shows that channels have even more to say for planar graphs.

Theorem 1.4 (Channeling 2s). *Let G be a planar graph. Then the number of distinct channels in G divides m_G^2 .*

Since the number of channels will always be a power of 2, Theorem 1.4 gives a lower bound on the power of 2 dividing m_G for any planar graph. We will prove this theorem in greater generality (for any graph admitting a Kasteleyn signing) in Theorem 3.5. We demonstrate the strength of this theorem in a number of examples throughout the article. In particular, we show 2-divisibility results related to graphs described above: that $2^{\frac{\gcd(n+1, m+1)-1}{2}}$ divides $m_{\mathcal{R}_{m \times n}}$ and that 2^{r-k} divides m_G for the graph described in Conjecture 1.1. These are Corollary 4.11 and Proposition 3.8, respectively.

Because of their utility, the majority of this paper is dedicated to studying the structure of channels and methods for finding them. Many of our results are tailored for subgraphs of the square lattice, where perfect matchings are equivalent to domino tilings of a region. (When possible, however, we will state results in greater generality.) Our most fascinating result is a characterization of channels in terms of dynamical systems. We state the result here for subgraphs of the square lattice and show the general case in Section 4.

Consider any simple cycle in the square lattice (considered as a grid graph). This cycle divides the square lattice into an interior and an (unbounded) exterior. Let G be the subgraph of the square lattice consisting of all vertices and edges in the cycle or its interior. In the dual language of domino tilings, graphs constructed in this way correspond roughly to simply connected subsets of a square grid. Since G is bipartite, we 2-color the vertices of G black and white. An example of such a graph is shown in Figure 3.




Figure 3: A graph G which is an induced subgraph of the square lattice bounded by a simple cycle.

Now we define a *billiard nest* on G to be any collection of paths traced out by billiard balls placed on black vertices of G and launched at 45 degree angles. When a billiard ball reaches a wall, it reflects at a 90 degree angle and proceeds in its new direction, continuing until it is caught by a corner or returns to its start position. (If the billiard brushes past a corner or hits one head-on, the situation is more complicated—the path splits into two paths continuing in different directions. See Section 4 for more examples and a precise definition.)




Figure 4: The three nonempty billiard nests in G .

Remarkably, channels and billiard nests are intrinsically connected. Let G' be the *inner subgraph* of G , the subgraph formed by removing all vertices of G which are incident to the unbounded face and all edges incident to those vertices.

Theorem 1.5. *Let G satisfy the assumptions described above, and let G' be the inner subgraph of G . Then the number of billiard nests in G is twice the number of channels in G' that contain only black vertices from G' .*

As a result, the number of billiard nests in G divides $2m_{G'}$. Moreover, if G has an equal number of black and white vertices, then G has exactly one nonempty billiard nest if and only if $m_{G'}$ is odd.

For example, for the graph G in Figure 3, the inner subgraph G' is shown in Figure 5. Since G has 4 billiard nests and satisfies the hypothesis of Theorem 1.5, it follows that $m_{G'}$ is divisible by 2. In this case we may simply count the perfect matchings of G' ; we find that $m_{G'} = 4$, which is divisible by 2 as anticipated by the theorem. In Theorem 4.7 and its corollary, we prove this billiard–channel correspondence in much greater generality (for *inner semi-Eulerian* graphs).




Figure 5: The inner subgraph G' of the graph G defined in Figure 3.

The connection between 2-divisibility, channels, and dynamical systems explains both the sensitivity and the regularity of perfect matching 2-divisibility. Small changes to G can result in entirely different billiard dynamics, with the effects visible in $m_{G'}$. For instance, if we take G to be the graph in Figure 6, then there is only one nonempty billiard nest.




Figure 6: A graph G which has only one nonempty billiard nest, along with its inner subgraph G' .

Since (counting the empty nest) G has 2 billiard nests, by Theorem 1.5 the inner subgraph G' has no nonempty channels on black vertices. Since G' has the same number of black and white vertices, this will imply (by Lemma 3.3) that G' has no nonempty channels at all. Thus by Proposition 1.3, G' has an odd number of perfect matchings—in this case 3.

The dynamics involved can also induce a regularity in the 2-divisibility of m_G . The well-known theory of arithmetic billiards describes billiard nests for rectangles in terms of divisibility properties of the rectangle side lengths. In Section 4.2, we exploit these results to explain the factor of $2^{\frac{\gcd(m+1, n+1)-1}{2}}$ dividing m_G for the $m \times n$ grid graph.

Billiard nests give a global explanation of channel structure for many graphs. Sometimes we are instead interested in local behavior. For instance, we may have a family of graphs that are globally similar but that differ locally. To relate these graphs, we introduce a set of channel-preserving graph operations and show that they may be applied repeatedly to reduce many graphs to a graph with no edges, whose channels are apparent.

The paper is organized as follows. The proof of Theorem 1.4 is algebraic and given in Sections 2 and 3. Sections 4, 5, and 6 are independent and may be read in any order. Section 4 describes

billiard nests for a large class of graphs called *inner semi-Eulerian graphs*. The results described in the introduction are applied to the rectangle grid graph, connecting its 2-divisibility to the theory of arithmetic billiards. Section 4 concludes with a fast algorithm that constructs the billiard nests and therefore the channels for certain graphs. In Section 5, we give a combinatorial proof of Proposition 1.3 for bipartite graphs. In the course of this proof we introduce a graph move called *channel routing*, which involves removing adjacent vertex pairs from a graph while tracking the effect on channels. Channel routing is not always well-behaved, but certain other graph moves always preserve the number of channels in a graph. These are called *channel-preserving moves* and are the focus of Section 6. Section 6 also introduces a useful graph move called *diagonal contraction*. We give results on diagonal contraction and channels which generalize a number of known domino tiling parity results. We wrap up in Section 7 with remarks and directions for future work.

2 Preliminaries

All graphs in this paper are undirected, finite, and (unless otherwise indicated) contain no self-loops. Until Section 6, we further assume that there is at most one edge between any pair of vertices. This convention is used only to simplify notation; using suitable definitions, all results hold without it. If a graph is bipartite, we will consider its vertices to be colored black and white. Additionally, all matchings discussed will be perfect matchings, and thus the word “perfect” will be omitted in the future for brevity. For a graph $G = (V, E)$, V denotes the vertex set, E denotes the edge set, and A denotes the adjacency matrix. We write \overline{vw} to denote an (undirected) edge between v and w ; this is to contrast with the notation $\{v, w\}$ indicating a set of two vertices. Given a vertex v , the neighborhood of v is

$$N(v) = \{v' \mid \overline{vv'} \in E\}.$$

For an edge e , we use the notation $G - e$ to denote the subgraph $(V, E - e)$. For a subset $S \subseteq V$, we use the notation $G - S$ to denote the subgraph of G induced on $V - S$. Recall that m_G is the number of matchings of G . As an exercise in this notation, we have the following proposition.

Proposition 2.1. *Let G be any graph, and let $e = \overline{v_1v_2}$. Then*

$$m_G = m_{G-e} + m_{G-\{v_1, v_2\}}.$$

Also, fix any vertex v . Then

$$m_G = \sum_{v' \in N(v)} m_{G-\{v, v'\}}.$$

Proof. For the first relation, notice that matchings of $G - e$ are just those matchings of G that do not use the edge e . The other matchings of G do use e , and therefore for these matchings the vertices v_1 and v_2 are never in an edge with any vertex other than each other. Thus such matchings are equivalent to matchings of $G - \{v_1, v_2\}$, plus the edge e , and the first equation is shown.

For the second relation, partition the set of matchings of G based on the vertex that pairs with v in the matching. By the same reasoning as in the last paragraph, the number of matchings in which v pairs with v' is $m_{G-\{v, v'\}}$. Summing over the possible pairings shows the claim. \square

Sometimes we will be interested in connected planar graphs G . Such graphs admit a dual graph, with vertices given by the faces of G and edges between faces separated by an edge in G . If the same face is on both sides of an edge of G , then that edge corresponds to a self-loop in the dual

graph. The *external face* of G is the face which is unbounded, and all other faces are *internal faces* of G . The *reduced dual graph* of G is the dual graph of G with the vertex corresponding to the external face of G removed. We say a vertex of G is *external* if it is incident to the external face, and we say it is *internal* otherwise.

We begin by recalling Lovász's original proof of Proposition 1.3. Later, in Section 5, we will reprove this combinatorially for bipartite graphs.

Proposition 2.2 ([12], Problem 5.18). *Let G be a finite graph with no self-loops. Then m_G is even if and only if there is a nonempty vertex set $C \subseteq V$ such that every vertex in G is adjacent to an even number of vertices in C .*

Proof. Let $A_2 = (a_{ij})_{i,j=1}^n$ be the adjacency matrix of G modulo 2. Note that vertex sets as described in the proposition exist if and only if the kernel of A_2 is nonzero (since an element of the kernel is a sequence of 1s and 0s indexed by vertices, which may be treated as an indicator function for a vertex set that will necessarily have the desired property).

The determinant of A_2 is

$$\det A_2 = \sum_{\sigma \in S_n} \operatorname{sgn}(\sigma) a_{1\sigma(1)} \cdots a_{n\sigma(n)} = \sum_{\sigma \in S_n} a_{1\sigma(1)} \cdots a_{n\sigma(n)},$$

where the sum is over all permutations of $\{1, \dots, n\}$. (We may ignore the sign since we work over $\mathbf{Z}/2\mathbf{Z}$.) Because the adjacency matrix is symmetric,

$$a_{1\sigma(1)} \cdots a_{n\sigma(n)} = a_{\sigma^{-1}(1)1} \cdots a_{\sigma^{-1}(n)n} = a_{1\sigma^{-1}(1)} \cdots a_{n\sigma^{-1}(n)}.$$

Thus we may remove from the sum all pairs $\sigma \neq \sigma^{-1}$ as such pairs yield terms that sum to 0. We are left with the sum over involutions

$$\sum_{\sigma = \sigma^{-1}} a_{1\sigma(1)} \cdots a_{n\sigma(n)}.$$

Now, because G has no self loops, if σ satisfies $\sigma(i) = i$ for some i , then the corresponding term in the sum contains $a_{ii} = 0$. Thus we are left with a sum over involutions of the vertices with no fixed points. For such an involution σ , $a_{1\sigma(1)} \cdots a_{n\sigma(n)}$ is 1 if there exists a matching of G that pairs vertex i with vertex $\sigma(i)$ for all i , and 0 otherwise. Hence we get the equality (in $\mathbf{Z}/2\mathbf{Z}$)

$$m_G = \det A_2.$$

Thus $m_G = 0$ in $\mathbf{Z}/2\mathbf{Z}$ if and only if $\ker A_2 \neq 0$, and we are done. \square

Now we discuss some important algebraic constructions. Our main tools for the remainder of the section are the Kasteleyn matrix of a graph and the Smith decomposition of a matrix.

Proposition 2.3 (Kasteleyn [7]). *Let G be a planar graph with adjacency matrix $A = (a_{ij})$. Then there exists a matrix K such that $K = (\pm a_{ij})$, and*

$$\det K = m_G^2.$$

The matrix K is called a *Kasteleyn matrix*, and if a (not necessarily planar) graph G admits such a K , then G is said to have a *Kasteleyn signing*.

For bipartite graphs we can be more specific.

Definition 2.4. Let G be a bipartite graph with adjacency matrix A . The *bipartite adjacency matrix* of G is the submatrix of A formed by selecting rows from A associated to white vertices and columns from A associated to black vertices.

When G is bipartite, the adjacency matrix A can be written in block matrix form as

$$A = \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix},$$

where B is the bipartite adjacency matrix. If in addition G is planar and has the same number of black and white vertices, then there is an analogue to the Kasteleyn matrix called the *bipartite Kasteleyn matrix*.

Proposition 2.5 (Percus [14]). *Let G be a bipartite planar graph with bipartite adjacency matrix $B = (b_{ij})$. If B is square, then there exists a matrix H , the bipartite Kasteleyn matrix, such that $H = (\pm b_{ij})$, and*

$$\det H = m_G.$$

We will want to diagonalize these matrices over the integers. The canonical tool for doing so is called *Smith normal form*.

Proposition 2.6. *Let A be a matrix over a principal ideal domain (PID) \mathcal{R} . Then there exist matrices S, D, T over \mathcal{R} with the following properties:*

- (i) $A = SDT$.
- (ii) S and T are invertible over \mathcal{R} . For $\mathcal{R} = \mathbf{Z}$, this means $\det S, \det T = \pm 1$.
- (iii) D is diagonal, with diagonal entries $\alpha_1, \dots, \alpha_n$ satisfying

$$\alpha_1 \text{ divides } \alpha_2 \text{ divides } \dots \text{ divides } \alpha_n.$$

The matrices S, D, T are called a Smith decomposition for A .

For our purposes, \mathcal{R} will be the integers or a finite field. Smith decompositions have many useful properties. See, e.g., [10] or [16] for more background and combinatorial applications of Smith decompositions. The next result follows directly from Proposition 2.6(i)–(ii).

Proposition 2.7. *Let $A = SDT$ be a Smith decomposition for a square matrix A over a PID \mathcal{R} . Then*

$$\ker A \cong \ker D \text{ as abelian groups and } \det A = u \det D \text{ for some unit } u \in \mathcal{R}.$$

We are now prepared to study 2-divisibility of m_G for graphs G with a Kasteleyn matrix K . As described above, planarity is a sufficient condition for G to have a Kasteleyn signing. Using the Smith normal form of K , we will find 2-divisibility results for such graphs that we develop further in the next section.

Definition 2.8. For an integral matrix A , define the *reduction of A modulo 2* to be the matrix A_2 over $\mathbf{Z}/2\mathbf{Z}$ given by reducing the entries of A modulo 2 and considering them as elements of $\mathbf{Z}/2\mathbf{Z}$.

Define $\ker_2 A$, the *2-kernel* of A , to be the kernel of A_2 as a vector space over $\mathbf{Z}/2\mathbf{Z}$.

Notice that for an adjacency matrix A and corresponding Kasteleyn matrix K , we have $A_2 = K_2$. This follows from the definition of K as a signed version of A . This is a key observation that will allow us to translate our algebraic results in this section into geometric results in Section 3. Before that, let us see what we can learn from reducing the Kasteleyn matrix mod 2. Let the *2-nullity* of a matrix be the dimension of its 2-kernel:

$$\text{null}_2 A := \dim \ker_2 A.$$

Lemma 2.9. *Let A be a square matrix with integer entries. Then*

$$2^{\text{null}_2 A} \mid \det A.$$

Proof. Let $A = SDT$ be a Smith decomposition of A . Reducing modulo 2 gives $A_2 = S_2 D_2 T_2$. One may check that this is a Smith normal form of A_2 . Then by Proposition 2.7, the kernel of A_2 is isomorphic to the kernel of D_2 . Set $k := \text{null}_2 A = \text{null}_2 D$. Since D_2 is diagonal, its kernel has a basis consisting of standard basis vectors that indicate columns where the diagonal entry is 0. Therefore there are exactly k such entries. Since 0 entries in D_2 correspond to even integral entries in D , there are exactly k even entries on the diagonal of D . Thus the determinant of D contains at least k factors of 2, and the result follows by Proposition 2.7. \square

Applying this lemma to the Kasteleyn matrix or the bipartite Kasteleyn matrix will let us use the 2-kernel to find powers of two in the number of matchings of a graph.

Theorem 2.10. *Let G be a graph with a Kasteleyn signing (e.g. a planar graph). If A is the adjacency matrix of G , then*

$$2^{\text{null}_2 A} \text{ divides } m_G^2.$$

Proof. As remarked above, $A_2 = K_2$, so in particular $\text{null}_2 A = \text{null}_2 K$. Additionally, by Proposition 2.3, the determinant of K is m_G^2 . Thus by Lemma 2.9,

$$2^{\text{null}_2 A} = 2^{\text{null}_2 K} \text{ divides } \det K = m_G^2. \quad \square$$

Theorem 2.11. *Let G be a bipartite graph with a Kasteleyn signing. If B is the bipartite adjacency matrix of G , then*

$$2^{\text{null}_2 B} \text{ divides } m_G.$$

Proof. If B is not square, then $m_G = 0$ and the claim holds trivially. Otherwise, the proof is the same as the previous theorem using the bipartite Kasteleyn matrix. \square

We therefore may deduce powers of 2 dividing m_G by finding elements of the 2-kernel of A . The remainder of the paper details how this can be done.

3 Channels

Let G be a graph with adjacency matrix A . Then a vector x in $\ker_2 A$ has entries in $\mathbf{Z}/2\mathbf{Z}$ and can be lifted to a vector \tilde{x} with entries $0, 1 \in \mathbf{Z}$. The condition $A_2 x = 0$ then becomes $A\tilde{x} = 2y$ for some integral vector y . Because each row of x corresponds to a vertex in G , we may interpret x as the indicator function for a vertex set C , where a row with a 1 indicates the vertex is in C and a row with a 0 indicates the vertex is not in C . This leads to the following interpretation of 2-kernel elements.

Definition 3.1. Let $G = (V, E)$ be any graph. A *channel* is a set C of vertices such that every vertex in G is adjacent to an even number of vertices in C . In other words, letting $N(v)$ denote the neighborhood of v , a channel satisfies

$$|N(v) \cap C| \text{ is even, for all } v \in V.$$

Let the set of channels in G be denoted $\mathcal{C}(G)$. If G is bipartite, then let $\mathcal{C}_B(G)$ (resp. $\mathcal{C}_W(G)$) be the subset of $\mathcal{C}(G)$ consisting of channels that use only black (resp. white) vertices from G .

The 2-kernel also has an additive structure as a $\mathbf{Z}/2\mathbf{Z}$ vector space. This transfers to $\mathcal{C}(G)$ by defining the sum of $C_1, C_2 \in \mathcal{C}(G)$ to be

$$C_1 \oplus C_2 := (C_1 \cup C_2) - (C_1 \cap C_2), \text{ the symmetric difference of } C_1 \text{ and } C_2.$$

Note that in the bipartite case, both $\mathcal{C}_B(G)$ and $\mathcal{C}_W(G)$ are subspaces of $\mathcal{C}(G)$.




Figure 7: A graph G with its three nonempty channels indicated by shading. Any two of these form a basis for the space $\mathcal{C}(G)$.

The discussion above then implies the following, which was hinted at in the proof of Proposition 2.2.

Lemma 3.2. *The spaces $\mathcal{C}(G)$ and $\ker_2 A$ are isomorphic vector spaces over $\mathbf{Z}/2\mathbf{Z}$.*

If G is additionally bipartite with bipartite adjacency matrix B , then we can refine Lemma 3.2 by accounting for vertex colors.

Lemma 3.3. *If G is bipartite, then there are $\mathbf{Z}/2\mathbf{Z}$ -vector space isomorphisms*

$$\mathcal{C}_B(G) \cong \ker_2 B \quad \text{and} \quad \mathcal{C}_W(G) \cong \ker_2 B^T.$$

Furthermore,

$$\mathcal{C}(G) = \mathcal{C}_B(G) \oplus \mathcal{C}_W(G).$$

Proof. The identification of channels using only black vertices with elements of $\ker_2 B$ follows the same lines as the discussion at the beginning of the section. The claim for white vertices is the same, noting that taking the transpose of B is equivalent to swapping the two vertex colors.

For the final claim, the set of black vertices in any channel C is also a channel: in a bipartite graph, the set of black vertices adjacent to a given vertex is either the empty set (which has even cardinality) if the vertex is black, or its entire neighborhood if the vertex is white. Similarly the white vertices in C also form a channel. As a result, any channel may be split into its black part and its white part. Since $\mathcal{C}_B(G)$ and $\mathcal{C}_W(G)$ have trivial intersection, this proves the direct sum decomposition. \square

Lemma 3.4. *Let $|V_B|$ and $|V_W|$ be the number of black vertices and white vertices in a bipartite graph G , respectively. Then*

$$\dim \mathcal{C}_B(G) - \dim \mathcal{C}_W(G) = |V_B| - |V_W|.$$

In particular, if G has the same number of black and white vertices, then

$$\dim \mathcal{C}_B(G) = \dim \mathcal{C}_W(G) = \frac{1}{2} \dim \mathcal{C}(G).$$

Proof. Let B be the bipartite adjacency matrix of G . We have

$$\dim \mathcal{C}_B(G) = \text{null}_2 B \tag{1}$$

and

$$\dim \mathcal{C}_W(G) = \text{null}_2 B^T = \dim \text{coker}_2 B = |V_W| - \dim \text{im}_2 B, \tag{2}$$

where coker_2 and im_2 refer to cokernel and image of B as a matrix over $\mathbf{Z}/2\mathbf{Z}$. The rank-nullity theorem for B may be stated as

$$\text{null}_2 B + \dim \text{im}_2 B = |V_B|.$$

The claim then follows by subtracting equation (1) from equation (2) and applying rank-nullity. \square

Combining the observations above with the 2-divisibility results from the last section, we can now prove our main theorem.

Theorem 3.5 (Channeling 2s). *If C_1, \dots, C_n are linearly independent channels in a graph G with a Kasteleyn signing, then*

$$2^n \text{ divides } m_G^2.$$

If additionally G is bipartite, and $C_1, \dots, C_n \in \mathcal{C}_B(G)$, then

$$2^n \text{ divides } m_G.$$

Proof. Let A be the adjacency matrix of G . By Lemma 3.2, C_1, \dots, C_n may be viewed as n linearly independent elements of $\ker_2 A$. It follows that $\text{null}_2 A \geq n$. Then the first claim results from applying Theorem 2.10.

The second claim follows similarly, using Lemma 3.3 and Theorem 2.11. \square

Remark 3.6. Despite the fact that Proposition 2.2 holds for an arbitrary graph, Theorem 3.5 does not. For example, the complete bipartite graph $K_{3,3}$ has $|\mathcal{C}_B(K_{3,3})| = 2^2$, but $m_{K_{3,3}} = 6$ is not divisible by 2^2 . Thus the assumption of a Kasteleyn signing for G cannot be weakened much further.




Figure 8: A basis for $\mathcal{C}_B(\mathcal{R}_{4 \times 9})$.

Example 3.7. Let $\mathcal{R}_{m \times n}$ denote the $m \times n$ rectangular grid graph. The shading in Figure 8 shows a basis for $\mathcal{C}_B(\mathcal{R}_{4 \times 9})$. By channeling 2s, we have that 2^2 divides $m_{\mathcal{R}_{4 \times 9}}$. And indeed, $m_{\mathcal{R}_{4 \times 9}} = 6336$ is divisible by 4. Note, however, that 6336 is also divisible by 2^6 —Theorem 3.5 gives only a *lower bound* on the power of 2 dividing m_G . \diamond

In Section 4, we shall employ billiard paths to count channels in a rectangle grid graph of arbitrary size. We already have the results we need, however, to give a lower bound supporting Pachter’s conjecture.

Proposition 3.8. Let $\mathcal{R}_{2r \times 2r}$ be the $2r \times 2r$ grid graph shown in Figure 1. If G is constructed from $\mathcal{R}_{2r \times 2r}$ by deleting any k of the highlighted edges in the figure, then

$$2^{r-k} \text{ divides } m_G.$$

Proof. For a $2r \times 2r$ square grid graph, we will find that the space $\mathcal{C}_B(\mathcal{R}_{2r \times 2r})$ has r independent channels which each intersect exactly one of vertex pairs removed from the step diagonal. Thus removing k of the step diagonal vertex pairs interrupts just k of the channels, so the other $r - k$ channels will still be present. The result then follows from channeling 2s once we construct these channels.

Consider Figure 9. To construct each channel, pick a black vertex b on the step diagonal. In the figure, these are the vertices along the diagonal from the bottom left to the top right. Our channel will consist of four (possibly empty) diagonal segments of vertices. Two segments will intersect the step diagonal transversely, one at b in the bottom left and one at b ’s mirror image in the top right. The other two segments will be parallel to the step diagonal and placed so that the channel vertices along the sides of the grid graph each have one intervening vertex between them. Each




Figure 9: A basis for $\mathcal{C}_B(\mathcal{R}_{8 \times 8})$. The diagonal segments which intersect the step diagonal are shown in purple, while the diagonal segments which are parallel to the step diagonal are in green. Note that the channels in the top left and bottom right are also channels of the graph with the two red edges removed.

such intervening vertex is adjacent to one endpoint from two diagonal segments. All other vertices are adjacent to either 0 or 2 vertices from each diagonal segment. Hence each vertex of the graph is adjacent to an even number of vertices from our diagonal segments, so these vertices form a channel.


By following this construction for each b lying along the lower r black vertices of the step diagonal, we create r channels that each contain a different vertex from the step diagonal, and therefore they must be independent. \square

4 Billiards and Channels


For an arbitrary graph, it is not at all clear how to construct or count its channels without solving the requisite linear system. In this section we give a geometric approach to channel construction based on a phenomenon that can be observed in the channels of a rectangle grid graph.

4.1 Billiard nests

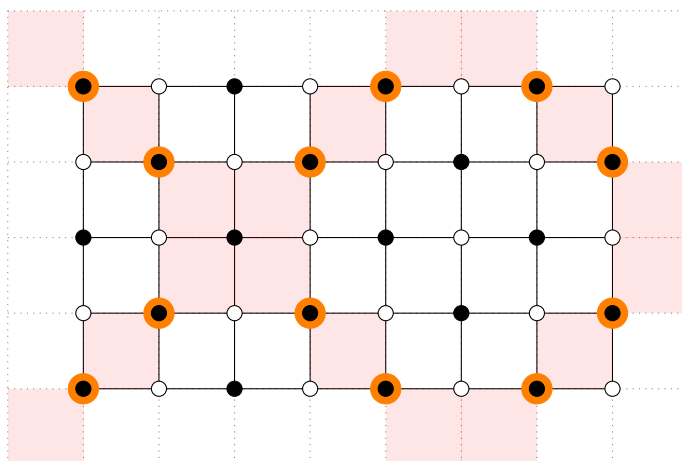
In the rectangle, we note that channels tend to form along diagonal lines as in the following figure. This pattern was studied by Tomei and Vieira in [19], where they described it in terms of polygonal tilings of the rectangle. We propose an alternative description.



By extending these diagonals, we find that these lines form a path which reflects off the edges of a larger rectangle, as shown below. The channel vertices are vertices in the interior of this larger rectangle which intersect exactly one line from this path.



Such paths either form loops or start and end on distinct corners. Notice that we may recover the path by remembering only the faces that it passes through and the fact that it passes through black vertices. The face information is shown in the next figure.



It turns out to be natural to consider arbitrary collections of billiard paths—what we will call a *billiard nest*. Viewing the nest as a set of faces will allow us to isolate the abstract properties of billiards which are of combinatorial importance. From this viewpoint, billiard nests can be defined on a large class of graphs, called *inner semi-Eulerian graphs*.

Definition 4.1. We say that a bipartite planar graph G is *inner semi-Eulerian* if every internal black vertex of G has even degree. Let G be inner semi-Eulerian and let F denote the set of internal faces of G . We say that a subset of faces $B \subseteq F$ of G is a *billiard nest* if the following hold:

- if b is an internal black vertex, then either all faces incident to b are in B , no faces incident to b are in B , or every second face incident to b is in B .

- if b is an external black vertex, then either all internal faces incident to b are in B or no internal faces incident to b are in B .

Denote the set of billiard nests in G by $\mathcal{B}(G)$.

When G is an induced subgraph of the square lattice with all internal faces being unit squares, this roughly agrees with the intuitive notion of billiard nests as (the faces containing) paths traced out by a collection of billiard balls since at an internal black vertex, the ball will pass from one face to the opposite one at that vertex. (At an external black vertex the billiard ball can be thought of as splitting into all possible directions.) Note that \emptyset and F are trivially billiard nests for every graph. As with channels, we may define the sum of two billiard nests to be their symmetric difference, making $\mathcal{B}(G)$ a vector space over $\mathbf{Z}/2\mathbf{Z}$.

There is a canonical basis for $\mathcal{B}(G)$ such that the basis billiard nests are mutually disjoint. Indeed, define a graph $G_{\mathcal{B}}$ with vertex set F and edges between f and f' if they satisfy one of the following:

- f and f' are incident to the same internal black vertex b and are separated by an even number of edges incident to b .
- f and f' are incident to the same external black vertex b .

Then the connected components of $G_{\mathcal{B}}$ are independent billiard nests that span $\mathcal{B}(G)$. This is called the *path basis* for $\mathcal{B}(G)$, and its elements are called *billiard paths*. Later we will demonstrate an efficient algorithm to find the path basis for certain graphs. This is particularly useful since, as we shall soon see, billiard nests in G are equivalent to channels in a certain subgraph of G .

Definition 4.2. Let G be inner semi-Eulerian. The *inner subgraph* of G , denoted G' , is the induced subgraph on the internal vertices of G . Similarly, the *outer subgraph* of G is the induced subgraph on the external vertices of G .

Given any inner semi-Eulerian graph H , an *outer completion* of H is an inner semi-Eulerian graph G such that $G' = H$ and such that the outer subgraph of G is a simple cycle.

We can always construct an outer completion for an inner semi-Eulerian graph H by taking a copy of the boundary of H , expanding the copy so that H lies within it, and adding edges between the two copies of the boundary as needed to make the graph inner semi-Eulerian. This is described in detail in the next proposition and the following example.


Proposition 4.3. *Any inner semi-Eulerian graph H has an outer completion.*

Proof. Without loss of generality, let H be connected. Set n to be the degree of the external face f of H in the dual graph H^\vee of H . The cyclic ordering on the faces adjacent to f in H^\vee induces a cyclic ordering on the external vertices of H (possibly with repetition). Construct a simple cycle Y with n vertices and edges, and embed Y in the plane such that H is in the interior of Y . Then the cyclic order of the external vertices of H surjectively assigns to each vertex of Y an external vertex of H . For each external vertex v of H , if v has odd degree, then place an edge between v and one of the vertices of Y to which v is assigned.


Call the resulting graph G . Then G is bipartite since we may color each vertex in Y with the color opposite of the vertex assigned to it. Also G is planar: the edges we introduced between Y and H may be embedded so that they do not cross since they connect to the boundary of H in the same order as they do to Y . Finally, G is inner semi-Eulerian since the construction of G makes

all external vertices of H have even degree, and all other internal black vertices of G are internal vertices of H . □

Example 4.4. Let H be the following graph.



The degree of the outer face f is 14, and the sequence of external vertices visited as we traverse the boundary is indicated by the numbering below. Construct a simple cycle Y of length 14 and embed it so H lies in the interior face of Y .



For each external vertex of H with odd degree, pick one of the vertices associated to it in Y and add an edge between them. The resulting graph G is inner semi-Eulerian.




Figure 10: An inner semi-Eulerian graph G with inner subgraph H .

◇

Remark 4.5. Often there is a more natural choice of outer completion for H than the construction described above. In particular, many of our examples that are subgraphs of the square lattice use an outer completion that is also a subgraph of the square lattice.

Given a billiard nest in G , we may construct an associated channel in G' as follows. Let $B \in \mathcal{B}(G)$ be a billiard nest. Define

$$\text{ch}: \mathcal{B}(G) \longrightarrow \mathcal{C}_B(G')$$

by setting the vertices in $\text{ch}(B)$ to be the internal black vertices b of G for which exactly half of the faces incident to b are in B .

Lemma 4.6. *The map ch is a group homomorphism from $\mathcal{B}(G)$ to $\mathcal{C}_B(G')$.*

Proof. Let $B \in \mathcal{B}(G)$. We first check that $\text{ch}(B)$ is in fact a channel. Let w be a white vertex of G' , i.e., an internal white vertex of G . Consider the edges incident to w . We wish to show that the number of these edges that contain a channel vertex is even. Let f_1, \dots, f_n be the incident faces to w in cyclic order. If f_i and f_{i+1} are both in B or both not in B , then the edge between them does not contain a vertex in $\text{ch}(B)$; otherwise, the edge contains a vertex in $\text{ch}(B)$. Thus the number of incident edges to w containing vertices from $\text{ch}(B)$ is the number of times f_i changes from being in B to not being in B or vice versa as we traverse the incident faces. Since after traversing all of the faces incident to w we must arrive back at the starting face, we must change state an even number of times. Thus $\text{ch}(B)$ is a channel.

Now we check that ch preserves symmetric differences. Let $B_1, B_2 \in \mathcal{B}(G)$. Consider the incident faces of b for an arbitrary black vertex $b \in G'$. If $b \in \text{ch}(B_1)$, then every other of these faces is in B_1 . If $b \notin \text{ch}(B_1)$, then all or none of these faces are in B_1 . The same holds for B_2 . Then the claim follows since the symmetric difference of the incident faces to b in B_1 with the incident faces to b in B_2 is every second face if and only if exactly one of B_1 and B_2 contains every second face incident to b , which occurs if and only if b is in exactly one of $\text{ch}(B_1)$ and $\text{ch}(B_2)$. □

We can now state our main result on billiard nests. Recall that the reduced dual graph of G is the dual graph with the vertex corresponding to the external face of G removed.

Theorem 4.7. *Let G be inner semi-Eulerian and let G' be the inner subgraph of G . Assume that the outer subgraph of G is a simple cycle. Then*

$$|\mathcal{B}(G)| = 2|\mathcal{C}_B(G')|.$$

Proof. First we show that

$$|\ker \text{ch}| = 2.$$

Let $B \in \mathcal{B}(G)$ be such that $\text{ch}(B) = 0$. Then B is locally constant, i.e., at each black vertex b , either all faces incident to b are in B or all faces incident to b are not in B . In a connected component of the reduced dual graph of G , there is a path between any two faces. Any edge between two faces on that path contains a black vertex, which enforces the locally constant condition. Thus B is constant along the entire path and therefore across any connected component of the reduced dual of G . So the claim follows if we can show that the reduced dual of G is connected. Indeed, since G is surrounded by a simple cycle, we may embed it such that it fills a convex region in the plane. Then a generic line connecting two internal faces will induce a path between them in the dual graph, which will avoid the external face by convexity. This proves that $B \in \mathcal{B}(G)$ if and only if B is constant, i.e., equal to either F or \emptyset . Hence $|\ker \text{ch}| = 2$.

To complete the proof, we show that ch is surjective. Let $C \in \mathcal{C}_B(G')$ be a channel. We will construct a billiard nest B such that $\text{ch}(B) = C$. Fix a spanning tree T of the reduced dual of G and some internal face f_0 of G . We decide whether an internal face f is in B as follows. If $f = f_0$, then $f \in B$. Otherwise, there is a unique path in T from f_0 to f . If the path crosses an even number of edges in G containing vertices in C , then $f \in B$. Otherwise, $f \notin B$.

Claim. If f_1, f_2 are two adjacent internal faces of G , then f_1 and f_2 are both in B or both not in B if and only if the edge e_1 separating f_1 and f_2 does not contain a channel vertex.

If the dual edge e_1^\vee is a part of T , then this follows by definition of B . Otherwise, adding e_1^\vee to T creates a simple planar cycle

$$Y = \{e_1^\vee, \dots, e_n^\vee\} \subseteq T \cup \{e_1^\vee\}.$$

We wish to show that the number of channel vertices in e_1 has the same parity as the total number of channel vertices in e_2, \dots, e_n . We will be done if we can show that the number of edges in

$$Y^\vee = \{e_1, \dots, e_n\}$$

that contain a channel vertex is even.

Let H be the induced subgraph of G on the vertices in the interior of Y . Each e_i connects a vertex $v_i \in H$ to a vertex $v'_i \in G \setminus H$. Set

$$\partial H = \{v_1, \dots, v_n\}.$$




Figure 11: The graph G near a dual graph cycle Y , shown as a gray hexagon. The vertices in the interior of Y form the subgraph H . The channel C is shown in orange.

Since H is bipartite and C uses only black vertices,

$$\sum_{b \in C \cap H} \deg_H b = \sum_{w \in H} \text{chdeg}_H w, \quad (4.1)$$

where the sum on the left is over black vertices, the sum on the right is over white vertices, and $\text{chdeg}_H w$ is the size of the neighborhood of w in $C \cap H$.

Note that for $b \in H \setminus \partial H$ we have

$$\deg_H b = \deg_G b \equiv_2 0$$

since G is inner semi-Eulerian and all vertices of H are internal to G . (Here we write $n \equiv_2 m$ to mean that $n - m$ is even.) Thus

$$\sum_{b \in C \cap H} \deg_H b \equiv_2 \sum_{b \in C \cap \partial H} \deg_H b \equiv_2 \#\{e_i \mid v_i \text{ is in } C\}.$$

The last equivalence follows since

$$\deg_H b = \deg_G b - \#\{e_i \mid b \in e_i\} \equiv_2 \#\{e_i \mid b \in e_i\}.$$

Working now with the other side of (4.1), for $w \in H \setminus \partial H$ we have that

$$\text{chdeg}_H w = \text{chdeg}_G w = \text{chdeg}_{G'} \equiv_2 0$$

by definition of a channel. Thus

$$\sum_{w \in H} \text{chdeg}_H w \equiv_2 \sum_{w \in \partial H} \text{chdeg}_H w \equiv_2 \#\{e_i \mid v'_i \text{ is in } C\}.$$

The last equivalence here follows since

$$\text{chdeg}_H w = \text{chdeg}_G w - \#\{e_i \mid w \in e_i \text{ and } v'_i \in C\} \equiv_2 \#\{e_i \mid w \in e_i \text{ and } v'_i \in C\}.$$

Substituting our results into (4.1), we find

$$\#\{e_i \mid v_i \text{ is in } C\} \equiv_2 \#\{e_i \mid v'_i \text{ is in } C\}.$$

Thus

$$\#\{e_i \mid e_i \text{ contains a vertex from } C\} = \#\{e_i \mid v_i \text{ is in } C\} + \#\{e_i \mid v'_i \text{ is in } C\} \equiv_2 0$$

as desired, proving the claim.

We may now verify that B is indeed a billiard nest such that $\text{ch}(B) = C$. Let $b \in G$ be a black vertex. If $b \in C$, then every edge incident to b contains a channel vertex. Thus, by the preceding claim, the faces incident to b must alternate between being in B and not being in B . If $b \notin C$, then every edge incident to b does not contain a channel vertex. Thus, by the claim, the internal faces incident to b are either all in B or all not in B . In particular, external black vertices are not in C (since they are not in G'), so B meets the conditions for being a billiard nest. Furthermore, it is clear from this description that $\text{ch}(B) = C$. Thus ch is a surjective homomorphism

$$\mathcal{B}(G) \longrightarrow \mathcal{C}_B(G').$$

Combining this with the size of the kernel computed earlier, we find that

$$|\mathcal{B}(G)| = 2|\mathcal{C}_B(G')|. \quad \square$$

Corollary 4.8. *Let G be inner semi-Eulerian and let G' be the inner subgraph of G . Assume that the outer subgraph of G is a simple cycle and that G has the same number of black and white vertices. Then*

$$\frac{1}{2}|\mathcal{B}(G)| \text{ divides the number of matchings of } G'.$$

Moreover, $m_{G'}$ is odd if and only if G has exactly one nonempty billiard nest.

Proof. By Theorem 4.7, we have

$$\frac{1}{2}|\mathcal{B}(G)| = |\mathcal{C}_B(G')|.$$

By channeling 2s (Theorem 3.5), we know that $|\mathcal{C}_B(G')|$ divides the number of matchings of G' . By Proposition 2.2, $m_{G'}$ is odd if and only if $\dim \mathcal{C}(G') = 2 \dim \mathcal{C}_B(G') = 0$, that is, if and only if $|\mathcal{C}_B(G')| = 1$. The claim follows. \square

Thus our study of channels in appropriate graphs H (in particular, by Proposition 4.3, all inner semi-Eulerian graphs) reduces to the study of billiard nests in an outer completion G . Billiard nests are considerably easier to work with since every face of G is contained in a unique billiard path in the path basis of G . In general there is no such basis for the channels of H ; vertices of H may be contained in no channel and there may be no channel basis for H with pairwise disjoint elements. However, any billiard path can be found by starting with a face of G and adding additional faces as required by the definition of billiard nests.

4.2 Arithmetic billiards

As an application, let us find the billiard paths for the rectangle grid graph $\mathcal{R}_{m+1 \times n+1}$, an outer completion of $\mathcal{R}_{m-1 \times n-1}$. For this graph, we may use our interpretation of billiard paths as the paths through black vertices traced out by billiard balls travelling at 45 degree angles. (Explicitly, a face in the billiard path is interpreted as a billiard ball traveling between the black vertices in that face.) We begin by straightening out the billiard paths; to do so, we tile the plane with copies of our rectangle.




Figure 12: A rectangle used to tile the plane. The coloring is provided as a visual indicator of the rectangle's orientation. Any billiard path on the left corresponds to a line of slope 1 on the right and vice versa.

We can then lift the billiard path to a straight line of slope 1 in the tessellation. A billiard path between two corners of $\mathcal{R}_{m+1 \times n+1}$ will be the diagonal of a square in the tessellation. Any such square must have a side length divisible by m and n (the side lengths of $\mathcal{R}_{m+1 \times n+1}$). The minimal square in the tessellation with corners from $\mathcal{R}_{m+1 \times n+1}$ then has side length given by the least common multiple of m and n . To determine the number of internal faces of $\mathcal{R}_{m+1 \times n+1}$ through which the path travels, we may count the number of unit squares through which the path travels in the tessellation. Since the straightened billiard path travels along the diagonal of a square with side length $\text{lcm}(m, n)$, this path travels through $\text{lcm}(m, n)$ unit squares.

For now assume that at least one of $m + 1$ and $n + 1$ is even. Then exactly two corners of $\mathcal{R}_{m+1 \times n+1}$ are black. Any two distinct billiard paths pass through distinct internal faces of $\mathcal{R}_{m+1 \times n+1}$. We shall count path basis elements by counting the internal faces through which they pass. There is one billiard path through the black corners. From the last paragraph, we know this path uses $\text{lcm}(m, n)$ internal faces. Now, every other path on black vertices uses *twice* as many internal faces. Indeed, since the other paths do not pass through a corner, they must end on their starting point. To reach their starting point in the tessellation, the paths must lift to the diagonal of a square of side length $2 \text{lcm}(m, n)$ since one of m and n is odd. Since every internal face is part of a unique path basis element, we may now count the billiard paths for $\mathcal{R}_{m+1 \times n+1}$.

Theorem 4.9. *The rectangle grid graph $\mathcal{R}_{m+1 \times n+1}$ with $(m + 1)(n + 1)$ even has a path basis of*

size

$$\frac{\gcd(m, n) + 1}{2}.$$

Proof. There are mn total internal faces in $\mathcal{R}_{m+1 \times n+1}$. From the above, $mn - \text{lcm}(m, n)$ of these are part of a billiard path not passing through a corner. Each such basis element uses $2 \text{lcm}(m, n)$ internal faces. Thus there are

$$\frac{mn - \text{lcm}(m, n)}{2 \text{lcm}(m, n)} = \frac{\gcd(m, n) - 1}{2}$$

non-corner billiard paths. Adding back the last billiard path gives the claim. \square

Corollary 4.10. *The rectangle grid graph $\mathcal{R}_{m-1 \times n-1}$ with $(m-1)(n-1)$ even has*

$$|\mathcal{C}_B(\mathcal{R}_{m-1 \times n-1})| = 2^{\frac{\gcd(m, n) - 1}{2}}.$$

Proof. The inner subgraph of $\mathcal{R}_{m+1 \times n+1}$ is $\mathcal{R}_{m-1 \times n-1}$, and the reduced dual graph of $\mathcal{R}_{m \times n}$ is connected. Thus the result follows by Theorem 4.7. \square

Corollary 4.11. *The number of matchings of $\mathcal{R}_{m-1 \times n-1}$ is divisible by*

$$2^{\frac{\gcd(m, n) - 1}{2}}.$$

Furthermore, the number of matchings is odd if and only if m and n are coprime.

Proof. If $m-1$ and $n-1$ are both odd, then $m_G = 0$ and the claim follows. Otherwise, the hypothesis of the previous corollary holds and we may channel 2s (apply Theorem 3.5) to get the divisibility statement. The last claim follows from Proposition 2.2. \square

In the next section, we apply the geometric interpretation of billiard paths for subgraphs of the square lattice to construct an algorithm for finding a path basis for such graphs.

4.3 Finding billiard paths in the square lattice

Let G be a subgraph of the square lattice $\mathbf{Z} \times \mathbf{Z}$ formed by taking a simple cycle and adding all edges and vertices in its interior. Assume we are given the boundary cycle (for instance, as an ordered list of edges). Let P be the set of exterior black vertices in G . From the input data we may determine in $O(|P|)$ time the set of internal faces in G which are incident to the boundary. Our algorithm for constructing a path basis will have complexity $O(|P| \log |P|)$. We shall construct an auxiliary graph A with vertex set P which will have connected components corresponding to billiard paths in G . To begin, given the coordinates (x, y) of a vertex $b \in P$ we compute two indices:

$$\begin{aligned} b^+ &= y - x, \\ b^- &= y + x, \end{aligned}$$

the positive and negative index, respectively. The positive index indicates which line of slope 1 the vertex lies on, and the negative index indicates which line of slope -1 the vertex lies on. Now sort the pairs

$$\{(b^+, b^-) \mid b \in P\}$$

lexicographically. Let b_1 and b_2 be two consecutive vertices in this list. If the vertex b_1 is incident to an internal face to its upper right, then put an edge between b_1 and b_2 in A . Iterate through the list and do this for each consecutive pair.

Next, sort the pairs

$$\{(b^-, b^+) \mid b \in P\}$$

lexicographically. Again, consider consecutive vertices b_1 and b_2 in this list. If b_1 is incident to an internal face to its upper left, then put an edge between b_1 and b_2 in A . Do this for all such consecutive pairs.

After this is complete, the connected components of A will be in correspondence with billiard paths of $\mathcal{B}(G)$. Specifically, a connected component of A corresponds to the minimal nest containing any face of G incident to a vertex in that component. If the number of connected components in A is d , then there are 2^d billiard nests in G , and if the reduced dual of G is connected, then there are 2^{d-1} channels on the black vertices of the inner subgraph of G .

Note that computing the indexes, adding edges to A , and finding the connected components of A takes $O(|P|)$ time, while the sorts take $O(|P| \log |P|)$ time. Thus the algorithm runs in almost linear time.

Example 4.12. Let G be the graph shown in Figure 13. Set A to be the graph on the external black vertices of G , with no edges.





Figure 13: The graph G with external black vertices labeled by their indexes in the form (b^+, b^-) . One of the billiard paths is shown in red.

We list out the index pairs in lexicographical order:

(b^+, b^-)	$(-4, 4)$	$(-4, 6)$	$(-2, 2)$	$(-2, 8)$	$(0, 0)$	$(0, 6)$	$(2, 2)$	$(2, 8)$	$(4, 4)$	$(4, 6)$
upper right face?	y	n	y	n	y	n	y	n	y	n


We connect each vertex with a y to the next vertex on the list. The following figure shows A after this step is completed.



The colors here mean nothing at the moment, but once we finish adding edges they will indicate the connected components of A . For the second phase, we reorder the index pairs, this time lexicographically based on (b^-, b^+) .

(b^-, b^+)	(0, 0)	(2, -2)	(2, 2)	(4, -4)	(4, 4)	(6, -4)	(6, 0)	(6, 4)	(8, -2)	(8, 2)
upper left face?	n	y	n	y	n	y	y	n	n	n

Again we connect each vertex with a y to the following vertex to complete the construction of A . The result is shown below, with connected components displayed in different colors.



There are two connected components of A , corresponding to the two billiard paths in $\mathcal{B}(G)$. Thus there are $2^{2-1} = 2$ channels on the black vertices of the inner subgraph G' of G . \diamond




Figure 14: The inner subgraph G' formed by deleting the external vertices of G , with its nonzero channel highlighted. Either of the billiard nests shown in the previous figure constructs this channel.

Remark 4.13. Let G be a subgraph of the square lattice formed by taking a simple cycle and adding all edges and vertices in its interior. A variant of the algorithm above may be used to compute the path basis for an outer completion of G (rather than for G itself, as done above). The path basis for an outer completion controls the channels (and therefore parity information) in G , rather than in the inner subgraph G' . This alteration then lets us understand the channels of graphs that do not have an outer completion fitting in the square lattice.

To compute the path basis for an outer completion H of G , with boundary cycle ∂H , make the following changes:

- The auxiliary graph A now has vertex set given by the internal faces of H that are incident to vertices in P (the exterior black vertices of G).
- Start off by connecting faces in A that are incident to the same black vertex in ∂H , as well as every second face surrounding a vertex in P .
- Instead of connecting the vertices of P , we connect the appropriate faces in A .
 - In the (b^+, b^-) table, if a vertex $v \in P$ has a y label, then connect the upper right face incident to v to the lower left face incident to the next vertex in the table.
 - In the (b^-, b^+) table, if a vertex $v \in P$ has a y label, then connect the upper left face incident to v to the lower right face incident to the next vertex in the table.

The connected components of A then are equinumerous with the path basis for H . Since an outer completion for G can be constructed in $O(|P|)$ time and can be chosen to add at most $|P|$ internal faces, the altered algorithm remains $O(|P| \log |P|)$.

5 Combinatorial arguments

In this section we give a combinatorial proof that existence of a nonempty channel in G is necessary and sufficient for m_G to be even (Proposition 2.2). In the course of this proof, we develop conditions under which deleting two adjacent vertices (a “vertex pair”) results in a graph with the same number

of channels. This is called *channel routing* and is described in Lemma 5.6; we apply it to deduce a result about rectangle grid graphs in addition to the main theorem of the section.

In Section 2 we described Lovász’s proof of Proposition 2.2. That argument was of a surprisingly different nature from the proof of Theorem 3.5, which relied upon a Kasteleyn signing and matrix normal forms and which could fail for graphs without a Kasteleyn signing (cf. Remark 3.6). Like Lovász’s argument, the combinatorial results in the remainder of this section also do not need a Kasteleyn signing. We do, however, require that our graph be bipartite. We prove the forward and backward directions of Proposition 2.2 separately in the two subsections.

5.1 Existence of a channel implies m_G even

Our approach for this direction will be to construct a fixed-point free involution on the set of matchings $\mathcal{M}(G)$, which will imply that $|\mathcal{M}(G)| = m_G$ is even. To construct such an involution, we employ a technique called *cycle flipping*. This is a commonly used method to build involutions on perfect matchings and show 2-divisibility results. See for example [1] or [13] to see this applied to graphs with reflective symmetry, or [9] for disjoint unions of two graphs.

Given a perfect matching $\mu \in \mathcal{M}(G)$, the idea is to find a cycle Y of edges in the graph such that every second edge in the cycle is in μ . We may then construct a new edge set μ' by replacing the edges of $\mu \cap Y$ with the edges of $Y - \mu$. Since each vertex in Y is contained in exactly one edge in either case, μ' is also a perfect matching.




Figure 15: Flipping the cycle Y from μ to μ' .

For the map this produces to be an involution, the same cycle has to be identified for both μ and μ' . We will produce this cycle using the following tool.

Definition 5.1. A *pairing function* across a vertex set $C \subseteq V$ is a collection of involutions f_v associated to each $v \in V$ which act on the edge set

$$\{\overline{vv'} \in E \mid v' \in C\}$$

such that $f_v \circ f_v = \text{id}$ and f_v has no fixed points.

The existence of a pairing function across C is a combinatorial realization of the statement that, for each vertex v , the neighborhood $N(v)$ contains an even number of elements from C . Applying this to the definition of a channel gives the following lemma.

Lemma 5.2. *Let G be a graph and let $C \subseteq V$ be any vertex set. Then C is a channel if and only if there exists a pairing function across C .*

We will use a pairing function to trace out a walk such that every second edge lies in our matching. Finiteness of the graph will force this path to eventually enter a cycle with the properties we require.




Figure 16: An illustration of the construction used in the proof of Theorem 5.3. The matching μ is shown in purple, and edges which are swapped by a relevant pairing functions are indicated by green lines. In this case, Theorem 5.3 produces the cycle $S(\mu) = \{e_2, \dots, e_{13}\}$.

Theorem 5.3. *Let G be a bipartite graph with a nonempty channel $C \in \mathcal{C}(G)$. Then to each matching μ of G , one can assign a nonempty set of edges $S(\mu) \subseteq E$ satisfying the following properties.*

- (i) $S(\mu)$ is a simple cycle of even length.
- (ii) Every second edge of $S(\mu)$ lies in μ .
- (iii) $S(\mu)$ depends only on the edges of μ containing a vertex from C .
- (iv) If μ' is a matching satisfying $S(\mu) = \mu \oplus \mu'$ then $S(\mu') = S(\mu)$.

Proof. Fix a vertex $v_0 \in C$ and a pairing function f_\bullet across C . We define a walk

$$e_0 = \overline{v_0 v_1}, \quad e_1 = \overline{v_1 v_2}, \quad e_2 = \overline{v_2 v_3}, \quad \dots$$

as follows. If n is even, take e_n to be the unique edge in μ incident to v_n . If n is odd, take $e_n = f_{v_n}(e_{n-1})$. Since f_\bullet is a pairing function across C , $v_{n-1} \in C$ implies $v_{n+1} \in C$. (An example of this walk is shown in Figure 16.)

Because G is finite, there exist n_0 and $p > 0$ such that $v_{n_0+p} = v_{n_0}$. Take p to be minimal and n_0 to be minimal for that choice of p . Note that p must be even since G is bipartite. Also n_0 must be even: if n_0 were odd, then both e_{n_0-1} and e_{n_0+p-1} would be edges in μ containing $v_{n_0+p} = v_{n_0}$. But this would imply $e_{n_0-1} = e_{n_0+p-1}$ and hence $v_{n_0-1} = v_{n_0+p-1}$, which would contradict minimality of n_0 .

We claim that

$$S(\mu) = \{e_n \mid n_0 \leq n < n_0 + p\}$$

satisfies the desired properties. Property (i) follows by construction as well as minimality and evenness of p ; (ii) is clear from construction. For (iii), the only dependence of the walk on μ involves edges incident to v_n for n even. Since $v_n \in C$ for n even, the cycles produced will be the same for matchings that only differ away from C .

For (iv), let $S(\mu) = \mu \oplus \mu'$, so that

$$\mu' = (\mu \setminus C) \cup (C \setminus \mu) = (\mu \setminus C) \cup \{e_{n_0+1}, e_{n_0+3}, \dots, e_{n_0+p-1}\}.$$

Denote the walk constructed for μ' by

$$e'_0 = \overline{v'_0 v'_1}, \quad e'_1 = \overline{v'_1 v'_2}, \quad e'_2 = \overline{v'_2 v'_3}, \quad \dots$$

Since μ and μ' agree outside of $S(\mu)$, $e_n = e'_n$ for $n < n_0$.

By definition, e'_{n_0} is the unique edge of μ' containing $v'_{n_0} = v_{n_0} = v_{n_0+p}$, so $e'_{n_0} = e_{n_0+p-1}$ and hence $v'_{n_0+1} = v_{n_0+p-1}$. Then

$$e'_{n_0+1} = f_{v'_{n_0+1}}(e'_{n_0}) = f_{v_{n_0+p-1}}(e_{n_0+p-1}) = e_{n_0+p-2}$$

and hence $v'_{n_0+2} = v_{n_0+p-2}$. Repeating this argument, we find that $e'_{n_0+k} = e_{n_0+p-1-k}$ for $0 \leq k < p$. In other words, the walk obtained for μ' is the same as that for μ except that the cycle $S(\mu)$ is traversed in the reverse direction. It follows that $S(\mu) = S(\mu')$, as desired. \square

Corollary 5.4. *Let G be a bipartite graph. If $\mathcal{C}(G)$ contains a nonempty channel, then m_G is even.*

Proof. If $C \in \mathcal{C}(G)$ is nonempty, then we claim the map on matchings of G given by

$$\mu \mapsto \mu' := \mu \oplus S(\mu)$$

is an involution with no fixed points. By our discussion at the start of the section, we just need to show that $S(\mu') = S(\mu)$. This follows directly from Theorem 5.3(iv). \square

Remark 5.5. It is interesting (and rather inconvenient) to note that the action of channels on matchings defined above cannot in general be extended to a group action of $\mathcal{C}(G)$ or $\mathcal{C}_B(G)$ since, for instance, the action of two distinct channels need not commute. Such a group action would be a very useful combinatorial tool. We give some thoughts on this at the end of the paper.

5.2 Even m_G implies existence of a channel

Proving the converse statement will take some different machinery, which will turn out to have more general applications. The following lemma describes how channels are affected by removal of an edge. We restrict to bipartite G for clarity—one can generalize the argument to all graphs if care is taken about how vertices can appear in channels, but for brevity we will not do so.

Suppose that G has an edge $e = \overline{bw}$ such that the edge-deleted graph $G - e$ has a channel containing b . The following lemma states that G and the vertex-deleted graph $G - \{b, w\}$ have the same number of channels on their black vertices. (Recall that $\mathcal{C}_B(G)$ denotes the set of channels containing only black vertices in a bipartite graph G .)

Lemma 5.6 (Channel Routing Lemma). *Let $G = (V, E)$ be a bipartite graph, and fix an edge $e = \overline{bw}$ (with vertices the corresponding colors). Define the subgraphs*

$$G^e = G - e$$

and

$$G' = G - \{b, w\}.$$

Then the following statements hold:

- (i) The channels of $\mathcal{C}_B(G)$ not containing b are exactly those of $\mathcal{C}_B(G^e)$ not containing b .
(ii) If there is a channel $B \in \mathcal{C}_B(G^e)$ such that $b \in B$, then there is a bijection

$$\mathcal{C}_B(G) \longleftrightarrow \mathcal{C}_B(G')$$

preserving channels which do not contain b .

Proof.

- (i) Removing/adding the edge e can only change $N(v) \cap C$ for a white vertex v if $v = w$ and $b \in C$. Thus it cannot affect any channel on black vertices that does not contain b .
(ii) We construct a bijection

$$\mathcal{C}_B(G) \longleftrightarrow \mathcal{C}_B(G')$$

by sending channels in $\mathcal{C}_B(G)$ according to

$$f: C \mapsto \begin{cases} C & \text{if } b \notin C, \\ B \oplus C & \text{if } b \in C. \end{cases}$$

This map is well-defined: for any white vertex $v \neq w$, $N(v)$ is the same in both G and G^e . Hence $|N(v) \cap C|$, $|N(v) \cap B|$, and $|N(v) \cap (B \oplus C)|$ are all even by the evenness constraint of channels and properties of symmetric difference, so $|N(v) \cap f(C)|$ is even. Since $f(C)$ never contains b , removing vertices b and w does not affect $|N(v) \cap f(C)|$, so $f(C)$ is a channel in $\mathcal{C}_B(G')$.

We define the inverse map similarly. For a channel C in $\mathcal{C}_B(G')$, define

$$g: C \mapsto \begin{cases} C & \text{if } |N(w) \cap C| \text{ is even,} \\ B \oplus C & \text{if } |N(w) \cap C| \text{ is odd,} \end{cases}$$

where $N(w)$ refers to the neighborhood of w in G . This map is again well-defined. To see this, note that at any white vertex $v \neq w$, $|N(v) \cap g(C)|$ is even as before. At w , if $g(C) = C$, then $|N(w) \cap g(C)| = |N(w) \cap C|$ is even. Otherwise, $|N(w) \cap C|$ is odd, but $|N(w) \cap B|$ is also odd since $N(w)$ differs from the neighborhood of w in G^e only in b , which lies in B . Thus again $|N(w) \cap g(C)| = |N(w) \cap (B \oplus C)|$ is even.

Finally, we verify that these maps are inverses. Any channel on black vertices that does not contain b satisfies the first condition in both definitions, so f and g act on them as the identity. For channels C in $\mathcal{C}_B(G)$ that do contain b , the set $B \oplus C$ has odd intersection with $N(w)$ since B does, and therefore $g \circ f$ is the identity. Channels in $\mathcal{C}_B(G')$ do not contain b , so those that have odd intersection with $N(w)$ map to a channel $B \oplus C$ which contains b . Thus $f \circ g$ is the identity as well, and the claim follows. \square

Channel routing is a versatile tool. To begin with, let us use it to prove constructively the claim titling this section.

Theorem 5.7. *If a bipartite graph G has an even number of matchings, then it has a nonempty channel.*

Proof. Assume without loss of generality that G has at least as many black vertices as white vertices. We will show that $\mathcal{C}_B(G) \neq 0$. First note that G is nonempty since the empty graph has one matching, which is odd. We proceed by induction on $|V| + |E|$. If every white vertex of G has even degree, then we can take the black vertices of V to be our channel. Otherwise, some white vertex w has odd degree. Then for some edge $e = \overline{bw}$, the subgraph $G - \{b, w\}$ has an even number of matchings, since otherwise, by Proposition 2.1,

$$m_G = \sum_{b: \overline{bw} \in E} m_{G - \{b, w\}}$$

would be the sum of an odd number of odd numbers, which would be odd, a contradiction. Define $G^e := G - e$ and $G' := G - \{b, w\}$. Now, because $m_{G'}$ is even, so is

$$m_G - m_{G'} = m_{G^e}.$$

Note G' cannot be the empty graph since that would imply G is a single edge, which would have an odd number of matchings, and likewise G^e is also nonempty. Thus by the inductive hypothesis both $\mathcal{C}_B(G')$ and $\mathcal{C}_B(G^e)$ have nonempty channels.

Let $B \in \mathcal{C}_B(G^e)$ be a nonempty channel. If $b \notin B$, then channel routing (Lemma 5.6(i)) implies B is a channel of G . Otherwise, $b \in B$. Let $C \in \mathcal{C}_B(G')$ be a nonempty channel. Then applying the map

$$g: C \mapsto \begin{cases} C & \text{if } |N(w) \cap C| \text{ is even} \\ B \oplus C & \text{if } |N(w) \cap C| \text{ is odd} \end{cases}$$

described in channel routing (Lemma 5.6(ii)) gives us a nonempty channel in $\mathcal{C}_B(G)$. In either case, $\mathcal{C}_B(G) \neq 0$. \square

Channel routing allows us to remove one vertex pair at a time from our graph while keeping track of the available channels. This is particularly useful when we have a channel in G^e containing b so that condition (ii) of channel routing holds. Let us see how this can be used for subgraphs of the square lattice.

Example 5.8. Refer to Figure 17. Let G be a subgraph of the square lattice such that each internal face is a unit square. Pick a diagonal of G that starts at a corner vertex v and ends at an opposing side vertex b , as in Figure 17a.


Let $e = \overline{bw}$ be the unique edge containing b that forms an obtuse angle with the diagonal. (If there is more than one such edge, then we are not in a situation where the method of this example applies.) Then the graph $G^e = G - e$ has a channel B given by the vertices on the diagonal between v and b .

In particular, $b \in B$. Thus, with $G' := G - \{b, w\}$, channel routing implies that


$$|\mathcal{C}_B(G)| = |\mathcal{C}_B(G')|.$$

This implies, for instance, that m_G and $m_{G'}$ have the same parity (by Theorem 3.5). \diamond


In some cases, repeated application of channel routing can reduce our graph to one with known properties. Rather than deleting each vertex pair individually and examining the channels of each intermediate graph, the following theorem allows us to check an analog of the channel routing condition on a single graph to remove all of the vertex pairs at once.



(a) The graph G with a diagonal selected.



(b) The graph G^e . The vertices on the diagonal form a channel for this graph.



(c) The graph G' . Channel routing implies this graph has the same number of channels as G .

Figure 17: Given a diagonal in a subgraph of the square lattice, we may apply channel routing to the edge e containing the last vertex b and meeting the diagonal at an obtuse angle.

Theorem 5.9. Let G be a bipartite graph with n vertex disjoint edges $e_1 = \overline{b_1 w_1}, \dots, e_n = \overline{b_n w_n}$ selected. Set

$$G^e = G - \{e_1, \dots, e_n\},$$

$$G' = G - \{b_1, w_1, \dots, b_n, w_n\}.$$

Suppose that, for all i , there exists a channel $B_i \in \mathcal{C}_B(G^e)$ such that

$$B_i \cap \{b_1, \dots, b_n\} \subseteq \{b_i\}.$$

Then m_G and $m_{G'}$ have the same parity. If equality holds for all i , then $|\mathcal{C}_B(G)| = |\mathcal{C}_B(G')|$.

Proof. If any nonempty channel in $\mathcal{C}_B(G^e)$ uses none of b_1, \dots, b_n , then it is also a channel in G and G' , so both m_G and $m_{G'}$ are even. Otherwise there exist channels B_1, \dots, B_n in $\mathcal{C}_B(G^e)$ such that $B_i \cap \{b_1, \dots, b_n\} = \{b_i\}$ for all i .

Since $b_1 \in B_1$, channel routing (Lemma 5.6(ii)) implies $|\mathcal{C}_B(G)| = |\mathcal{C}_B(G - \{b_1, w_1\})|$. Since $b_1 \notin B_i$ for $i > 1$, the channels B_2, \dots, B_n are also channels in $\mathcal{C}_B(G - \{b_1, w_1\})$. Thus we may proceed by induction to find $|\mathcal{C}_B(G)| = |\mathcal{C}_B(G')|$, as desired. \square

As one consequence, we can use Theorem 5.9 to determine the parity of the number of domino tilings of a rectangle grid graph.




Figure 18: The first two diagrams show G^e and its associated channels in orange. By Theorem 5.9, the number of matchings of G' (the third figure) has the same parity as that of the original graph $G = \mathcal{R}_{4 \times 7}$. Any matching of G' will use the purple edges, so this result continues to hold if we remove those vertex pairs.

Proposition 5.10. *The $m \times n$ rectangular grid graph has an odd number of matchings if and only if $\gcd(m + 1, n + 1) = 1$.*

Proof. If $m = n = 0$, then $m_G = 1$. If $m = n > 0$, then m_G is even, either by Proposition 3.8 for n even or since $m_G = 0$ for n odd. Thus we may assume without loss of generality that $m < n$.

Refer to Figure 18. Declare the lower left vertex of $\mathcal{R}_{m \times n}$ to be black. Let e_1, \dots, e_r (where $r = \lceil m/2 \rceil$) be the edges between the black vertices in the m th column and the white vertices in the $(m + 1)$ th column. Set $G^e = G - \{e_1, \dots, e_r\}$. The channels constructed in Proposition 3.8 give r independent channels in $\mathcal{C}_B(\mathcal{R}_{m \times m})$. These are also valid channels of G^e which satisfy the hypothesis of Theorem 5.9. Thus by that theorem, we may remove black vertices in column m and white vertices in column $m + 1$ while preserving the parity of m_G . We are left with G' , a graph consisting of two rectangle grid graphs connected by bridges as in the third row of Figure 18. Since there are an equal number of white and black vertices in each rectangle, any matching of G' must use the middle edge on each bridge. Thus we may remove the rest of the vertices in those columns without changing the number of matchings. The resulting graph will be the disjoint union of a $m \times (m - 1)$ rectangle and a $m \times (n - m - 1)$ rectangle. The result then follows by induction on the size of the rectangle, since

$$\gcd(m + 1, n - m - 1 + 1) = \gcd(m + 1, n + 1) \text{ and } \gcd(m + 1, m - 1 + 1) = 1. \quad \square$$

6 Graph operations (or, how to dig channels)


Lemma 5.6 shows that in some cases deleting a vertex pair preserves the number of channels in a graph. In this section, we describe a set of local graph moves which unconditionally preserve channels. In many cases, these operations will allow us to compute the number of channels by reducing to a graph with no edges.

The nature of these graph operations requires that we allow multiple edges between a pair of vertices. We continue to use $N(v)$ to denote the neighborhood of v ; however, it may now be a multiset in which a vertex appears with multiplicity equal to the number of edges connecting it to v . In this case, for a channel C , $N(v) \cap C$ denotes the sub-multiset of $N(v)$ consisting of all vertices which appear in C (with the same multiplicity as they appear in $N(v)$). All other set operations that appear involve standard sets.


We begin by introducing our operations of interest.

6.1 Channel-preserving moves


A *2-valent vertex contraction* may be applied to any vertex v of degree two that is adjacent to distinct vertices v_1, v_2 . The resulting graph is formed by contracting the edges incident to v and deleting self-loops if they occur.



A *doubled edge deletion* may be applied to any pair of edges e_1, e_2 that share the same endpoints. This operation removes e_1 and e_2 from the graph.



A *forced vertex pair removal* may be applied to distinct adjacent vertices v_1, v_2 such that v_1 has degree one. The resulting graph is formed by removing v_1, v_2 and all edges incident to these vertices.



Definition 6.1. A *channel-preserving move* is one of:

- (VC) 2-valent vertex contraction,
- (ED) doubled edge deletion, or
- (FV) forced vertex pair removal.

As the name suggests, applying a channel-preserving move to a graph preserves the number of channels in that graph. In the following, we write $n \equiv_2 m$ to mean $n - m$ is even.

Lemma 6.2. *Let G be a graph and let G' be the result of applying a channel-preserving move to G . Then*

$$|\mathcal{C}(G)| = |\mathcal{C}(G')|.$$

If additionally G is bipartite, then

$$|\mathcal{C}_B(G)| = |\mathcal{C}_B(G')| \text{ and } |\mathcal{C}_W(G)| = |\mathcal{C}_W(G')|.$$

Proof. First, ED moves clearly preserve the parity of $|N(v) \cap C|$ for all vertices v and vertex sets C , implying the claim. We will show the result for VC moves; the argument for FV moves is similar.

Let v be a vertex of degree two with adjacent vertices v_1 and v_2 . Call the resulting contracted vertex w in G' . If $C \in \mathcal{C}(G)$, then in order for the evenness condition to hold at v , it follows that $v_1 \in C$ if and only if $v_2 \in C$. Thus we may define

$$C' = C \setminus \{v_1, v_2\} \cup W, \quad \text{where } W = \begin{cases} \{w\} & \text{if } v_1, v_2 \in C, \\ \emptyset & \text{otherwise.} \end{cases}$$

This preserves the neighborhood size of all unchanged vertices by replacing any occurrence of v_1 or v_2 in a neighborhood with w . Thus evenness holds everywhere except possibly at w . To see that $|N(w) \cap C'|$ is even, note that

$$|N(w) \cap C'| = |N(v_1) \cap C| + |N(v_2) \cap C| - 2|\{v\} \cap C| - k|\{v_1, v_2\} \cap C|,$$

where k is the number of edges between v_1 and v_2 . Since $|\{v_1, v_2\} \cap C| = |N(v) \cap C|$ is even, it follows that

$$|N(w) \cap C'| \equiv_2 |N(v_1) \cap C| + |N(v_2) \cap C|$$

is even and hence $C' \in \mathcal{C}(G')$.

Conversely, if $C' \in \mathcal{C}(G')$, then $|N(w) \cap C'|$ is even. Since

$$|N(w) \cap C'| = |N(v_1) \cap C'| + |N(v_2) \cap C'|,$$

it follows that $|N(v_1) \cap C'| \equiv_2 |N(v_2) \cap C'|$. Again let k denote the number of edges between v_1 and v_2 . We may now define $C = C' \setminus \{w\} \cup V_1 \cup V_2$, where

$$V_1 = \begin{cases} \{v_1, v_2\} & \text{if } w \in C', \\ \emptyset & \text{otherwise,} \end{cases} \quad \text{and } V_2 = \begin{cases} \{v\} & \text{if } |N(v_1) \cap C'| + k|\{w\} \cap C'| \text{ is odd,} \\ \emptyset & \text{otherwise,} \end{cases}$$

which is the inverse to the map $C \mapsto C'$ described above. Again this preserves the neighborhood size of all unchanged vertices. The definition of C ensures that $|N(v_1) \cap C| \equiv_2 |N(v_2) \cap C|$ is even by adding v to C if necessary. Also $|N(v) \cap C| = 2$ if $w \in C'$, and 0 otherwise. Thus evenness holds at all vertices, and $C \in \mathcal{C}(G)$.

Further, notice that if G is bipartite, then w has the same color as v_1 and v_2 . So if C uses only vertices of a single color, then C' only uses vertices of the same color. The converse also holds, since for singly colored channels, at most one of $|N(v_1) \cap C'| \neq 0$ and $w \in C'$ can hold. Hence we have a bijection $\mathcal{C}(G) \leftrightarrow \mathcal{C}(G')$ that descends to \mathcal{C}_B and \mathcal{C}_W for bipartite graphs. \square

We will call a graph *reducible* if it can be reduced to a set of degree 0 vertices using only channel-preserving moves.

Theorem 6.3. *Let G be a reducible graph. Then the number of degree 0 vertices remaining after G has been fully reduced is the dimension of $\mathcal{C}(G)$. In particular, this number is independent of the choice of channel-preserving moves used to reduce the graph.*

Proof. Since channel-preserving moves preserve channels, we just need to show that a set of n vertices of degree 0 has 2^n channels. This is clear, since any subset of these vertices is a valid channel. \square

Example 6.4. The graph G shown in Figure 7 is a reducible graph. Figure 19 shows a possible sequence of channel-preserving moves. Because G reduces to two vertices of degree 0, G must have 2^2 channels. This is indeed the case; the three nonempty channels are shown in Figure 7. \diamond




Figure 19: A reduction of the graph shown in Figure 7.

Example 6.5. Not all graphs are reducible. Figure 20 shows a planar bipartite graph that admits no channel-preserving moves. \diamond




Figure 20

As this example indicates, it is not clear at first if reducibility occurs often enough to be useful—we would like to have a simple structural property that will imply reducibility. The following lemma will be useful for identifying potential VC or FV moves. Note that the degree of a face in a planar graph is its degree as a vertex in the dual graph.

Lemma 6.6 (Corner identification). *Let G be a connected planar graph with at least two vertices such that each internal vertex and internal face have degree at least 4. Let b be the number of external vertices in G . Then the average degree of the external vertices in G is at most*

$$3 - \frac{4}{b}.$$

In particular, there is an external vertex of degree less than 3.

Proof. By assumption, the degree of each vertex is at least one. We will use the following notation:

- b = the number of external vertices,
- i = the number of internal vertices,
- D = the total degree of all external vertices,
- e = the number of edges,
- f = the number of internal faces.

By planarity of G , we have that

$$b + i - e + f = 1. \tag{6.1}$$

Because the total degree of all internal and external vertices is $2e$, and the degree of each internal vertex is at least 4, it follows that

$$2e \geq 4i + D.$$

The total degree of all internal and external faces is also $2e$. Since the degree of the external face is the total number of vertices counted with multiplicity as one proceeds through its boundary cycle, this number is at least b . Thus the total degree of all internal faces is at most $2e - b$, and consequently

$$2e - b \geq 4f.$$

Adding the two inequalities and utilizing Euler's formula (6.1) then gives

$$D \leq 3b - 4.$$

Dividing by b reveals the claim. \square

A planar graph is called *inner Eulerian* if all internal vertices have even degree. The following results show the utility of this class of graphs.

Lemma 6.7. *Channel-preserving moves preserve the property of being inner Eulerian.*

Proof. Assume G is inner Eulerian. Let v be a vertex of G of degree 2. Let the degrees of the two vertices v_1 and v_2 adjacent to v be d_1 and d_2 , respectively, and let the number of edges connecting v_1 and v_2 be n . If we perform a VC move on v , then we will be left with a vertex of degree $d_1 + d_2 - 2n - 2$. If $d_1 + d_2$ is even, then we are done. Otherwise, one of the two vertices adjacent to v has odd degree and must therefore be incident to the external face since G is inner Eulerian. Thus the vertex resulting from contracting v is also external and therefore has unconstrained degree. Thus VC moves preserve being inner Eulerian.

Let v now be a vertex of degree one. Then v is external, so the vertex v' adjacent to v is also external. If we apply an FV move to v , then v and v' will be removed. This will change only the degree of vertices adjacent to v' . However, once v' is removed all of these vertices will be adjacent to the external face and therefore have unconstrained degree. Thus FV moves preserve being inner Eulerian.

Finally, ED moves preserve the parity of the degree of every vertex and thus also preserve the property of being inner Eulerian. \square

Theorem 6.8. *Let G be an inner Eulerian bipartite graph. Then G is reducible.*

Proof. By the previous lemma, we are free to perform any channel-preserving move while remaining inner Eulerian. If we can show that it is always possible to perform a channel-preserving move on an inner Eulerian bipartite graph G with at least one edge, then the result will follow by induction on the number of edges.

Without loss of generality, we may assume that G is connected, so every vertex has degree at least one. If any internal vertex has degree 2, then we can perform a VC move. Otherwise every internal vertex has degree at least 4 (it must be even as G is inner Eulerian). If any internal face has degree 2, then we can perform an ED move. Otherwise every internal face has degree at least 4 (it must be even as G is bipartite).

Thus if neither a VC or ED move is possible in the graph interior, then by Lemma 6.6, there exists an external vertex of degree 1 or 2 to which an FV or VC move can be applied. \square

6.2 Contracting diagonals

When our graph is a suitably nice subgraph of the square lattice, there is often a useful sequence of channel-preserving moves available called a *diagonal contraction*. Pick a degree 2 vertex v that is a corner of the graph. Then v defines a unique diagonal passing through it, as in the top of Figure 21.

We say that the diagonal is *contractible* if each internal vertex and each internal face it intersects have degree 4. Such a diagonal can be contracted as follows. Consider all vertices on the diagonal from v to w , the last vertex on the diagonal before it reaches the external face. For each such vertex




Figure 21: Contracting the highlighted diagonal by deleting the diagonal vertices and merging the vertices immediately opposite.

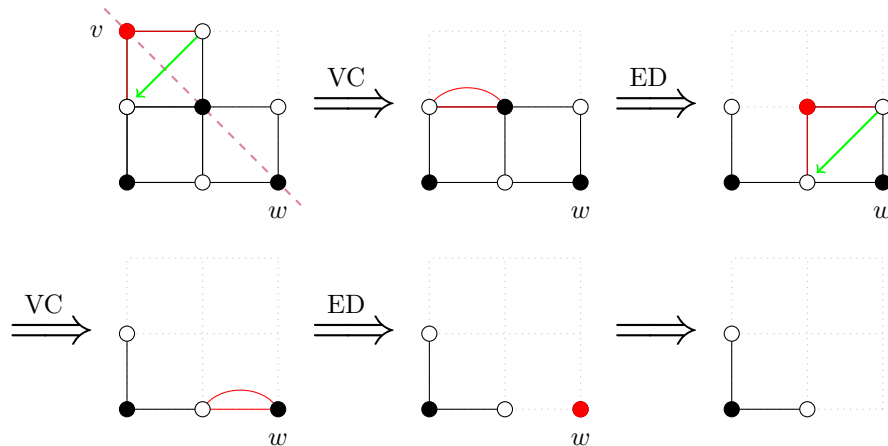


Figure 22: Diagonal contraction (from v to w) that ends on a degree 2 vertex. Removing the degree 0 vertex w removes a basis channel for G .

v_i , delete v_i and combine each neighbor of v_i with its mirror image across the diagonal, as shown in Figure 21. If a vertex is to combine with a missing vertex (denoted by a red \times in the figure), then that vertex is instead deleted.

Theorem 6.9. *Let G be a subgraph of the square lattice. Let G' be the result of applying a diagonal contraction to a contractible diagonal from a black corner vertex v to a vertex w . If w has degree 2, then*

$$|\mathcal{C}_B(G)| = 2|\mathcal{C}_B(G')| \text{ and } |\mathcal{C}_W(G)| = |\mathcal{C}_W(G')|.$$

Otherwise,

$$|\mathcal{C}_B(G)| = |\mathcal{C}_B(G')| \text{ and } |\mathcal{C}_W(G)| = |\mathcal{C}_W(G')|.$$

Proof. We will show that a diagonal contraction move consists of a sequence of channel-preserving moves. If $v = w$, then diagonal contraction is just a VC move. Otherwise, the diagonal passes through an internal face which v shares with exactly one other black vertex v' and two white vertices v_1 and v_2 . Applying a VC move to v combines v , v_1 , and v_2 . Since v_1 and v_2 were both adjacent to v' , there are now two edges between v and v' . Thus we can apply an ED move to this edge pair. Now, if v' has degree 2 at this point, then we may repeat this argument starting with v' .

Otherwise, v' now has degree 1 or 0, so it originally had degree 3 or 2, so it must have been w , the last vertex on the diagonal. We can either apply an FV move to the degree 1 vertex or remove the degree 0 vertex from the graph. (The possible ending scenarios are shown in Figures 22, 23, and 24.) Our resulting graph is the diagonal contraction G' essentially by definition. All of our moves were channel preserving, except for removing the degree 0 vertex w ; this occurs if and only if w had degree 2 in G . Removing a degree 0 black vertex from the graph halves the number of channels on black vertices and preserves the number of channels on white vertices. Thus the claim is shown. \square

Example 6.10. Let us apply diagonal contraction to a well-known class of graphs. The *Aztec diamond* G_n of rank n is a diamond of side length n in the square lattice. The Aztec diamonds G_n for $n = 1, 2$, and 2 are shown in Figure 25.




Figure 23: Diagonal contraction (from v to w) that ends on a degree 3 vertex. In this case we end with an FV move that removes the vertex adjacent to w . Channels are preserved.




Figure 24: Diagonal contraction (from v to w) that ends on a degree 4 vertex. Channels are preserved.





Figure 25: The Aztec diamonds of rank 1, 2, and 3.

We will show that $|\mathcal{C}(G_n)| = 2^{2^n}$ by induction on n . The rank 1 Aztec diamond has 4 channels. For the rank n Aztec diamond with $n > 1$, we perform the following diagonal contractions:




This produces the Aztec diamond of rank $n - 1$. Since both diagonal contractions ended on a vertex of degree 2,

$$|\mathcal{C}(G_n)| = 2^2 |\mathcal{C}(G_{n-1})|.$$

The result follows by induction. Since there are 2^{2n} channels in G_n , by Theorem 3.5 it follows that 2^n divides the number of matchings of G_n . Indeed, it is well-known that G_n has $2^{\binom{n+1}{2}}$ matchings. \diamond

The fact that diagonal contraction preserves channels has some implications that have been noticed previously in the literature. For instance, Tenner [17] proves a Tiling Parity Theorem and uses it to great effect. When stated in our language, the Tiling Parity Theorem is a statement about diagonal contraction for certain diagonals that do not end on a corner. Here we describe a generalization of Tenner's theorem. (In her language, we cover the case $k \geq 3$, though the cases $k < 3$ follow by a similar argument.) Recall that $n \equiv_2 m$ means that $n - m$ is even.






Figure 26: A graph to which the Parity Theorem is applicable. By Theorem 6.11,

$$41 = m_G \equiv_2 m_{G'_e} + 2m_{G'_v} \equiv_2 m_{G'_e} = 11.$$

Theorem 6.11 (Parity Theorem). *Let G be a subgraph of the square lattice and let v_1, v_2, v_3, v_4 be collinear consecutive external vertices such that v_2 and v_3 have degree 3. Let $G_e = G - \overline{v_1 v_2}$ and $G_v = G - \{v_1, v_2\}$ be the edge-deleted and vertex-deleted subgraphs of G , so that v_2 defines a unique diagonal in G_e and v_3 defines a unique diagonal in G_v . If both of these diagonals are contractible, then*

$$m_G \equiv_2 2^{\delta_e} m_{G'_e} + 2^{\delta_v} m_{G'_v},$$

where G'_e and G'_v are the graphs resulting from contracting the diagonals in G_e and G_v respectively, and δ_i is 1 or 0 depending on if the diagonal in G_i ends on a vertex of degree 2 or of degree at least 3, respectively.

Proof. By Proposition 2.1,

$$m_G = m_{G_e} + m_{G_v}.$$

By Theorem 6.9, channels are preserved by diagonal contractions if and only if the diagonal ends on a vertex of degree at least 3. Specifically,

$$|\mathcal{C}(G_e)| = 2^{\delta_e} |\mathcal{C}(G'_e)| \text{ and } |\mathcal{C}(G_v)| = 2^{\delta_v} |\mathcal{C}(G'_v)|.$$

Now by Proposition 2.2, m_{G_i} is even if and only if $|\mathcal{C}(G_i)| > 1$. Thus

$$m_{G_e} \equiv_2 2^{\delta_e} m_{G'_e} \text{ and } m_{G_v} \equiv_2 2^{\delta_v} m_{G'_v}.$$

The claim follows. □

Remark 6.12. In fact, if G has at least one matching, then

$$m_G \equiv_2 m_{G'_e} + m_{G'_v}.$$

This holds because G has a matching only if G has an equal number of black and white vertices, whereas a diagonal contraction that ends on a degree 2 vertex removes unequal numbers of white and black vertices. Thus $\delta_i = 1$ only if $m_{G'_i} = 0$.

We conclude this section by showing that diagonal contraction is always possible for certain subgraphs of the square lattice.

Theorem 6.13. *Let G be a subgraph of the square lattice such that every internal face of G is a unit square and such that every edge bounds an internal face. Then G has a contractible diagonal starting at a degree 2 vertex.*

Proof. Each internal vertex is incident only to internal faces and therefore to exactly four unit squares. Hence all internal faces and vertices have degree 4, so any diagonal in G will be contractible. Any degree 2 vertex in G will be a corner, since by assumption the two incident edges bound a unit square. Thus we just need to show that there is a degree 2 vertex in G . But this follows from Lemma 6.6 since every internal face and vertex have degree four. \square

7 Conclusion

As we have seen, channels provide an effective lower bound on the power of two dividing a matching count. In addition, when there are no nonempty channels, they tell us that the number of matchings is odd. It would be nice to find exact powers of two more generally. This prompts a natural question.

Problem 7.1. *Determine when the number of channels is the exact power of two dividing m_G^2 . Even better, determine how many additional powers of two are carried by each channel more generally. Alternatively, provide a method for determining an upper bound on powers of two dividing m_G .*

For more on how additional powers of two are distributed in the Smith normal form of the Kasteleyn matrix (and therefore among channels), see [10]. Additional powers of two may be associated to a result such as Ciucu's Factorization Theorem (Proposition 1.2). Indeed, graphs where this theorem applies tend to have additional powers of two beyond what channels would predict (for instance the Aztec Diamond in Example 6.10, cf. [2]). One possible route for approaching Problem 7.1 is to study equivalence classes of perfect matchings under the action of channels described in Theorem 5.3. Another route may arise by solving the following problem.

Problem 7.2. *Find a combinatorial proof of Theorem 3.5.*

Since Theorem 3.5 requires the Kasteleyn signing of G , such a proof would likely invoke planarity. As mentioned in Remark 5.5, one possible approach to this is constructing a free action of $\mathcal{C}_B(G)$ on the set of matchings of G . (For the general, non-bipartite case, we would want an action of $\mathcal{C}(G)$ on pairs of matchings.) Because the definition of channels involves neighborhoods of even size, searching for an action that uses properties of Eulerian circuits may yield productive results. A related problem is to construct an action of billiard nests on matchings, for graphs which admit them. This could be more tractable due to the canonical path basis for the space of nests.

The reducible graphs described in Section 6 may provide a tractable entry point to these problems. For such graphs, the problem of constructing an action of channels on matchings reduces to understanding how such an action plays with the channel-preserving moves. Since vertex contraction and forced vertex deletion both preserve matchings, this further reduces to studying the impact of doubled edge deletion on matchings.

Acknowledgments

The first author was supported by a grant from the North Carolina State University Office of Undergraduate Research. The second author was partially supported by grants from the National Science Foundation, DMS-1700302 and CCF-1900460.

References

- [1] Mihai Ciucu, *Enumeration of Perfect Matchings in Graphs with Reflective Symmetry*, Journal of Combinatorial Theory, Series A **77** (January 1997), no. 1, 67–97.
- [2] ———, *Perfect matchings and perfect powers*, Journal of Algebraic Combinatorics **17** (January 2003), no. 3, 335–375.
- [3] Henry Cohn, *2-adic behavior of numbers of domino tilings*, Electronic Journal of Combinatorics **6** (August 1999), no. 1, #R14.
- [4] David Eppstein and Vijay V. Vazirani, *NC Algorithms for Computing a Perfect Matching, the Number of Perfect Matchings, and a Maximum Flow in One-Crossing-Minor-Free Graphs*, The 31st ACM on Symposium on Parallelism in Algorithms and Architectures - SPAA '19, 2019, pp. 23–30.
- [5] Chris Fraser, Thomas Lam, and Ian Le, *From dimers to webs*, Transactions of the American Mathematical Society **371** (January 2019), no. 9, 6087–6124.
- [6] P. W. Kasteleyn, *The statistics of dimers on a lattice. i. the number of dimer arrangements on a quadratic lattice*, Physica **27** (December 1961), no. 12, 1209–1225.
- [7] ———, *Dimer statistics and phase transitions*, Journal of Mathematical Physics **4** (February 1963), no. 2, 287–293.
- [8] Richard W. Kenyon, James G. Propp, and David B. Wilson, *Trees and matchings.*, The Electronic Journal of Combinatorics **7** (2000), no. 1, #R25.
- [9] Eric H. Kuo, *Applications of graphical condensation for enumerating matchings and tilings*, Theoretical Computer Science **319** (2004), no. 1-3, 29–57.
- [10] Greg Kuperberg, *Kasteleyn cokernels*, Electronic Journal of Combinatorics **9** (August 2002), no. 1, #R29.
- [11] Thomas Lam, *Dimers, webs, and positroids*, Journal of the London Mathematical Society **92** (2015), no. 3, 633–656.
- [12] L. Lovász, *Combinatorial problems and exercises, second edition*, Elsevier Science, 1993.
- [13] Lior Pachter, *Combinatorial approaches and conjectures for 2-divisibility problems concerning domino tilings of polyominoes*, Electronic Journal of Combinatorics **4** (1997), #R29.
- [14] Jerome K. Percus, *One more technique for the dimer problem*, Journal of Mathematical Physics **10** (October 1969), no. 10, 1881–1884.
- [15] James Propp, *Enumeration of matchings: Problems and progress*, New Perspectives in Geometric Combinatorics **38** (April 1999), 255–291.
- [16] Richard P. Stanley, *Smith normal form in combinatorics*, Journal of Combinatorial Theory. Series A **144** (November 2016), 476–495.
- [17] Bridget Eileen Tenner, *Tiling parity results and the holey square solution*, Electronic Journal of Combinatorics **11** (2005), no. 2, #R17.
- [18] ———, *Domino tiling congruence modulo 4*, Graphs and Combinatorics **25** (November 2009), no. 4, 625–638.
- [19] Carlos Tomei and Tania Vieira, *The kernel of the adjacency matrix of a rectangular mesh*, Discrete and Computational Geometry **28** (August 2002), no. 3, 411–425.