Matchings in Lattice Graphs

Claire Kenyon[†] Dana Randall[‡] Alistair Sinclair[§]

TR-93-019

March 1993

Abstract

We study the problem of counting the number of matchings of given cardinality in a d-dimensional rectangular lattice. This problem arises in several models in statistical physics, including monomer-dimer systems and cell-cluster theory. A classical algorithm due to Fisher, Kasteleyn and Temperley counts perfect matchings exactly in two dimensions, but is not applicable in higher dimensions and does not allow one to count matchings of arbitrary cardinality. In this paper, we present the first efficient approximation algorithms for counting matchings of arbitrary cardinality in (i) d-dimensional "periodic" lattices (i.e., with wrap-around edges) in any fixed dimension d; and (ii) two-dimensional lattices with "fixed boundary conditions" (i.e., no wrap-around edges). Our technique generalizes to approximately counting matchings in any bipartite graph that is the Cayley graph of some finite group.

[†]CNRS, Ecole Normale Supérieure de Lyon, France. Part of this work was done while this author was visiting ICSI, Berkeley. E-mail: `kenyon@lip.ens-lyon.fr`.

[‡]Department of Computer Science, University of California at Berkeley. Supported in part by an AT&T PhD Fellowship and NSF grant CCR88-13632. E-mail: `randall@cs.berkeley.edu`.

[§]University of Edinburgh and International Computer Science Institute, Berkeley. Supported in part by grant GR/F 90363 of the UK Science and Engineering Research Council. E-mail: `sinclair@icsi.berkeley.edu`.

1 Summary

1.1 Background and motivation

This paper is concerned with the following computational problem: given a finite lattice graph in some fixed number of dimensions, and some number of dominoes, determine the number of ways of placing dominoes on the edges of the graph so that no two dominoes overlap at a vertex. Equivalently, we can think of dominoes as covering a pair of adjacent squares (cubes) in the dual lattice.

This is a classical problem in statistical physics, first introduced by Fowler and Rushbrooke in 1937 [3], and is the earliest example of a large class of problems concerned with computing the number of non-overlapping arrangements of figures of various shapes on a lattice (see, e.g., [11, 16] for a survey). The problem arises in several physical models. For example, in two dimensions the lattice represents the surface of a crystal and the dominoes diatomic molecules (or *dimers*), and the number of domino arrangements is the number of ways in which a given number of dimers can attach themselves onto the surface; from this information, most of the thermodynamical properties of the system can be computed. In three dimensions, the same problem occurs in the theory of mixtures of molecules of different sizes and in the cell-cluster theory of the liquid state. For further background information, see [4, 11] and the references given there.

The problem also has inherent combinatorial interest: clearly a domino arrangement is simply a matching, so we are actually being asked for the number of matchings of specified cardinality in the lattice graph. Counting matchings is a central problem in computer science and has received much attention since the seminal work of Valiant [15], who proved that it is #P-complete for general graphs. The enumeration of perfect matchings (where the dominoes are required to completely cover the graph) is equivalent to computing the permanent of a 0-1 matrix, a long-studied problem in its own right [12]. This paper investigates the complexity of these problems in the important special case of lattice graphs.

1.2 Previous work

We are interested in two classes of lattice graphs: the first class are graphs with *fixed boundary conditions*, in which the lattice is not perfectly regular but has distinguished boundary vertices. Thus, we consider the d-dimensional *rectangular* (or cartesian) lattice $L(n,d)$, where the vertices are the n^d integer lattice points in $[1,n]^d$, and two points x,y are connected by an edge iff they are unit distance apart. The second class is graphs with *periodic boundary conditions*, in which the lattice includes wrap-around edges to make it toroidal; that is, we augment $L(n,d)$ with an edge between $(x_1,\ldots,x_{i-1},n,x_{i+1},\ldots,x_d)$ and $(x_1,\ldots,x_{i-1},1,x_{i+1},\ldots,x_d)$, for each i. We will write $\widetilde{L}(n,d)$ for the periodic lattice.

We will always assume that n is even, so that both $L(n,d)$ and $\widetilde{L}(n,d)$ always contain a perfect matching of size $m=\frac{1}{2}n^d$. For any graph, we let $\mathcal{M}$ be the set of perfect matchings and $\mathcal{N}$ the set of *near-perfect* matchings, i.e., matchings with exactly two unmatched vertices. We refer to the set of unmatched vertices as *holes*, and let $\mathcal{N}(u,v)$ be the set of near-perfect matchings with holes u and v.

A beautiful classical result due to Fisher, Kasteleyn and Temperley [1, 9, 14] provides an analytic (closed-form) expression for the number of perfect matchings in the two-dimensional lattice $L(n,2)$, for any n. This is actually based on a more general result which shows that the number of perfect matchings in any *planar* graph can be computed in polynomial time by evaluating a suitable determinant. In fact, the technique extends to counting perfect matchings in any family of graphs with fixed genus [10]. Thus, in particular, the periodic two-dimensional lattice $\widetilde{L}(n,2)$, which has genus 1 for all n, can also be handled.

However, the technique breaks down completely in three or more dimensions, about which very little is known. Moreover, it yields no information about the number of matchings of cardinality smaller than m, even in the planar case. In fact, in both of these cases the problem is known to be #P-complete for general graphs: counting perfect matchings in a graph [15], and matchings of specified cardinality in a planar graph [5] are both #P-complete problems. Thus it is extremely unlikely that the Fisher, Kasteleyn and Temperley technique for fixed genus graphs can be extended to answer these questions for all lattices.

In recent years, there have been considerable advances in the design of efficient *approximation* algorithms for counting matchings. A widely accepted notion of efficient approximability for combinatorial enumeration problems is the existence of a fully-polynomial randomized approximation scheme (see, e.g., [8, 13]). The definition is as follows:

Definition. A *fully-polynomial randomized approximation scheme* (*fpras*) for a non-negative real-valued function f is a probabilistic algorithm which, on input x and $\epsilon > 0$, outputs a number $\tilde{f}(x)$ such that $\Pr\{f(x)(1+\epsilon)^{-1} \le \tilde{f}(x) \le f(x)(1+\epsilon)\} \ge 1-\delta$, and runs in time polynomial in $|x|$, ϵ^{-1}, and $\lg \delta^{-1}$. □

Jerrum and Sinclair [6] showed the existence of a fpras for counting matchings of any cardinality up to $(1-\alpha)k$ (for constant $0 < \alpha < 1$) in arbitrary graphs containing a k-matching. For larger matchings (including perfect matchings), an extra condition is required: there exists a fpras for counting matchings of every cardinality in any family of $2m$-vertex graphs satisfying

$$\frac{|\mathcal{N}|}{|\mathcal{M}|} \le q(m), \tag{1}$$

for some fixed polynomial q. This condition expresses the fact that the number of near-perfect matchings should not exceed the number of perfect matchings by too much. (Note that in any graph with $2m$ vertices, $|\mathcal{N}| \ge m|\mathcal{M}|$, since the removal of any edge from a perfect matching yields a unique near-perfect matching.)

Condition (1) is known to hold for all dense graphs, all graphs with sufficiently good expansion properties, and almost every random bipartite graph in the $\mathcal{B}(n,p)$ model for any density p above the threshold value for existence of a perfect matching [6]. However, these results shed no light on the special case of lattice graphs since (1) is not known to hold for them (even in two dimensions). Moreover, the technique used to establish (1) in the above cases is not applicable here since it involves constructing short augmenting paths for near-perfect matchings; such paths do not exist in lattice graphs, which have large diameter. Our work consists in exhibiting a new technique that allows condition (1) to be verified for the lattices we are interested in.

1.3 Results

The main contribution of this paper is to establish the existence of polynomial time approximation algorithms for counting arrangements of any given number of dominoes (i.e., matchings of any given cardinality) in periodic lattices of any dimension. More precisely, we prove the following result:

Theorem 1 *There exists a fpras for counting matchings of any cardinality in the d-dimensional periodic lattice $\widetilde{L}(n,d)$ for any fixed dimension d.*

This theorem extends previously known results in two ways. First, and most significantly, we are now able to count perfect matchings in lattices of dimension greater than two. This of course includes the three-dimensional case, which is of greatest physical interest. Second, we are able to count matchings with any specified number of holes, a problem which was not approachable by the results of Fisher, Kasteleyn and Temperley, even in two dimensions.

In the planar case we can also handle lattices with fixed boundaries, as the following theorem states.

Theorem 2 *There exists a fpras for counting matchings of any cardinality in the two-dimensional lattice with fixed boundaries $L(n,2)$.*

This theorem extends the results of Fisher, Kasteleyn and Temperley by allowing us to count matchings with any specified number of holes. Given $2c$ fixed holes, for any constant c, the Fisher, Kasteleyn and Temperley technique can be used to count the number of perfect matchings in the graph formed by removing the holes, since this graph is still planar. When c is small, we could use this method to count the total number of matchings with $2c$ holes by considering all possible positions for the holes; but this approach fails if c is allowed to grow with n, and is inefficient even for quite small fixed values of c. Our results let us count the number of matchings with holes directly, for any number of holes.

Finally, we can extend our results to any bipartite graph which is the Cayley graph of some finite group. This includes other commonly studied lattices, such as the hexagonal lattice with periodic boundary conditions. More precisely:

Theorem 3 *There exists a fpras for counting matchings of any cardinality in any bipartite graph which is the Cayley graph of some finite group.*

The proofs of the above theorems, presented in the next two sections, are elementary and rely on a novel translation technique: the strong symmetry properties of the lattice (and of arbitrary Cayley graphs) allow any matching to be translated, which in turn permits the symmetry to be broken. We conjecture that this technique may shed more light on other quantities related to matchings on the lattice, such as the correlation between holes at two specified vertices.

2 Rectangular Lattices

The algorithms presented in this paper all rely on the following result, which says that the number of matchings of any cardinality can be approximated efficiently provided $|\mathcal{N}|$ is not too much larger than $|\mathcal{M}|$. The approximation algorithm referred to in the theorem is based on simulation of a rapidly mixing Markov chain whose states are matchings [6, Theorem 5.3].

Theorem 4 (Jerrum and Sinclair) *There exists a fpras for counting the number of perfect matchings, $|\mathcal{M}|$, in any family of $2m$-vertex graphs that satisfies $|\mathcal{N}|/|\mathcal{M}| \le q(m)$, for some fixed polynomial q.*

The remark following this result in [6] points out that this polynomial relationship between near-perfect and perfect matchings also allows one to approximately count matchings of arbitrary cardinality in polynomial time.

In all of the cases which follow, we consider $2m$-vertex bipartite graphs where the m vertices on one side of the bipartition are colored white and those on the other side are colored black. Thus, in any near-perfect matching, one hole is white and the other black. (In the case of the two-dimensional lattice, this coloring corresponds to the usual black and white coloring of the checker-board squares which form the dual graph.)

The technique that we use in our proofs relies on the structure of the union of two matchings in a graph. Consider the subgraph C consisting of the union of the edges in two perfect matchings M_1 and M_2. If we color the edges from M_1 red and those from M_2 blue, we find that every vertex is adjacent to exactly one red edge and one blue edge, so C is the union of even-length cycles, each of which alternates colors. (Some of these cycles may be trivial, consisting of a single edge colored both red and blue.) Clearly the converse is also true, i.e., any covering of the graph with even-length cycles which alternate colors defines two perfect matchings: the set of red edges and the set of blue edges.

Similarly, suppose we have two near-perfect matchings, N_1 with holes u and v, and N_2 with holes u' and v', where u, u', v and v' are distinct vertices. Then in the subgraph C defined by the union of the red edges N_1 and the blue edges N_2, vertices u, u', v and v' all have degree one and all other vertices have degree two. So C consists of even-length alternating cycles, plus two alternating paths whose endpoints are u, u', v and v'. Moreover, either both of these paths have even length or both have odd length. See figure 1.

Our proofs rely on the observation that, if u' is a neighbor of u and v' is a neighbor of v, then by augmenting C with edges (u, u') and (v, v'), we can ensure that every vertex has degree two. When the graph is bipartite, the resulting subgraph must consist solely of even-length cycles, and therefore the cycle containing u and u' must also contain v and v'. By recoloring some of the edges on this new cycle, we can force it to alternate colors so that the cycle cover defines two perfect matchings. This allows us to define a mapping from the set of pairs $\mathcal{N}(u, v) \times \mathcal{N}(u', v')$ to the set of pairs $\mathcal{M} \times \mathcal{M}$ that is *injective*, which in turn, by virtue of the symmetry properties of the lattice, implies that $|\mathcal{N}|$ is not much larger than $|\mathcal{M}|$.

We now recall the statement of the first theorem before proving it:

Theorem 1 *There exists a fpras for counting matchings of any cardinality in the d-dimensional periodic lattice $\widetilde{L}(n, d)$ for any fixed dimension d.*

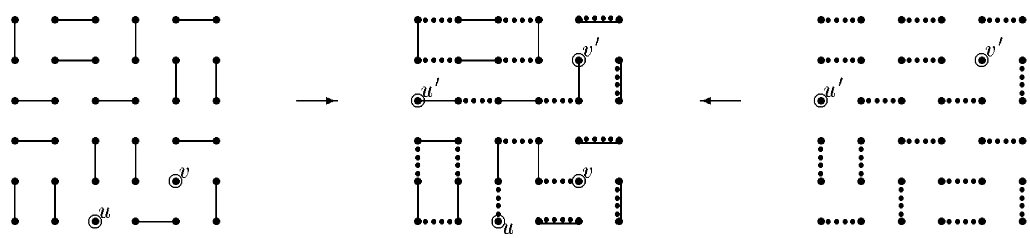

Figure 1: The union of two near-perfect matchings.

Proof. Let $\mathcal{M}$ and $\mathcal{N}$ be the sets of perfect and near-perfect matchings respectively in $\widetilde{L}(n,d)$. By Theorem 4, it suffices to show that $|\mathcal{N}|/|\mathcal{M}| \leq q(m)$ for some polynomial q, where $m = \frac{1}{2}n^d$ is the number of edges in a perfect matching. First we fix two holes, u and v. Let u' be the neighbor one to the right of u, i.e., $u' = u + (1,0,...,0) \bmod n$. Similarly, let v' be the neighbor one to the right of v.

We proceed to construct an injection ϕ from $\mathcal{N}(u,v) \times \mathcal{N}(u',v')$ into $\mathcal{M} \times \mathcal{M}$. To do this, let $N_1 \in \mathcal{N}(u,v)$ and $N_2 \in \mathcal{N}(u',v')$, and consider the subgraph C of $\widetilde{L}(n,d)$ defined by the union of red edges N_1, blue edges N_2 and *special edges* (u,u') and (v,v'). If we color the special edges red, then u' and v' are each adjacent to two red edges, and every other vertex is adjacent to one edge of each color; if we now flip the colors of the edges along one of the paths from u' to v', every vertex will be adjacent to exactly one edge of each color. To avoid ambiguity, we choose the path from u' to v' which does not pass through u. As we saw earlier, the sets of colored edges now define two perfect matchings. See figure 2.

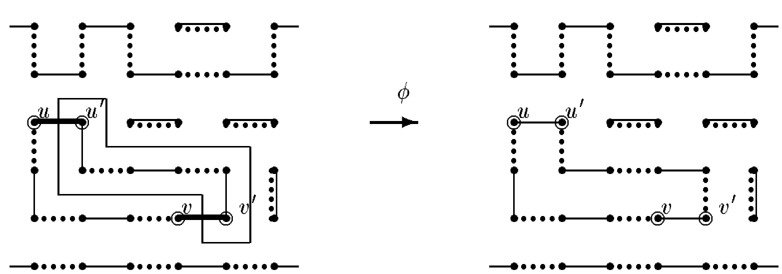

Figure 2: Mapping two near-perfect matchings to two perfect matchings.

We need to check that this map ϕ is injective: given any pair of perfect matchings (M_1, M_2) in the image of the map, we show that we can uniquely reconstruct the pair of near-perfect matchings, one with holes u and v and the other with holes u' and v', that are mapped by ϕ to (M_1, M_2). Note that the union of any pair of matchings in the image of ϕ always contains an alternating cycle that includes the edges (u, u') and (v, v'). Now color the edges of the matching containing (u, u') red, and the edges of the other matching blue. By flipping the colors of the edges along the path from u' to v' (again choosing the path which avoids u, for consistency), we make u' adjacent to two red edges. Since u' and v' are the holes of some near-perfect matching, they lie on opposite sides of the bipartition and any path between them must have odd length. Therefore, after the flipping operation v' must be adjacent to two red edges as well, while all other vertices are still adjacent to one edge of each color. If we now remove the edges (u, u') and (v, v'), the colored edges must correspond to the two near-perfect matchings that are mapped by ϕ to (M_1, M_2).

The above construction demonstrates that $|\mathcal{N}(u,v)|\,|\mathcal{N}(u',v')| \leq |\mathcal{M}|^2$. To finish the proof, we use the structure of the lattice $\widetilde{L}(n,d)$: in a periodic lattice, the operation of shifting a matching one position to the right is a bijection between the sets $\mathcal{N}(u,v)$ and $\mathcal{N}(u',v')$, so $|\mathcal{N}(u,v)| = |\mathcal{N}(u',v')|$. Thus the above relationship gives $|\mathcal{N}(u,v)|^2 \leq |\mathcal{M}|^2$, which implies $|\mathcal{N}(u,v)| \leq |\mathcal{M}|$. Summing over all choices of a black vertex u and a white vertex v, we find that $|\mathcal{N}| \leq n^{2d}|\mathcal{M}|/4$. The proof is completed by appealing to Theorem 4. □

Remark. It should be clear from the above proof that Theorem 1 generalizes to "hybrid" lattices that have fixed boundary conditions in some dimensions provided there exists at least one dimension in which the lattice has periodic boundary conditions. It also holds in more general bipartite rectangular lattices of size $n_1 \times n_2 \times \ldots \times n_d$ with periodic boundary conditions (i.e., for any dimension i in which the boundary is periodic, n_i must be even). □

The following theorem extends the above technique to handle two-dimensional lattices with fixed boundaries. We again show that in these lattices the number of near-perfect matchings cannot be too large compared to the number of perfect matchings.

Theorem 2 *There exists a fpras for counting matchings of any cardinality in the two-dimensional lattice with fixed boundaries* $L(n, 2)$.

Proof. We will prove the theorem for the slightly more general case of $n_1 \times n_2$ lattices with fixed boundaries, where n_1 is even. Let τ be a map which shifts the lattice $L(n, 2)$ one position to the right in $\mathbb{Z}^2$; that is, for a vertex $w = (w_1, w_2)$, define $\tau(w) = (w_1 + 1, w_2)$. We extend this map to matchings in the natural way: if N is a matching, then $\tau(N) \in [2, n_1 + 1] \times [1, n_2]$ is defined by $(\tau(x), \tau(y)) \in \tau(N)$ iff $(x, y) \in N$.

Let $\mathcal{M}$ and $\mathcal{N}$ be the sets of perfect and near-perfect matchings respectively in the lattice $L(n, 2)$. As in the last proof, we will fix holes u and v and show that $|\mathcal{N}(u,v)| \leq |\mathcal{M}|$. We again define an injection $\phi : \mathcal{N}(u,v) \times \mathcal{N}(u,v) \hookrightarrow \mathcal{M} \times \mathcal{M}$ as follows. Let $N_1, N_2 \in \mathcal{N}(u,v)$ be two near-perfect matchings. Consider the subgraph C obtained by taking the union of N_1 with a shifted version of N_2 and adding the two special edges as before, i.e., $C = N_1 \cup \tau(N_2) \cup \{(u,u'), (v,v')\}$, where $u' = \tau(u)$ and $v' = \tau(v)$. Then all the vertices in the leftmost column 1 and the rightmost column $n+1$ have degree one in C, and all other vertices have degree two. Thus C is the union of even-length cycles and paths with each endpoint in either the first or $(n+1)$st column (see figure 3). Color the edges from N_1 red and the edges from N_2 blue.

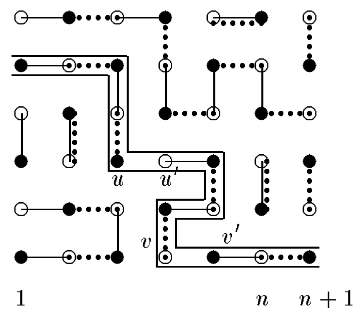

Figure 3: Union of N_1 and $\tau(N_2)$.

We will argue that, because the two-dimensional lattice is planar, any path or cycle which passes through u and u' must also pass through v and v'. This is immediate if u and u' lie on a cycle, so we focus on the case where u and u' lie on a path. The proof is by contradiction, and there are two cases to consider (see figure 4).

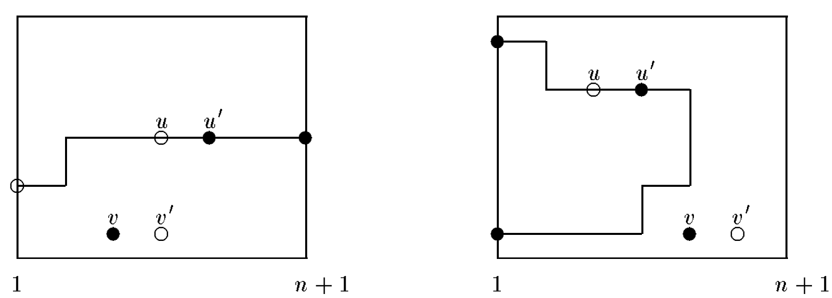

Figure 4: Proof of Theorem 2.

First, suppose that we have a path P from the first column to the $(n+1)$st column which passes through u and u', and not through v and v'. Without loss of generality we can assume that v and v' lie below P. Then P starts with a red edge, ends with a blue edge, and has one special edge, so it has odd length. It follows that if P starts at a black vertex then it ends at a white vertex, and conversely. Therefore, the number of vertices in the first column above P has opposite parity to the number of vertices in the $(n+1)$st column above P. (Since n is even, corresponding vertices in each of these columns fall on the same side of the black-and-white bipartition.) But consider the set of all vertices that lie above the path P. There must be an even number of these vertices lying in the first through nth columns, since these vertices are matched in N_1, and an even number lying in the second through $(n+1)$st columns, since these vertices are matched in N_2. This is a contradiction.

Second, suppose that P, the path going through u and u', starts and ends in the first column. By interchanging the roles of u, u' and v, v' if necessary, we may assume without loss of generality that v and v' lie outside the cycle defined by the path P and the first column. Now P starts and ends with a red edge and has one special edge, so it must have even length. If it starts at a black vertex it ends at a black vertex, and conversely, so there are an odd number of vertices in the first column that lie between these endpoints. Let S be the set of points that lie strictly inside the path P. Then $|S|$ must be even since N_1 matches all the vertices in S. But N_2 matches all the vertices in S except those which lie in the first column, a contradiction since this number is odd.

Therefore we can conclude that u, u', v, v' all lie on the same even-length cycle or the same path. In either case we can proceed as in the proof of Theorem 1: color the special edges red and then flip the colors of the edges along the path between u' and v' (in the case of a cycle, where this is ambiguous, we always choose the path which does not pass through u). The sets of colored edges then define two perfect matchings M_1 and $\tau(M_2)$.

Furthermore, given any two matchings in the image of the map ϕ we can uniquely reconstruct the pair of near-perfect matchings which are their preimage, so ϕ is injective. To see this, note that any element in the image of ϕ consists of two perfect matchings M_1 and M_2 such that $M_1 \cup \tau(M_2)$ contains a cycle or path which passes through all of u, u', v, v', and from here we can reconstruct N_1 with holes u and v and $\tau(N_2)$ with holes u' and v' by reversing the color flipping operation as in the proof of Theorem 1. Thus we have $|\mathcal{N}(u,v)| \leq |\mathcal{M}|$, and summing over choices of u, v we get $|\mathcal{N}| \leq n_1^2 n_2^2 |\mathcal{M}|/4$. Combining this result with Theorem 4, we see that there exists a fpras for counting matchings of any cardinality in $L(n,2)$. □

3 Other Lattices

The following theorem extends the techniques from the last section to handle other lattices. More precisely, we can, in polynomial time, approximately count the number of matchings of any size in any $2m$-vertex bipartite graph which is the Cayley graph of some finite group. Recall that the *Cayley graph* of a group G with a given set of generators is defined by identifying vertices with words in G and connecting vertices x and y by an edge in the graph if, for some generator a of G, $xa = y$. This class of graphs includes any finite hexagonal lattice which has periodic boundaries around some fundamental domain. One group which generates this lattice is $\langle a, b, c \mid a^2, b^2, c^2, (abc)^2, (ab)^i, (bc)^i, (ca)^i \rangle$, for any integer i, where a, b and c are the generators and the words which follow are relators equivalent to the identity in the group. See figure 5.

Theorem 3 *There exists a fpras for counting the number of matchings of any cardinality in any bipartite graph which is the Cayley graph of some finite group.*

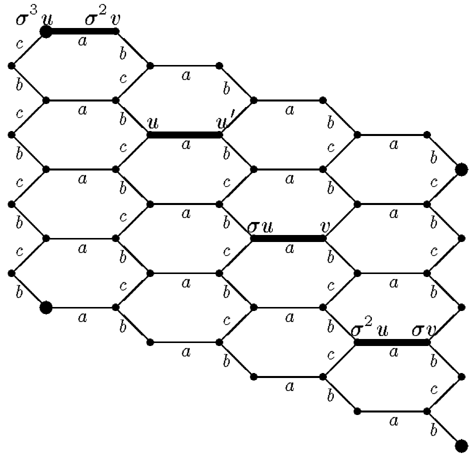

Figure 5: Proof of Theorem 3 for the hexagonal lattice
$\langle a, b, c \mid a^2, b^2, c^2, (abc)^2, (ab)^4, (bc)^4, (ca)^4 \rangle$

Proof. Given a group G, we consider its Cayley graph, which we assume to be bipartite. We will show that there cannot be many more near-perfect matchings than perfect matchings in the graph and then appeal again to Theorem 4. Choose a vertex and label it with e, the identity element of G. This determines a label in G for every vertex in the graph, corresponding to the product of the labels along any path leading to it from the identity vertex.

Fix a pair of holes u and v in the graph. Let $u' = ua$ be the neighbor of u defined by some fixed generator a. Let $\sigma = vu'^{-1}$ be the word in G which maps u' to v by multiplication on the left. Then, since any action on the graph preserves neighbors, σu is a neighbor of v. Moreover, since the group is finite, there exists some k such that $\sigma^k = e$. We will show that $|\mathcal{N}(u,v)| \leq |\mathcal{M}|$ by exhibiting an injection ϕ from the cartesian product $\prod_{i=1}^{k} [\mathcal{N}(\sigma^{i-1}u, \sigma^{i-1}v)]$ into $\mathcal{M}^k$.

We define the map ϕ in $k-1$ stages. Let $N_i \in \mathcal{N}(\sigma^{i-1}u, \sigma^{i-1}v)$, for $1 \leq i \leq k$, be a set of near-perfect matchings. Stage one maps the pair (N_1, N_2) into (M_1, N_2'), where M_1 is a perfect matching and N_2' is an "auxiliary" near-perfect matching. In stage j, for $2 \leq j \leq k-2$, we map the pair (N_j', N_{j+1}) into (M_j, N_{j+1}'), where N_j' is the auxiliary near-perfect matching from the previous stage. The $(k-1)$st stage maps (N_{k-1}', N_k) into the final pair of perfect matchings (M_{k-1}, M_k).

In the first stage we consider the subgraph $C_1 = N_1 \cup N_2 \cup \{(\sigma u, v)\}$. We color the edges from $N_1 \cup \{(\sigma u, v)\}$ red and those from N_2 blue. Then all vertices have degree two except u and σv, each of which has degree one, and v is the only vertex that has two edges of the same color incident to it. By flipping the colors of the edges along the portion of the path from v to σv, we can force the path from u to σv to have alternating colors. Because the graph is bipartite, the two vertices of degree one, u and σv, will both be adjacent to a blue edge.

Thus the blue edges form a perfect matching, M_1, and the red edges form the first auxiliary near-perfect matching, N'_2, with holes u and σv.

At the beginning of stage j in this mapping, we have already mapped $\prod_{i=1}^{k} [\mathcal{N}(\sigma^{i-1}u, \sigma^{i-1}v)]$ into $\mathcal{M}^{j-1} \times \mathcal{N}(u, \sigma^{j-1}v) \times \prod_{i=j+1}^{k} [\mathcal{N}(\sigma^{i-1}u, \sigma^{i-1}v)]$. Stage j itself will consist of an injection from $\mathcal{N}(u, \sigma^{j-1}v) \times \mathcal{N}(\sigma^j u, \sigma^j v)$ into $\mathcal{M} \times \mathcal{N}(u, \sigma^j v)$. In particular, we will map the pair (N'_j, N_{j+1}) to (M_j, N'_{j+1}), as follows.

If we consider the subgraph $C_j = N'_j \cup N_{j+1} \cup \{(\sigma^j u, \sigma^{j-1}v)\}$, we get an odd-length path from u to $\sigma^j v$. By flipping the colors of the edges along the portion of the path from $\sigma^{j-1}v$ to $\sigma^j v$, we get a perfect matching, M_j, and a near-perfect matching, N'_{j+1}, with holes u and $\sigma^j v$.

At the $(k-1)$st stage the mapping terminates. Here, we are mapping $N'_{k-1} \in \mathcal{N}(u, \sigma^{k-2}v)$ and $N_k \in \mathcal{N}(\sigma^{k-1}u, \sigma^{k-1}v)$. But $v = \sigma u'$, so $\sigma^{k-1}v = \sigma^k u' = u'$, since σ^k is the group identity. Therefore, the subgraph $N'_{k-1} \cup N_k \cup \{(\sigma^{k-1}u, \sigma^{k-2}v), (u, u')\}$ consists only of even-length cycles. By flipping colors along one of the paths from $\sigma^{k-2}v$ to u' (choosing the path which passes through u, to avoid ambiguity), we get even cycles with alternating colors: again this follows because the Cayley graph is bipartite. The two sets of colored edges now define the final two perfect matchings M_{k-1} and M_k.

Given the labels of the holes u and v, the vertex $u' = ua$ is uniquely determined, as is the word $\sigma = vu'^{-1}$ and its inverse. We can then invert the map ϕ by working backwards in stages, each stage being similar to the proof of Theorem 1. This shows that ϕ is injective, and therefore $\prod_{i=1}^{k} |\mathcal{N}(\sigma^{i-1}u, \sigma^{i-1}v)| \leq |\mathcal{M}|^k$.

The last step is to see that, for any word σ, translation by σ^i is a bijection between matchings $\mathcal{N}(u, v)$ and matchings $\mathcal{N}(\sigma^i u, \sigma^i v)$. More precisely, we extend σ^i to matchings by defining $(\sigma^i x, \sigma^i y) \in \sigma^i(N)$ iff $(x, y) \in N$, where $N \in \mathcal{N}(u, v)$. This is valid since if x and y are neighbors in the Cayley graph then there is some generator b such that $y = xb$, so $\sigma^i x$ and $\sigma^i y = \sigma^i xb$ are also neighbors. And, since u and v are unmatched in N, $\sigma^i u$ and $\sigma^i v$ are the unmatched vertices in $\sigma^i(N)$, so $\sigma^i(N) \in \mathcal{N}(\sigma^i u, \sigma^i v)$. Thus $|\mathcal{N}(u, v)| = |\mathcal{N}(\sigma^i u, \sigma^i v)|$ for any i. Combining this with the fact that ϕ is injective, we see that $|\mathcal{N}(u, v)|^k \leq |\mathcal{M}|^k$. Hence $|\mathcal{N}(u, v)| \leq |\mathcal{M}|$, which implies that $|\mathcal{N}| \leq n^2 |\mathcal{M}|$, where n is the size of the Cayley graph. $\square$

4 Open Problems

We have used simple techniques to show that the number of near-perfect matchings is polynomially related to the number of perfect matchings in any bipartite Cayley graph. Can similar techniques be used to show that the same relationship holds if we relax the bipartite condition and consider arbitrary Cayley graphs? This would allow us to handle several other lattices of interest in statistical physics, such as the triangular lattice and the face- and body-centered cubic lattices.

Our method also breaks down in the case of lattices with fixed boundary condition in dimensions higher than two. Techniques similar to those we have presented can be used to reduce the question of relating the number of near-perfect matchings to perfect matchings to that of showing the *local* property that the number of matchings with fixed holes u and v is polynomially related to the number of matchings with holes u' and v', where u' is a neighbor of u and v' is a neighbor of v. However, we have been unable to use this observation to obtain a proof for fixed boundary conditions in the general case.

It is not clear how the positions of the two holes u and v affect the number of near-perfect matchings in a lattice graph, a quantity which is also studied in the context of monomer-dimer systems [2]. We conjecture that the injections we have established between near-perfect matchings with fixed holes and perfect matchings might shed light on this correlation.

Finally, and somewhat speculatively, we conjecture that condition (1) on the ratio of the number of near-perfect matchings to the number of perfect matchings suggests a promising approach to designing a fpras for the permanent in the general case. It is already known that almost every graph satisfies (1) for a certain fixed polynomial q, and that the condition can be efficiently tested for an arbitrary graph [6]. Perhaps any graph that does not satisfy (1) can be efficiently decomposed in such a way that the resulting components satisfy (1), and hence fall within the scope of Theorem 4; this idea was used in [7] to obtain an approximation scheme for the general permanent whose running time, though still exponential, improves substantially on naïve deterministic methods. We hope that the methods of the present paper will contribute to a better understanding of condition (1) for general graphs.

References

[1] Fisher, M.E. Statistical mechanics of dimers on a plane lattice. *Physics Review* **124** (1961), pp. 1664–1672.

[2] Fisher, M.E. and Stephenson, J. Statistical mechanics of dimers on a plane lattice II. Dimer correlations and monomers. *Physics Review* **132** (1963), pp. 1411-1431.

[3] Fowler, R.H. and Rushbrooke, G.S. Statistical theory of perfect solutions. *Transactions of the Faraday Society* **33** (1937), pp. 1272–1294.

[4] Heilmann, O.J. and Lieb, E.H. Theory of monomer-dimer systems. *Communications in Mathematical Physics* **25** (1972), pp. 190–232.

[5] Jerrum, M.R. Two-dimensional monomer-dimer systems are computationally intractable. *Journal of Statistical Physics* **48** (1987), pp. 121–134.

[6] Jerrum, M.R. and Sinclair, A.J. Approximating the permanent. *SIAM Journal on Computing* **18** (1989), pp. 1149–1178.

[7] Jerrum, M.R. and Vazirani, U.V. A mildly exponential approximation algorithm for the permanent. *Proceedings of the 33rd IEEE Symposium on Foundations of Computer Science* (1992), pp. 320–326.

[8] Karp, R.M., Luby, M. and Madras, N. Monte-Carlo approximation algorithms for enumeration problems. *Journal of Algorithms* **10** (1989), pp. 429–448.

[9] Kasteleyn, P.W. The statistics of dimers on a lattice I. The number of dimer arrangements on a quadratic lattice. *Physica* **27** (1961), pp. 1209–1225.

[10] Kasteleyn, P.W. Dimer statistics and phase transitions. *Journal of Mathematical Physics* 4 (1963), pp. 287–293.

[11] Kasteleyn, P.W. Graph theory and crystal physics. In *Graph Theory and Theoretical Physics* (F. Harary ed.), Academic Press, London, 1967, pp. 43–110.

[12] Minc, H. *Permanents.* Addison-Wesley, Reading MA, 1978.

[13] Sinclair, A.J. *Algorithms for random generation and counting: a Markov chain approach.* Monograph Series Progress in Theoretical Computer Science, Birkhäuser, Boston, 1993.

[14] Temperley, H.N.V. and Fisher, M.E. Dimer problem in statistical mechanics—an exact result. *Philosophical Magazine* **6** (1961), pp. 1061–1063.

[15] Valiant, L.G. The complexity of computing the permanent. *Theoretical Computer Science* **8** (1979), pp. 189–201.

[16] Welsh, D.J.A. The computational complexity of some classical problems from statistical physics. In *Disorder in Physical Systems* (G. Grimmett and D. Welsh eds.), Clarendon Press, Oxford, 1990, pp. 307-321.