Approximating the Number of Monomer-Dimer Coverings of a Lattice*

Claire Kenyon† Dana Randall‡ Alistair Sinclair§

March 19, 1998

---

*A preliminary version of this paper appeared under the title "Matchings in lattice graphs" in the *Proceedings of the 25th ACM Symposium on Theory of Computing*, San Diego, May 1993, ACM Press, pp. 738–746.

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

‡Department of Computer Science, Princeton University. This work was done while the author was at the University of California at Berkeley. Supported in part by an AT&T PhD Fellowship and by NSF grant CCR92-01092. E-mail: `randall@cs.princeton.edu`.

§Computer Science Division, University of California at Berkeley. Part of this work was done while the author was visiting ICSI, Berkeley, on leave from the University of Edinburgh. Supported in part by grant GR/F 90363 of the UK Science and Engineering Research Council, by Esprit Working Group "RAND", and by NSF grant CCR95-05448. E-mail: `sinclair@icsi.berkeley.edu`.

Abstract

The paper studies the problem of counting the number of coverings of a d-dimensional rectangular lattice by a specified number of monomers and dimers. This problem arises in several models in statistical physics, and has been widely studied. A classical technique due to Fisher, Kasteleyn and Temperley solves the problem exactly in two dimensions when the number of monomers is zero (the dimer covering problem), but is not applicable in higher dimensions or in the presence of monomers. This paper presents the first provably polynomial time approximation algorithms for computing the number of coverings with any specified number of monomers in d-dimensional rectangular lattices with periodic boundaries, for any fixed dimension d, and in two-dimensional lattices with fixed boundaries. The algorithms are based on Monte Carlo simulation of a suitable Markov chain, and, in contrast to most Monte Carlo algorithms in statistical physics, have rigorously derived performance guarantees that do not rely on any assumptions. The method generalizes to counting coverings of any finite vertex-transitive graph, a class which includes most natural finite lattices with periodic boundary conditions.

Keywords: monomer-dimer problem, dimer coverings, lattice statistics, Monte Carlo methods, relaxation time, mixing time, approximation algorithm, Fisher-Kasteleyn-Temperley algorithm, perfect matchings, monomer-dimer correlations, vertex-transitive graphs.

1 Introduction

1.1 Historical background

A fundamental problem in lattice statistics is the *monomer-dimer* problem, in which the sites of a regular lattice are covered by a non-overlapping arrangement of monomers (molecules occupying one site) and dimers (molecules occupying two sites that are neighbors in the lattice). The three-dimensional problem occurs classically in the theory of mixtures of molecules of different sizes [11] and the cell-cluster theory of the liquid state [4]. In two dimensions, the problem serves as a model for the adsorption of diatomic molecules on a crystal surface [36]. (In this last example, 'monomers' correspond to empty sites.)

Most thermodynamic properties of the system can be deduced from knowledge of the number of ways of covering the lattice with given numbers of monomers and dimers. Suppose the lattice has $2m$ sites, and consider coverings consisting of s dimers and $2(m-s)$ monomers; the ratio $p = s/m$ is the *dimer density.* The essential problem is to compute the number of coverings as the lattice size m increases, for various values of p. Considerable effort has been invested in this problem over the past sixty years. In the remainder of this subsection we present a rather incomplete survey; for further information see, for example, [17, 29, 42] and the references given there.

The monomer-dimer problem gained prominence in 1937 through the early paper of Fowler and Rushbrooke [8]. A breakthrough was achieved in 1961, when, independently, Fisher, Kasteleyn and Temperley provided an analytic solution for the case of *dimer coverings* (i.e., arrangements with dimer density 1) on a two-dimensional rectangular lattice [6, 27, 39]. The key idea is to express the number of dimer coverings as a Pfaffian, which in turn can be evaluated as the square root of an associated determinant. These calculations give precise asymptotics for $f(n)$, the number of dimer coverings of an $n \times n$ rectangular lattice (with n even); specifically,

$$\frac{1}{n^2} \ln f(n) \to \lambda \text{ as } n \to \infty, \text{ where } \lambda = \frac{1}{\pi} \sum_{r \geq 0} \frac{(-1)^r}{(2r+1)^2} = 0.29156...$$

Moreover, since the problem is reduced to evaluation of a determinant, the quantity $f(n)$ can be computed numerically for any value of n in an efficient manner. In fact, this technique is more general and allows the number of dimer coverings of any *planar* graph (or indeed, of any family of graphs with fixed genus) to be computed efficiently [28].

Unfortunately, these methods do not extend to two-dimensional lattices with dimer density less than 1, or to lattices in higher dimensions even when the dimer density remains 1. This limitation was formalized by Hammersley et al [15] and, in a different sense, by Jerrum [18]. Indeed, the three-dimensional dimer covering problem, which asks for the number, $f(n)$, of ways of filling an $n \times n \times n$ rectangular lattice with dimers, is one of the classical unsolved problems of solid-state chemistry. A few facts are known: for example, $\ln(f(n))/n^3$ tends to a finite limit λ as n tends to infinity [12]. Hammersley [13] proved the lower bound $\lambda \geq 0.418347$, while the early paper by Fowler and Rushbrooke [8] showed the upper bound $\lambda \leq 0.54931$. It has been conjectured that λ lies between 0.43 and 0.45. In other work, Bhattacharjee et al [1] studied the phase transition behavior of the three-dimensional model. Notwithstanding these efforts, no reliable method is known for computing $f(n)$ to good accuracy. A similar lack of rigorous results holds for the problem at dimer densities less than 1, even in two dimensions. Notable exceptions are series expansions valid at low densities [10] and lower bounds on the free energy [2, 16].

1.2 Results

This paper makes progress on the monomer-dimer problem in cases where the technique of Fisher, Kasteleyn and Temperley fails. Specifically, we give a polynomial time algorithm for computing, to arbitrary precision, the number of coverings of a rectangular lattice in *any* dimension with *any* specified dimer density. Our algorithm applies to the case of *periodic* boundary conditions, i.e., the edges of the lattice are "wrapped around" to make it toroidal.

To make the behavior of the algorithm precise, for a fixed dimension d, let $f(n,s)$ denote the number of coverings of the d-dimensional rectangular lattice $[1,\ldots,n]^d$ (with periodic boundary conditions) by s dimers and $n^d - 2s$ monomers.

Definition. A *fully polynomial randomized approximation scheme* (*fpras*) for f is a probabilistic algorithm which, on inputs n, s and $\epsilon, \delta \in (0,1)$, always runs in time polynomial in n, ϵ^{-1}, and $\log \delta^{-1}$, and outputs a number A (a random variable) that, with probability at least $1-\delta$, satisfies

$$f(n,s)(1+\epsilon)^{-1} \le A \le f(n,s)(1+\epsilon). \qquad \square$$

The probabilistic algorithm can be thought of as performing a suitable Monte Carlo experiment (see Section 1.4 below) and using it to obtain a statistical estimate of $f(n,s)$. The parameter ϵ determines the accuracy required of the estimate, while δ controls the confidence level. A fpras provides an efficient means of numerically computing f, in the sense that its running time grows only *slowly* (i.e., polynomially) with the lattice size n, the accuracy parameter ϵ, and the confidence parameter δ. This is widely accepted in Computer Science as a robust criterion for algorithmic efficiency; for a justification of this definition, see, for example, [9, 26]. (By contrast, note that naïve algorithms based on exhaustive enumeration of coverings have a running time that is *exponential* in n, and are thus computationally useless in practice unless n is very small.)

The main result of this paper is a fpras for computing the above function $f(n,s)$ for rectangular lattices of any dimension d. This extends previous computational techniques in two ways. First, it enables one to compute the number of dimer coverings in lattices in three and higher dimensions. And second, it enables one to count coverings with dimer density less than 1, a problem that was not approachable by the methods of Fisher, Kasteleyn and Temperley even in two dimensions.

Our algorithm provides a feasible approach to numerical computation of such quantities as $f(n)$, the number of dimer coverings of an $n \times n \times n$ rectangular lattice in three dimensions. This is apparently the first such method whose running time provably grows only polynomially with n. We should, however, inject three qualifying remarks here. First, the running time of the algorithm, though polynomial, is not quite small enough to be genuinely practical; nonetheless, we strongly suspect that careful honing of the algorithm and its analysis will lead to a practical method. Secondly, the algorithm provides only statistical estimates of f, rather than precise values; we stress, however, that the error bars on these estimates can be made arbitrarily small, and, in contrast to previous Monte Carlo approximation methods, are completely rigorous and require no assumptions of any kind. Thirdly, although the algorithm allows $f(n)$ to be computed efficiently for each n, we do not provide bounds on the time required to compute the asymptotics of $f(n)$ as n tends to infinity, and therefore the entropy $\lim_{n\to\infty} \ln(f(n))/n^3$. This would require, in addition, bounds on the rate of convergence of this series.

As we have stated, the above result holds for lattices with *periodic* boundary conditions. In the two-dimensional case, the method extends to lattices with *fixed* boundaries: i.e., we again get a fpras for computing the number of coverings with any specified dimer density. This result again goes beyond the technique of Fisher, Kasteleyn and Temperley for planar graphs, which holds only for dimer density 1.[1]

Finally, we can extend the above results to a much broader class of lattices, or to any family of graphs with sufficiently strong symmetry properties. Specifically, we get a fpras for counting coverings, with any specified dimer density, of any finite vertex-transitive graph.[2] (A graph G is vertex-transitive if, for any pair of vertices u and v, there is an automorphism of G that maps u to v.) This class includes most other commonly studied lattices with periodic boundary conditions, such as the triangular lattice, the hexagonal lattice[3], and the body- and face-centered cubic lattices.

1.3 Techniques

The algorithms mentioned above are all based on a Monte Carlo procedure due to Jerrum and Sinclair [20, 38], originally inspired by the approach of Broder [3], for approximating the number of matchings in a graph. A *matching* in a $2m$-vertex graph $G = (V, E)$ is any subset M of the edge set E such that no two edges in M have a common endpoint. Clearly, matchings of cardinality s correspond precisely to monomer-dimer arrangements in G with s dimers and $2(m-s)$ monomers. The classical monomer-dimer problem discussed in the previous two

---

[1] If the number of monomers is some fixed constant $2c$ (so that the dimer density tends to 1 as $n \to \infty$) then the Fisher, Kasteleyn and Temperley technique can in principle be used, as follows. For each possible set of $2c$ positions for the monomers, use the technique to count dimer coverings in the graph formed by removing these sites from the lattice: this works because the graph remains planar. Now sum over all possible positions for the monomers. However, this approach no longer runs in polynomial time if c is allowed to grow with n, and is extremely inefficient in practice even for quite small fixed values of c.

[2] For graphs, as opposed to rectangular lattices, the definition of fpras must be modified so that the input is a $2m$-vertex graph G together with a number s, and the running time is a polynomial function of m, ϵ^{-1} and $\log \delta^{-1}$.

[3] An analytic solution to the dimer covering problem for this lattice has been known for some time [41, 28]. In contrast to the rectangular lattice, the assumption of periodic boundary conditions is important here: Elser [5] has solved the dimer covering problem on a hexagonal lattice with fixed boundaries, and shown that the result depends significantly on the shape of the boundary.

subsections is the special case in which G is the d-dimensional rectangular lattice $[1, \ldots, n]^d$ for some d. The Monte Carlo algorithm simulates a reversible Markov chain whose state space is the set of all matchings in the graph, and which converges to an equilibrium distribution in which each matching has a certain natural weight (see section 1.4 for details).

To quantify the behavior of the Monte Carlo procedure, we require some further terminology. Matchings in G of cardinality m are called *perfect matchings*, and those of cardinality $m-1$ are called *near-perfect matchings*: these correspond respectively to dimer coverings and coverings with precisely two monomers. We define the quantity $\alpha(G)$ to be the ratio of the number of near-perfect matchings to the number of perfect matchings in G (assuming the latter is non-zero). An application of the Monte Carlo procedure of Jerrum and Sinclair yields the following result.

Theorem 1 (Jerrum and Sinclair [20, Theorem 5.3]) *There exists a fpras for the number of matchings of any cardinality in any family of $2m$-vertex graphs G that satisfies $\alpha(G) \leq q(m)$, for a fixed polynomial q.* $\square$

In fact, the running time of the fpras is intimately related to $\alpha(G)$, and actually depends *linearly* on it. We shall sketch the algorithm and explain this connection in Section 1.4. For the moment, however, we merely note that a good upper bound on $\alpha(G)$ will give us an efficient approximation algorithm for the number of monomer-dimer coverings.

The ratio $\alpha(G)$ measures the factor by which the number of near-perfect matchings in G exceeds the number of perfect matchings. Note that this ratio is always at least m, since the removal of any edge from a perfect matching yields a unique near-perfect matching. For an efficient algorithm, we want the ratio to be not too much larger than m: in particular, for a fpras it must be bounded above by a polynomial function of m for the family of graphs in question. Note that this is not a trivial property: it is not hard to construct a family of $2m$-vertex graphs, $m = 1, 2, \ldots$, for which the ratio grows exponentially with m. We will discuss this issue in more detail in Section 4.

Our main technical contribution in this paper is to prove that the ratio $\alpha(G)$ is small for lattices and, more generally, for any family of graphs with sufficiently strong symmetry properties. Specifically, we show that if G is the d-dimensional rectangular lattice $[1, \ldots, n]^d$ with periodic boundary conditions (so that $m = \frac{1}{2}n^d$), then $\alpha(G) \leq m^2 = \frac{1}{4}n^{2d}$. This ensures that the Monte Carlo algorithm is in fact a fpras: i.e., its running time grows only polynomially with n for any fixed dimension d. A similar bound holds for arbitrary vertex-transitive graphs: namely, if G is a vertex-transitive graph with $2m$ vertices, then $\alpha(G) \leq 4m^3$, and $\alpha(G) \leq m^2$ if in addition G is bipartite.

We stress that our Monte Carlo algorithm differs from earlier ones for monomer-dimer systems (see, e.g., [14]), and indeed for many other problems in statistical physics, in that it is guaranteed (independent of any heuristic arguments) to provide statistically reliable estimates in a running time that grows only polynomially with the number of lattice sites. Monte Carlo algorithms with this property have recently been devised for various other problems, such as the Ising model [21] and the ice model [32]. In all cases, the key to the analysis is to prove a good bound on the rate of convergence to equilibrium of the Markov chain being simulated (often called the *relaxation time*, or *mixing time*, of the chain). For recent surveys of the analytical technology developed for estimating mixing rates of Markov chains, see, e.g., [22, 24, 38].

Our proofs of the above bound for lattices and general vertex-transitive graphs, presented in the next two sections, are elementary and construct explicit injections from pairs of near-perfect matchings to pairs of perfect matchings. The proofs also make crucial use of the strong symmetry properties of the lattice (and of arbitrary vertex-transitive graphs), which allow any matching (monomer-dimer configuration) to be translated. We conjecture that this technique may shed more light on other quantities related to monomer-dimer systems, and in particular the correlation between monomers at two specified sites, as studied in two dimensions by Fisher and Stephenson [7].

The remainder of the paper is organized as follows. In Section 1.4, for the sake of completeness, we sketch the Monte Carlo algorithm of [20] upon which this paper is based. Sections 2 and 3 contain our technical contributions: in Section 2 we prove upper bounds on $\alpha(G)$ for rectangular lattices with periodic boundary conditions in any dimension, and with fixed boundaries in two dimensions; in Section 3 we extend our techniques to handle arbitrary finite vertex-transitive graphs. Finally, in Section 4 we conclude with some further remarks on the physical and combinatorial significance of the ratio $\alpha(G)$, together with some open problems.

1.4 Overview of the algorithm

This subsection is devoted to a sketch of the Monte Carlo algorithm of Jerrum and Sinclair [20], which counts monomer-dimer coverings in any family of graphs G for which $\alpha(G)$ is small. This material is not essential to the technical development of the paper, but is included for the benefit of the reader interested in the background to Theorem 1.

Let $G = (V, E)$ be a graph with $|V| = 2m$ that contains a perfect matching. Let μ be any positive real number, and associate with each matching M in G a *weight* $w(M) = \mu^{|M|}$, where $|M|$ denotes the cardinality of matching M. Define the *monomer-dimer partition function* (or *generating function*) of G by

$$Z_G(\mu) = \sum_M w(M) = \sum_{s=0}^{m} a_s \mu^s,$$

where the coefficient a_s is the number of matchings in G of cardinality s. Thus, in the monomer-dimer problem on G, we are trying to compute the coefficients a_s for various values of s. In what follows, when G is understood we shall suppress the subscript and write Z in place of Z_G.

The Monte Carlo method described in [20, Section 4] simulates a Metropolis-style Markov chain whose state space is the set of matchings in G and whose equilibrium distribution π_μ is

$$\pi_\mu(M) = \frac{\mu^{|M|}}{Z(\mu)}.$$

Thus the equilibrium probability of a matching M is proportional to its weight, and the normalizing factor is the partition function. In this chain, a transition is possible between two matchings if and only if they differ up to the addition, deletion or exchange of one edge. The standard Metropolis acceptance rule is applied to proposed transitions so as to obtain the desired equilibrium distribution π_μ. Thus, more precisely, transitions from any matching M are made according to the following rule:

1. with probability $\frac{1}{2}$ let $M' = M$; otherwise,
2. select an edge $e = \{u, v\} \in E$ u.a.r. and set

$$M' = \begin{cases} M - e & \text{if } e \in M; \\ M + e & \text{if both } u \text{ and } v \text{ are unmatched in } M; \\ M + e - e' & \text{if exactly one of } u \text{ and } v \text{ is matched in } M \\ & \quad\quad \text{and } e' \text{ is the matching edge;} \\ M & \text{otherwise;} \end{cases}$$

3. go to M' with probability $\min\{1, \pi_\mu(M')/\pi_\mu(M)\}$.

Note that implementing one step of this process is simple; in particular, the probability in Step 3 involves only the *ratio* of the weights of M' and M, namely $\mu^{|M'|-|M|}$, which is easy to compute. It is not hard to verify that this Markov chain is ergodic and reversible, with equilibrium distribution π_μ.

By simulating the above Markov chain for sufficiently many steps until it reaches equilibrium, and observing its final state, one is effectively able to sample from the distribution π_μ. By repeated independent sampling at suitable values of μ, good statistical estimates of any desired coefficient a_s can be computed, as follows.

First, we show how to reduce the problem of computing a_s to that of computing the partition function itself at a suitable value $\mu = \widehat{\mu}$. So suppose we know $Z(\widehat{\mu})$. Then we sample from the distribution $\pi_{\widehat{\mu}}$ and observe the proportion of matchings of cardinality s in the sample; clearly, this is an unbiased estimator of the quantity $a_s\widehat{\mu}^s/Z(\widehat{\mu})$, enabling us to read off the value of a_s. The value $\mu = \widehat{\mu}$ is chosen so as to ensure that the proportion we are trying to estimate is not prohibitively small: that this is always possible follows from the fact that the coefficients a_s are *log-concave*, i.e., $a_{s-1}a_{s+1} \leq a_s^2$ [17, 20]. This fact implies that there is a value $\mu = \widehat{\mu}$ for which $a_s\widehat{\mu}^s = \max_{s'}\{a_{s'}\widehat{\mu}^{s'}\}$, i.e., the matchings of cardinality s have largest aggregated weight in the distribution $\pi_{\widehat{\mu}}$. (The ideal value is $\widehat{\mu} = a_{s-1}/a_s$; in practice a suitable value can be determined experimentally by adjusting μ until the observed distribution π_μ peaks around matchings of cardinality s.) This means that the proportion we are trying to estimate is at least $(m+1)^{-1}$, so by a routine variance calculation a sample of size only $\mathrm{O}(m)$ suffices for a good statistical estimate.[4]

It remains to describe how to compute the partition function itself at a given point $\mu = \widehat{\mu}$. The strategy is to express $Z(\widehat{\mu})$ as the product

$$Z(\widehat{\mu}) = \frac{Z(\mu_r)}{Z(\mu_{r-1})} \times \frac{Z(\mu_{r-1})}{Z(\mu_{r-2})} \times \cdots \frac{Z(\mu_2)}{Z(\mu_1)} \times \frac{Z(\mu_1)}{Z(\mu_0)} \times Z(\mu_0),$$

[4]For simplicity, in this sketch we omit the dependence of the various quantities on the accuracy ϵ and confidence δ, and treat these as constants.

where $0 = \mu_0 < \mu_1 < \cdots < \mu_r = \widehat{\mu}$ is a suitably chosen sequence of values. Note that $Z(\mu_0) = Z(0) = 1$. We can then estimate each factor $Z(\mu_i)/Z(\mu_{i-1})$ in this product by sampling from the distribution π_{μ_i} (obtained from the Markov chain in equilibrium with $\mu = \mu_i$), as follows. Consider the random variable $f_i(M) = (\frac{\mu_{i-1}}{\mu_i})^{|M|}$, where M is a matching chosen from the distribution π_{μ_i}. The expectation of this random variable is

$$\mathbf{E}\, f_i = \sum_M \left(\frac{\mu_{i-1}}{\mu_i}\right)^{|M|} \frac{\mu_i^{|M|}}{Z(\mu_i)} = \frac{1}{Z(\mu_i)} \sum_M \mu_{i-1}^{|M|} = \frac{Z(\mu_{i-1})}{Z(\mu_i)}.$$

Thus the ratio $Z(\mu_{i-1})/Z(\mu_i)$ can be estimated by sampling matchings from the distribution π_{μ_i} and computing the sample mean of f_i. The sequence of values μ_i needs to be chosen to make the expectation $Z(\mu_{i-1})/Z(\mu_i)$ not too small, so that the sample size required for a statistically good estimate is not too large. In practice, a suitable sequence can be determined by experiment; however, by analyzing the variance of the estimator it is possible to prove that the sequence $\mu_1 = (2|E|)^{-1}$ and $\mu_i = (1 + \frac{1}{m})\mu_{i-1}$ for $1 < i < r$, with a sample size of only about $\widetilde{\mathrm{O}}(m)$ for each ratio, suffices. (The tilde in the O-expression hides small logarithmic factors as well as constants.) In this case, the number of ratios, r, is also only $\widetilde{\mathrm{O}}(m)$. For the details, see [22, Section 4] (or [20, Section 5] for an alternative, slightly less efficient approach).

From the above description, it should be clear that a modest number (at most $\widetilde{\mathrm{O}}(m^2)$) of independent samples from the Markov chain in equilibrium are enough for a good statistical estimate of any coefficient a_s. However, in order to obtain each sample, the Markov chain must be simulated for sufficiently many steps that it is close to equilibrium: indeed, this turns out to be the dominant factor in the running time of the algorithm. To quantify this, one must prove a good *a priori* bound on the *relaxation time*, or *mixing time* of the chain, which is a challenging task. This was done in [20, 38], where the mixing time was shown to be a polynomial function of m and the parameter μ, specifically $\widetilde{\mathrm{O}}(\mu'^4 m|E|^2)$, where $\mu' = \max\{\mu, 1\}$. This bound has since been substantially improved to $\widetilde{\mathrm{O}}(\mu' m^2|E|)$: see [22, Section 4].

Examining the algorithm described above we see that, in order to compute the coefficient a_s, we have to use the Markov chain with various parameters $\mu \le a_{s-1}/a_s$. (By log-concavity, this ratio increases with s.) Hence the time required for each sample will be no more than $\widetilde{\mathrm{O}}(m^2|E|a_{s-1}/a_s)$. Thus the sampling time will be largest for the highest coefficient a_m (i.e., the number of dimer coverings of G), and is bounded by $\widetilde{\mathrm{O}}(m^2|E|\alpha(G))$. Putting this together with our earlier analysis, we deduce that the coefficient a_m (and hence any coefficient a_s) can be computed in time $\widetilde{\mathrm{O}}(m^4|E|\alpha(G))$.[5] This expression will be polynomial in m iff $\alpha(G) \le q(m)$ for some polynomial q, which explains Theorem 1. Note also that the running time of the approximation algorithm is linear in $\alpha(G)$, as we claimed in the previous subsection.

[5]Note that the quantity $\alpha(G)$ appears only by virtue of the highest coefficient a_m. For lower coefficients a_s, it is replaced by the smaller ratio a_{s-1}/a_s. It follows that one can obtain all the coefficients a_s with $s \le (1-\xi)m$ in time polynomial in $m^{1/\xi}$, regardless of the value of $\alpha(G)$. However, this running time grows exponentially with $(1-p)^{-1}$, where $p = 1 - \xi$ is the maximum dimer density. Note also that the above algorithm provides a fpras for the entire partition function Z_G — though of course not for all its coefficients — regardless of the value of $\alpha(G)$.

2 Rectangular Lattices

We begin by introducing some definitions and notation concerning lattices. We will be interested in two classes of lattices: the first class are those 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 lattices 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.

Throughout the remainder of the paper, we shall adopt the terminology of graphs and matchings introduced in Section 1.3. Thus we view $L(n, d)$ and $\widetilde{L}(n, d)$ as graphs with $2m = n^d$ vertices, and we always assume that n is even, so that both $L(n, d)$ and $\widetilde{L}(n, d)$ contain a perfect matching (dimer covering). In the next section, we will be working with general vertex-transitive graphs with an even number, $2m$, of vertices; it is known that all such graphs contain a perfect matching [31]. For any graph G, we will denote by $\mathcal{M}$ the set of perfect matchings in G and by $\mathcal{N}$ the set of near-perfect matchings (monomer-dimer coverings with exactly two monomers). In any matching (monomer-dimer covering), we refer to the set of unmatched vertices in the graph as *holes*, and we write $\mathcal{N}(u, v)$ for the set of near-perfect matchings with holes u and v.

In this section, all the graphs we consider will be bipartite, with m vertices on each side of the bipartition. It will sometimes be convenient to view the vertices on one side of the bipartition as being colored white, and those on the other side 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.) Note that in any near-perfect matching, one hole is white and the other black.

Recall that our aim is to construct efficient approximation algorithms for the number of monomer-dimer coverings of various lattice graphs with any specified number of dimers. This notion of efficient approximation algorithm is formalized as a *fully-polynomial randomized approximation scheme* (fpras), as defined in Section 1.2. All our algorithms will appeal to Theorem 1 of Section 1.3, which says that it is sufficient to demonstrate a suitable upper bound on the quantity $\alpha(G) = |\mathcal{N}|/|\mathcal{M}|$, the ratio of the number of near-perfect matchings to the number of perfect matchings in the lattice graph G. (Note that α is well-defined for all the graphs we consider because $|\mathcal{M}| > 0$.) Specifically, to get a fpras we need to show that $\alpha(G) \leq q(m)$ for $2m$-vertex lattices G, where q is a polynomial.

We now proceed to prove that such a relationship holds for families of rectangular lattice graphs, and (in the next section) for more general vertex-transitive graphs. 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.

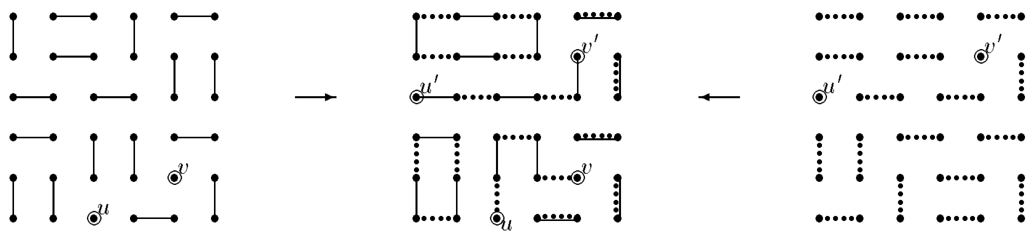

Figure 1: The union of two near-perfect matchings.

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. We use this observation 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 are now in a position to state our first result.

Theorem 2 *For the d-dimensional periodic lattice $\widetilde{L}(n,d)$, the ratio $\alpha(\widetilde{L}(n,d))$ is bounded above by $n^{2d}/4$.*

Before proving this theorem, we combine it with Theorem 1 to obtain the following immediate corollary.

Corollary 3 *There exists a fpras for the number of monomer-dimer coverings with any specified number of dimers in the d-dimensional periodic lattice $\widetilde{L}(n,d)$, for any fixed dimension d.*

Proof of Theorem 2. Let $\mathcal{M}$ and $\mathcal{N}$ be the sets of perfect and near-perfect matchings respectively in $\widetilde{L}(n,d)$. First we fix two holes, u and v. We will show that, regardless of the choice of u and v, it is the case that $|\mathcal{N}(u,v)| \le |\mathcal{M}|$. Summing over all choices of a white hole u and a black hole v, this implies that $|\mathcal{N}| \le n^{2d}|\mathcal{M}|/4$, so that $\alpha(\widetilde{L}(n,d)) = |\mathcal{N}|/|\mathcal{M}| \le n^{2d}/4$, as claimed.

To prove the above bound on $|\mathcal{N}(u,v)|$, note first that we may assume that u and v are not adjacent in G: if they are then trivially $|\mathcal{N}(u,v)| \le |\mathcal{M}|$, since any matching in $\mathcal{N}(u,v)$ can be extended to a perfect matching by adding the edge $\{u,v\}$. Now, 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. Our assumption that u and v are not adjacent ensures that u, u', v, v' are all distinct.

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).

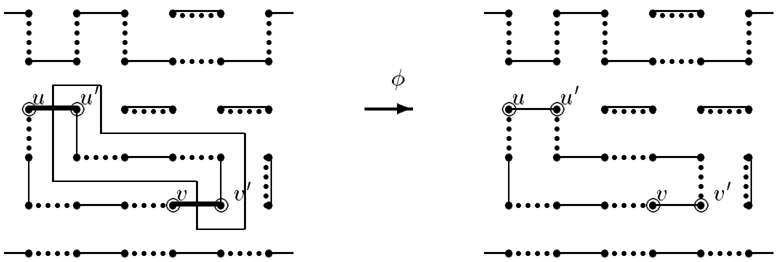

Figure 2: Mapping two near-perfect matchings to two perfect matchings.

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 both of 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 $|\mathcal{N}(u,v)|\,|\mathcal{N}(u',v')| \le |\mathcal{M}|^2$. 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 $|\mathcal{N}(u,v)|^2 \le |\mathcal{M}|^2$, which implies $|\mathcal{N}(u,v)| \le |\mathcal{M}|$ as required. □

Remark. It should be clear from the above proof that Theorem 2 (and hence Corollary 3) 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 (thus allowing shifting to the right). 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. Again we show that in these lattices the number of near-perfect matchings cannot be too large compared to the number of perfect matchings, and then appeal to Theorem 1.

Theorem 4 *For the two-dimensional lattice with fixed boundaries $L(n,2)$, the ratio $\alpha(L(n,d))$ is bounded above by $n^4/4$.*

Corollary 5 *There exists a fpras for the number of monomer-dimer coverings with any specified number of dimers in the two-dimensional lattice with fixed boundaries $L(n,2)$.* □

Proof of Theorem 4. Let τ be a map which shifts the lattice $L(n,2)$ one position to the right in $\mathbf{Z}^2$; 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 in $L(n,2)$, then $\tau(N)$ is the matching in $[2, n+1] \times [1, n]$ 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 $|\mathcal{N}(u,v)| \le |\mathcal{M}|$. We do this by defining 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 cycles of even length 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 $\tau(N_2)$ blue.

We will argue that any path or cycle which passes through u and u' must also pass through v and v'. Since C is bipartite, this is immediate if u and u' lie on a cycle, so we focus on the

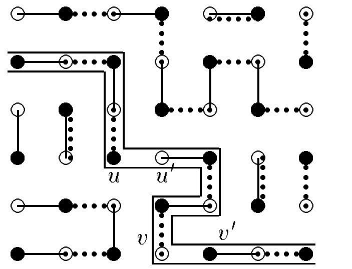

Figure 3: Union of N_1 and $\tau(N_2)$.

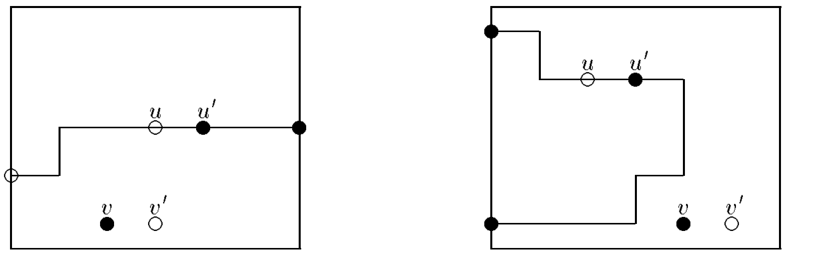

Figure 4: Proof of Theorem 4.

case where u and u' lie on a path; here the planarity of $L(n, 2)$ will play a crucial role. The proof is by contradiction, and there are two cases to consider (see figure 4).

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 (respectively, white) vertex then it ends at a white (respectively, black) vertex. 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 (respectively, white) vertex then it must end at a black (respectively, white) vertex, 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 2: 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 2. Thus we have $|\mathcal{N}(u,v)| \leq |\mathcal{M}|$. Summing over choices of u and v, we get $|\mathcal{N}| \leq n^4|\mathcal{M}|/4$, which yields the required bound on $\alpha(G)$. $\square$

Remark. The above proof, and hence Theorem 4 and Corollary 5, extends in obvious fashion to $n_1 \times n_2$ lattices with fixed boundaries where n_1 is even. $\square$

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 monomer-dimer coverings with any specified number of dimers in any finite vertex-transitive graph.

Recall that a graph G is *vertex-transitive* if, for any two vertices u and v in G, there exists an automorphism of G (i.e., a bijection from the vertex set to itself that preserves adjacency) which maps u to v. This class of graphs includes most other commonly studied lattices with periodic boundary conditions, such as the triangular lattice, the hexagonal lattice, and the body- and face-centered cubic lattices, as well as all finite Cayley graphs. It does not, however, include lattices with *fixed* boundary conditions.

Theorem 6 *Let G be a vertex-transitive graph with $2m$ vertices. Then $\alpha(G) \leq 4m^3$, and $\alpha(G) \leq m^2$ if in addition G is bipartite.*

Note that this theorem, applied to the bipartite vertex-transitive graph $\widetilde{L}(n,d)$, yields precisely the same bound on α as Theorem 2.

Corollary 7 *There exists a fpras for the number of monomer-dimer coverings with any specified number of dimers in any finite vertex-transitive graph.* $\square$

Proof of Theorem 6. Let x, y be a pair of vertices in G such that $|\mathcal{N}(x,y)| = \max_{u,v} |\mathcal{N}(u,v)|$. Now consider any pair of holes u and v, and let v' be a neighbor of v such that $d(u,v') < d(u,v)$, where $d(\cdot\,,\cdot)$ denotes distance in G. Since G is vertex-transitive, there exists a mapping of the vertex set which sends x to v' and preserves adjacency in G; let u' be the image of y under this mapping. Clearly the mapping is a bijection between $\mathcal{N}(x,y)$ and $\mathcal{N}(u',v')$, so $|\mathcal{N}(u',v')|$ is also maximum.

We first consider the simpler case when G is bipartite. We will construct an injection ϕ from $\mathcal{N}(u,v) \times \mathcal{N}(u',v')$ to $\mathcal{N}(u,u') \times \mathcal{M}$, implying that

$$|\mathcal{N}(u,v)|\,|\mathcal{N}(u',v')| \;\le\; |\mathcal{N}(u,u')|\,|\mathcal{M}|,$$

and hence that $|\mathcal{N}(u,v)| \le |\mathcal{M}|$ since $|\mathcal{N}(u',v')|$ is maximum. Summing over all pairs (u,v) we get $|\mathcal{N}| \le m^2|\mathcal{M}|$, which verifies the second claim in the theorem.

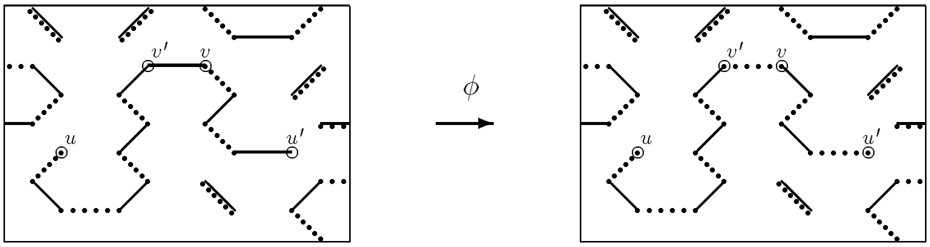

Figure 5: Definition of ϕ (shown on a region of the periodic hexagonal lattice).

The injection ϕ is defined as follows: let $N_1 \in \mathcal{N}(u,v)$ and $N_2 \in \mathcal{N}(u',v')$, and consider the subgraph $C = N_1 \cup N_2 \cup \{\{v,v'\}\}$. Color the edges of N_1 red and the edges of N_2 blue. In C, all vertices have degree two except for u and u', which both have degree one. Since G is bipartite all cycles in C have even length, so the edge $\{v,v'\}$, which is adjacent to one red edge and one blue edge, cannot close a cycle. Therefore $\{v,v'\}$ must lie on the path P with endpoints u and u'. By coloring the edge $\{v,v'\}$ red, and then flipping the colors of the edges along the portion of P between u' and v', we can ensure that P has alternating colors, starting and ending with blue edges. Now the blue edges form a perfect matching, and the red edges form a near-perfect matching with holes u and u', i.e., an element of $\mathcal{N}(u,u')$ (see figure 5). This completes the definition of ϕ. Injectivity follows by an argument very similar to that in the proof of Theorem 2.

In the general case, where G is not necessarily bipartite, the above injection breaks down because we cannot assume that the edge $\{v,v'\}$ lies on the path between u and u'. However, as we shall see in a moment, it is possible to construct an injection ψ from the set $\mathcal{N}(u,v) \times$

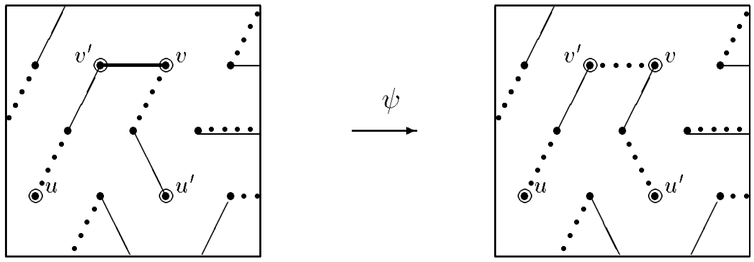

Case 1: u, u', v and v' lie on a single path.

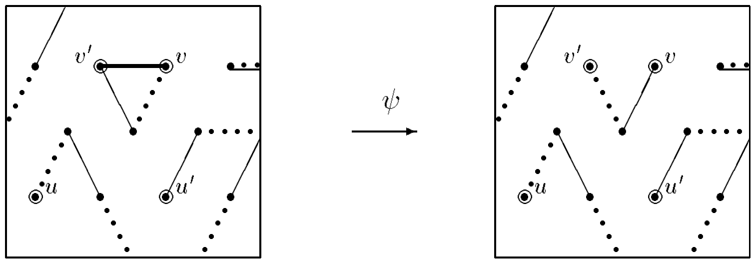

Case 2: v and v' lie on a cycle.

Figure 6: Definition of ψ (shown on a region of the periodic triangular lattice).

$\mathcal{N}(u', v')$ into the somewhat larger set $\Big(\mathcal{N}(u,u') \times \mathcal{M}\Big) \cup \Big(\mathcal{N}(u,v') \times \mathcal{N}(u',v)\Big)$. This implies that

$$|\mathcal{N}(u,v)|\,|\mathcal{N}(u',v')| \;\le\; |\mathcal{N}(u,v')|\,|\mathcal{N}(u',v)| \,+\, |\mathcal{M}|\,|\mathcal{N}(u,u')|,$$

and hence, since $|\mathcal{N}(u',v')|$ is maximum,

$$|\mathcal{N}(u,v)| \;\le\; |\mathcal{N}(u,v')| + |\mathcal{M}|. \tag{1}$$

Now from equation (1) it is easy to deduce that

$$|\mathcal{N}(u,v)| \;\le\; d(u,v)|\mathcal{M}|. \tag{2}$$

To see this, use induction on $d(u,v)$: the claim is immediate when $d(u,v) = 1$; when $d(u,v) > 1$, the induction hypothesis gives $|\mathcal{N}(u,v')| \le d(u,v')|\mathcal{M}| = (d(u,v)-1)|\mathcal{M}|$, which together with (1) implies that $|\mathcal{N}(u,v)| \le d(u,v)|\mathcal{M}|$. Finally, summing (2) over all pairs (u,v), we obtain $|\mathcal{N}| \le \sum_{u,v} 2m|\mathcal{M}| \le 4m^3|\mathcal{M}|$, and hence the first claim of the theorem.

It remains only to construct the injection ψ. Let $N_1 \in \mathcal{N}(u,v)$ and $N_2 \in \mathcal{N}(u',v')$, and again consider the subgraph $C = N_1 \cup N_2 \cup \{\{v,v'\}\}$, with the edges of N_1 colored red and those of N_2 colored blue. If the edge $\{v,v'\}$ lies on the path with endpoints u and u', we perform the same construction as in the bipartite case and obtain an element of $\mathcal{N}(u,u') \times \mathcal{M}$. If, on the other hand, the edge $\{v,v'\}$ lies on a cycle, we remove it and flip the colors of all edges along the path from v to v'; the effect of this is to create a red matching with holes u and v', and a blue matching with holes u' and v, i.e., an element of $\mathcal{N}(u,v') \times \mathcal{N}(u',v)$ (see figure 6). Checking that ψ is injective is again similar to the proof of Theorem 2. $\square$

4 Concluding Remarks and Open Problems

In this paper, we have used elementary combinatorial techniques to show that, for any vertex-transitive graph G, the quantity $\alpha(G)$ is small, i.e., the number of near-perfect matchings (monomer-dimer coverings with two monomers) exceeds the number of perfect matchings (dimer coverings) in G by only a small polynomial factor. This allowed us to deduce rigorous polynomial time bounds for a Monte Carlo algorithm for counting coverings in such graphs with any specified number of dimers.

Our results show that, for a vertex-transitive graph G with $2m$ vertices, the quantity $\alpha(G)$ lies in the range $[m, m^2]$ when G is bipartite, and $[m, 4m^3]$ in general. (The upper bounds come from Theorem 6, while the lower bound is trivial—see Section 1.3.) It would be interesting to know whether either of these bounds can be improved for general vertex-transitive graphs, and to determine the precise value of α for the d-dimensional rectangular lattice $\widetilde{L}(n,d)$. Apart from their inherent interest, these bounds would affect the efficiency of the Monte Carlo algorithm since the quantity $\alpha(G)$ enters into the running time as explained in Section 1.4.

Our technique appears to break down in the case of lattices with fixed boundary conditions (in dimensions higher than two). Arguments similar to those we have presented can be used to reduce the question of bounding α (at least in the bipartite case) to that of establishing the *local* property that the number of near-perfect 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 higher dimensions.

One can go further and ask for a precise characterization of those families of graphs for which the ratio α is polynomially bounded, and hence for which the monomer-dimer problem is tractable using the above Monte Carlo approach. This question is also of considerable combinatorial interest, since counting perfect matchings (dimer coverings) in a bipartite graph is equivalent to computing the *permanent* of a 0-1 matrix [33]. This is a widely studied problem in combinatorics for which the existence of an efficient approximation algorithm is an important open question in the theory of computation [40]. The Monte Carlo algorithm sketched above runs in polynomial time for a wider class of graphs than any other currently known algorithm, so it is of interest to establish precisely which graphs are amenable to it. (For other simpler, but apparently less widely applicable approximation algorithms, see [25, 19] and [35].) Moreover, it is conceivable that any graph G for which $\alpha(G)$ is large can be efficiently decomposed in such a way that the resulting components have a small value of α, and hence fall within the scope of the Monte Carlo algorithm; this idea was used in [23] to obtain an approximation scheme for general graphs whose running time, though still exponential, improves substantially on naïve deterministic methods.

The question of whether α is polynomially bounded for a given family of graphs is apparently rather subtle. It is not hard to construct "bad" examples. Consider, for example, the family of graphs $\{G_n : n = 1, 2, \ldots\}$ defined in figure 7, where G_n has $2m = 4n + 2$ vertices. It is easy to see that G_n has only one perfect matching but more than $2^n = 2^{(m-1)/2}$ near-perfect matchings (consider just those with holes at u and v), so the ratio $\alpha(G_n) > 2^n$ is exponentially large. On the other hand, α is known to be polynomially bounded for all sufficiently dense graphs, all graphs with sufficiently good "expansion" properties, and almost every random graph in a suitable model [20,34]. Interestingly, the technique used to prove this property in all these cases is not applicable to lattices since it involves constructing short augmenting paths for near-perfect matchings; such paths do not exist in lattice graphs, which have large diameter. The injective mapping technique presented in this paper is therefore a new approach, and we hope that it will lead to a better understanding of the behavior of the ratio α in general graphs.

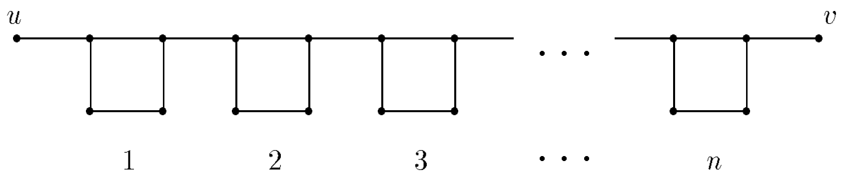

Figure 7: The "bad" graph G_n

Finally, in the case of lattices, we conjecture that the explicit mappings we have exhibited between near-perfect matchings with two fixed holes and perfect matchings might shed light on the behavior of the number of near-perfect matchings as a function of the positions of the holes. In physical terms, this corresponds to the correlation between a pair of monomers in a sea of dimers, a quantity for which partial results were obtained in two dimensions by Fisher and Stephenson [7]. For example, our techniques immediately yield a simple and rigorous proof that, for the rectangular lattice $\widetilde{L}(n, d)$ in any dimension d, the number of configurations with two monomers at any fixed pair of vertices u, v is bounded by n^{-d} times the number of configurations with two adjacent monomers. A more careful analysis may enable one to make more precise statements about this correlation.

Acknowledgements

We are very grateful to Mark Jerrum for permission to include his version of Theorem 6, which substantially simplifies and strengthens our original version in [30, Theorem 3]. We are also grateful to Mike Luby and Jim Propp, and especially to Stu Whittington, for helpful comments on the ideas presented here, and to Peter Cameron and Dominic Welsh for pointing out reference [31].

References

[1] Bhattacharjee, S.M., Nagle, J.F., Huse, D.A. and Fisher, M.E. Critical behaviour of a three-dimensional dimer model. *Journal of Statistical Physics* **32** (1983), pp. 361–374.

[2] Bondy, J.A. and Welsh, D.J.A. A note on the monomer dimer problem. *Proceedings of the Cambridge Philosophical Society* **62** (1966), pp. 503–505.

[3] Broder, A.Z. How hard is it to marry at random? (On the approximation of the permanent). *Proceedings of the 18th ACM Symposium on Theory of Computing* (1986), pp. 50–58. Erratum in *Proceedings of the 20th ACM Symposium on Theory of Computing* (1988), p. 551.

[4] Cohen, E.G.D., de Boer, J. and Salsburg, Z.W. A cell-cluster theory for the liquid state II. *Physica* **XXI** (1955), pp. 137–147.

[5] Elser, V. Solution of the dimer problem on a hexagonal lattice with boundary. *Journal of Physics A: Math. Gen.* **17** (1984), pp. 1509–1513.

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

[7] 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.

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

[9] Garey, M.R. and Johnson, D.S. *Computers and intractability.* Freeman, San Francisco, 1979.

[10] Gaunt, D.S. Exact series-expansion study of the monomer-dimer problem. *Physics Review* **179** (1969), pp. 174–179.

[11] Guggenheim, E.A. *Mixtures.* Clarendon Press, Oxford, 1952.

[12] Hammersley, J.M. Existence theorems and Monte Carlo methods for the monomer dimer problem. In *Research papers in statistics: Festschrift for J. Neyman* (F.N.David, ed.), John Wiley, New York, 1966, pp. 125–146.

[13] Hammersley, J.M. An improved lower bound for the multidimensional dimer problem. *Proceedings of the Cambridge Philosophical Society* **64** (1968), pp. 455–463.

[14] Hammersley, J.M. Calculation of lattice statistics. *Proceedings of the Computational Physics Conference* 1970, London, Institute of Physics and the Physical Society, pp. 1–8.

[15] Hammersley, J.M., Feuerverger, A., Izenman, A. and Makani, K. Negative finding for the three-dimensional dimer problem. *Journal of Mathematical Physics* **10**(3) (1969), pp. 443–446.

[16] Hammersley, J.M. and Menon, V.V. A lower bound for the monomer-dimer problem. *Journal of the Institute of Mathematics and its Applications* **6** (1970), pp. 341–364.

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

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

[19] Jerrum, M.R. An analysis of a Monte Carlo algorithm for estimating the permanent. *Proceedings of the 3rd Conference on Integer Programming and Combinatorial Optimization* (1993), CORE, Louvain-la-Neuve, Belgium, pp. 171–182.

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

[21] Jerrum, M. R. and Sinclair, A. J. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on Computing* **22** (1993), pp. 1087–1116.

[22] Jerrum, M. R. and Sinclair, A. J. Markov chain Monte Carlo Algorithms. In *Approximation Algorithms for NP-hard Problems* (D. Hochbaum, ed.), PWS Publishing, Boston, 1996, to appear.

[23] 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.

[24] Kannan, R. Markov chains and polynomial time algorithms. *Proceedings of the 35th IEEE Conference on Foundations of Computer Science* (1994), pp. 656–671.

[25] Karmarkar, N., Karp, R.M., Lipton, R., Lovász, L. and Luby, M. A Monte Carlo algorithm for estimating the permanent. *SIAM Journal on Computing* **22** (1993), pp. 284–293.

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

[27] 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.

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

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

[30] Kenyon, C., Randall, D. and Sinclair, A.J. Matchings in lattice graphs. *Proceedings of the 25th ACM Symposium on Theory of Computing* (1993), pp. 738–746. **4** (1963), pp. 287–293

[31] Little, C.H.C., Grant, D.G. and Holton, D.A. On defect-d matchings in graphs. *Discrete Mathematics* **13** (1975), pp. 41–54. Erratum in *Discrete Mathematics***14** (1976), p. 203.

[32] Mihail, M. and Winkler, P. On the number of Eulerian orientations of a graph. *Proceedings of the 3rd ACM-SIAM Symposium on Discrete Algorithms* (1992), pp. 138–145.

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

[34] Motwani, R. Expanding graphs and the average-case analysis of algorithms for matchings and related problems. *Proceedings of the 21st ACM Symposium on Theory of Computing* (1989), pp. 550–561.

[35] Rasmussen, L.E. Approximating the permanent: a simple approach. *Random Structures and Algorithms* **5** (1994), pp. 349–361.

[36] Roberts, J.K. Some properties of adsorbed films of oxygen on tungsten. *Proceedings of the Royal Society of London A* **152** (1935), pp. 464–480.

[37] Sinclair, A.J. Improved Bounds for Mixing Rates of Markov Chains and Multicommodity Flow. *Combinatorics, Probability and Computing* **1** (1992), pp. 351–370.

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

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

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

[41] Wannier, G.H. Antiferromagnetism: the triangular Ising net. *Physics Review* **79** (1950), pp. 357–364.

[42] 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.