

ALGORITHMIC PIROGOV-SINAI THEORY

TYLER HELMUTH, WILL PERKINS, AND GUUS REGTS

ABSTRACT. We develop an efficient algorithmic approach for approximate counting and sampling in the low-temperature regime of a broad class of statistical physics models on finite subsets of the lattice \mathbb{Z}^d and on the torus $(\mathbb{Z}/n\mathbb{Z})^d$. Our approach is based on combining contour representations from Pirogov–Sinai theory with Barvinok’s approach to approximate counting using truncated Taylor series. Some consequences of our main results include an FPTAS for approximating the partition function of the hard-core model at sufficiently high fugacity on subsets of \mathbb{Z}^d with appropriate boundary conditions and an efficient sampling algorithm for the ferromagnetic Potts model on the discrete torus $(\mathbb{Z}/n\mathbb{Z})^d$ at sufficiently low temperature.

CONTENTS

1. Introduction	1
2. Cluster expansions, Taylor series, and approximate counting	9
3. Contour models	16
4. Convergence of the cluster expansion	24
5. Sampling	27
6. The torus	32
7. Conclusions	36
Acknowledgements	38
References	38

1. INTRODUCTION

For a wide class of equilibrium lattice statistical mechanics models it is known that there is a phase transition from a high-temperature disordered state to a low-temperature ordered state. In many cases this transition is reflected in the dynamical and algorithmic behavior of these models. For example, a simple Markov chain (the Glauber dynamics) provides an efficient means of sampling from many models on finite subsets of \mathbb{Z}^d at high temperatures but is often known to be inefficient at low temperatures [15]. For many models there are no known efficient sampling algorithms at low temperatures, e.g., this is the case for the well-studied hard-core model and for the ferromagnetic q -state Potts model when q and d are greater than 2. See Sections 1.1 and 1.2 for definitions of these models.

Date: 6/13/2019.

Our main contribution is to rectify this by providing efficient approximate counting and sampling algorithms at low temperatures on subsets of \mathbb{Z}^d and on the torus $\mathbb{T}_n^d = (\mathbb{Z}/n\mathbb{Z})^d$. Our results apply to a wide class of statistical mechanics models, including the hard-core and ferromagnetic Potts models. The following theorem is representative of our results.

Theorem 1.1. *For all $d \geq 2$ and $q \geq 2$ there exists $\beta^* = \beta^*(d, q)$ such that for all inverse temperatures $\beta > \beta^*$ and all $c > 0$ there is a polynomial-time algorithm to sample from the q -state Potts model on \mathbb{T}_n^d within n^{-c} total variation distance.*

To the best of our knowledge, this is the first provably efficient sampling algorithm for the q -state Potts model on the torus \mathbb{T}_n^d below the critical temperature for $q, d \geq 3$. We are also able to give an efficient algorithm to approximate the partition function of the model, see Theorem 1.2 below.

Before describing our full results for the Potts and hard-core models we briefly recall the motivation for, and intuition behind, our work.

There are two natural computational problems associated to the Potts model and other discrete models from statistical physics. Given a graph G and an inverse temperature β the *counting* problem is to compute the partition function $Z(G, \beta)$ of the model, and the *sampling* problem is to produce a sample distributed according to the probability law of the model on G . If we take the graph G as our input, the algorithmic problem of computing $Z(G, \beta)$ can be $\#P$ -hard in general, and so research has focused on providing *approximate counting algorithms* that return values close to $Z(G, \beta)$ and *approximate sampling algorithms* that produce samples close in distribution to the given model. For many problems, namely those that are *self-reducible* [36, 55], the existence of an efficient approximate counting algorithm implies the existence of an efficient approximate sampling algorithm, and vice-versa.

The existence of efficient algorithms for these computational tasks is often known in the high-temperature regime of statistical physics models. In contrast, algorithms are often lacking in the low-temperature regime, even on restricted classes of graphs like lattices. This often reflects the existence of *phase transitions* in these models on certain infinite graphs, e.g., the infinite regular tree or \mathbb{Z}^d .

At the same time, the low-temperature regime of many discrete statistical physics models is fairly well-understood at a probabilistic level when the graph considered is a nice subset of \mathbb{Z}^d or the torus, see, e.g., [23, Chapter 7]. One might therefore hope that the algorithmic tasks of sampling and counting are tractable when restricted to these settings. Theorem 1.1 and our other results confirm that this is the case. While we focus in this paper on the Potts and hard-core models as they are two of the most studied lattice spin models, our main results (Theorems 3.1 and 5.5) are much more general and apply to many discrete statistical physics models e.g., the Widom–Rowlinson model, the Blume–Capel model, and many of the H -coloring models described in [14].

The most systematic probabilistic understanding of the low-temperature regime of discrete lattice spin models is based on Pirogov–Sinai theory. Roughly speaking, this is a significantly more sophisticated development of the Peierls’ contour argument. The main idea of our algorithms is to make use of Pirogov–Sinai theory to express the logarithm of the partition function as a convergent cluster expansion, where terms of the expansion correspond to overlapping clusters of contours. We then use the approach of Barvinok

to approximate the logarithm of the partition function, i.e., we truncate its Taylor series expansion and compute the initial coefficients exactly by using the cluster expansion representation. We describe this in more detail in Section 1.3 below.

Contour arguments have also been used to prove the slow mixing of Markov chains on lattices [15, 51, 16, 11], and our results can counterintuitively be phrased as saying that a contour-based proof that a Markov chain on \mathbb{Z}^d mixes slowly implies the existence of an efficient sampling algorithm at low enough temperatures.

In the next two sections we present our results for the Potts and hard-core models in detail, but first we give precise definitions for our notions of approximation. In this introduction we only define approximation for non-negative parameters though our main counting algorithms (Theorems 1.2–1.5 below) in fact apply for complex parameters. Readers interested in complex parameters should consult the more general Theorem 3.1.

We define fully polynomial-time approximation schemes in terms of the approximate evaluation of polynomials since many counting problems can be recast as the evaluation of a univariate polynomial. For a positive number p , we say \hat{p} is an ϵ -relative approximation to p if $e^{-\epsilon}\hat{p} \leq p \leq e^{\epsilon}\hat{p}$.

Definition 1. A fully polynomial time approximation scheme (FPTAS) for approximating the evaluation of a polynomial $p(z)$ with nonnegative coefficients at $z > 0$ is an algorithm that for any $\epsilon > 0$ produces an ϵ -relative approximation to $p(z)$ and runs in time bounded by a polynomial in $\deg(p)$ and $1/\epsilon$.

We use the total variation distance to measure the quality of an approximate sample.

Definition 2. An ϵ -approximate sample from a probability measure μ is a configuration drawn according to a probability measure $\hat{\mu}$ with

$$\|\hat{\mu} - \mu\|_{TV} < \epsilon.$$

Definition 3. Suppose (μ_n) is a sequence of probability measures indexed by n . An efficient sampling algorithm is a randomized algorithm that returns an ϵ -approximate sample to μ_n and runs in time polynomial in n and $1/\epsilon$.

1.1. The Potts model. The q -state Potts model on a finite graph $G = (V, E)$ is the probability distribution over assignments of q colors to the vertices V of G given by

$$\mu_{G,q,\beta}(\sigma) := \frac{\exp\left[\beta \sum_{\{i,j\} \in E} \mathbf{1}_{\sigma_i = \sigma_j}\right]}{Z_{G,q}(\beta)}$$

where

$$Z_{G,q}(\beta) := \sum_{\sigma \in [q]^V} \exp\left[\beta \sum_{\{i,j\} \in E} \mathbf{1}_{\sigma_i = \sigma_j}\right]$$

is the *partition function*. We have written $[q] := \{1, 2, \dots, q\}$ for the set of colours. In what follows we assume $\beta > 0$, i.e., that the model is *ferromagnetic*, meaning that it prefers configurations with more monochromatic edges. The case $q = 2$ of the Potts model is also called the *Ising model*.

The Potts model is a simple model of a magnetic material and in classical statistical physics it is studied on the d -dimensional lattice \mathbb{Z}^d . For the remainder of this discussion we will consider d and the number of colors q to be fixed. The Potts model on \mathbb{Z}^d is

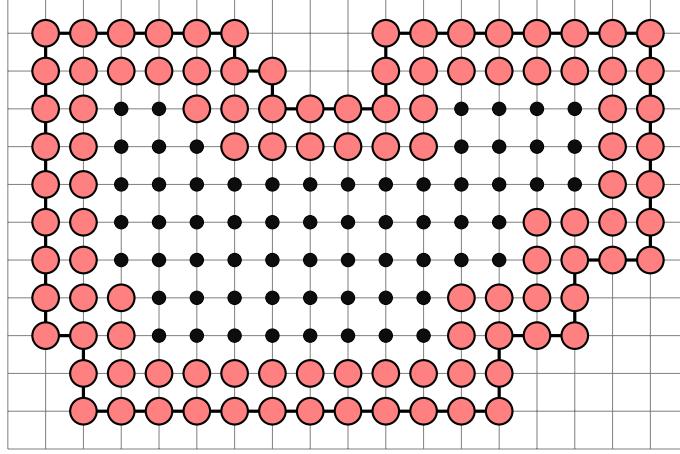


FIGURE 1. Red padded boundary conditions for the Potts model on a region Λ . The thick black line passes through the interior vertex boundary $\partial^{\text{in}}\Lambda$ of Λ . Vertices determined by the boundary condition have been drawn red. Solid black vertices indicate where the configuration is not determined by the boundary conditions.

defined by taking a sequence of finite graphs $\Lambda_n \subset \mathbb{Z}^d$ so that $\Lambda \rightarrow \mathbb{Z}^d$, and infinite volume measures are obtained as weak limits of finite volume measures $\mu_{\Lambda_n, q, \beta}$. If for a given choice of β only one infinite volume measure exists the model is said to be in the *uniqueness regime*. Otherwise, when multiple infinite volume measures are possible, the Potts model is said to exhibit *phase coexistence*. The transition between uniqueness and coexistence as β changes is a *phase transition* and occurs at a critical point $\beta_c(d, q)$ (see, e.g., [30]).

To state our results precisely requires two definitions. Let Λ be a subgraph of \mathbb{Z}^d . We write $E(\Lambda) \subset E(\mathbb{Z}^d)$ for the edge set of Λ , and by a slight abuse of notation, we write Λ in place of $V(\Lambda)$ for the vertex set of Λ . A finite subgraph Λ is a *region* if Λ^c is connected under the adjacency relation derived from the distance function $d_\infty(x, y) := \max_{i=1}^d |x_i - y_i|$. For a color $\varphi \in [q]$, the set of allowed configurations with *padded monochromatic boundary conditions* are:

$$\Omega_\Lambda^\varphi := \{\sigma \in [q]^\Lambda : d_\infty(i, \Lambda^c) \leq 2 \implies \sigma_i = \varphi\}.$$

See Figure 1. The corresponding partition function is

$$Z_{q, \Lambda}^\varphi(\beta) := \sum_{\omega \in \Omega_\Lambda^\varphi} \exp \left[\beta \sum_{\{i, j\} \in E(\Lambda)} \mathbf{1}_{\sigma_i = \sigma_j} \right].$$

Theorem 1.2. *For all $d \geq 2, q \geq 2$, there exists $\beta^* = \beta^*(d, q) > 0$ so that for all $\beta > \beta^*$, there is an efficient sampling algorithm and an FPTAS for the q -state Potts model on any finite region Λ of \mathbb{Z}^d with padded monochromatic boundary conditions.*

The running time of these algorithms is $(n/\epsilon)^{O(\log d)}$ where n is the number of vertices in the region Λ . While this is polynomial in n and $1/\epsilon$, it would be desirable to improve the running time, perhaps to something close to linear in n . See Section 7.3 for more.

On the torus \mathbb{T}_n^d a great deal of work has gone into understanding the mixing times of different Markov chains. When $d = 2$ a great deal is known: the Glauber dynamics and Swendsen–Wang dynamics mix rapidly (in polynomial time) for $\beta < \beta_c$, and the Swendsen–Wang dynamics mix rapidly for $\beta > \beta_c$ [15, 16, 60, 28, 12]. More generally the Swendsen–Wang dynamics are thought to be rapidly mixing for all d and q when $\beta \neq \beta_c$.

Our results hold on the torus for a slightly weaker notion of approximation.

Theorem 1.3. *For all $d \geq 2$ and $q \geq 2$ there exists $\beta^* = \beta^*(d, q)$ and $c = c(d, q) > 0$ so that for all $\beta > \beta^*$ and all $\epsilon \geq e^{-cn}$ there is an algorithm to obtain an ϵ -relative approximation of the partition function and an ϵ -approximate sampling algorithm both running in time polynomial in n and $1/\epsilon$ for the q -state Potts model on \mathbb{T}_n^d .*

1.1.1. Related results. Recall that an FPRAS is a randomized algorithm that returns an ϵ -relative approximation with probability at least $2/3$ and runs in time polynomial in the instance size and $1/\epsilon$. An FPRAS for the ferromagnetic Ising models on general graphs was given by Jerrum and Sinclair [35]. Randall and Wilson [52] showed that this algorithm can be used to sample efficiently from the model. Recently, Guo and Jerrum [31] gave an alternative sampling algorithm, based on a Markov chain associated to the random cluster model. For $q \geq 3$, the complexity of approximating the ferromagnetic Potts model partition function on general graphs is unknown. It is #BIS-hard (as hard as approximately counting the number of independent sets in a bipartite graph, see Section 1.2) to do so even on bounded degree graphs [29, 25].

By making use of Theorems 3.1 and 5.1 of the present article Jenssen, Keevash, and Perkins have proven a variant of Theorem 1.2 for the low temperature q -state Potts model on bounded degree expander graphs [34]. Subsequent to the initial posting of the present article to the arXiv, Barvinok and Regts have given an algorithm for approximating the partition function of the q -state Potts model at low temperatures on a variety of graphs [5]. Their main hypotheses concerns the existence of a nice set of generators for the cycle space of the graph, and for finite simply connected subsets of \mathbb{Z}^d they obtain estimates for $\beta_0(q)$ that are better than those implicitly given by Theorem 1.2.

1.2. Hard-core model. The *hard-core model* on a finite graph G is a random independent set \mathbf{I} from the set $\mathcal{I}(G)$ of all independent sets of G according to the distribution

$$\mu_{G, \lambda}(I) := \mathbb{P}[\mathbf{I} = I] = \frac{\lambda^{|I|}}{Z_G(\lambda)},$$

where $\lambda > 0$ is the *fugacity* and where the partition function is

$$Z_G(\lambda) := \sum_{I \in \mathcal{I}(G)} \lambda^{|I|}.$$

Our main result for the hard-core model is that if we take subsets of \mathbb{Z}^d with appropriate boundary conditions, then there are in fact efficient counting and sampling algorithms at high fugacities. To state our results, we recall that a vertex $i \in \mathbb{Z}^d$ is *even*

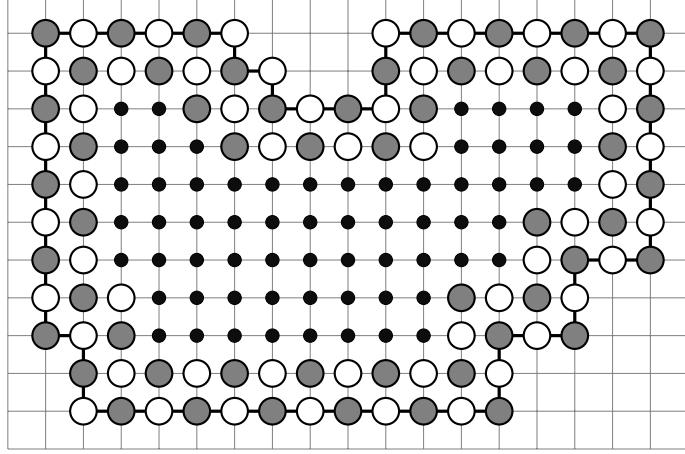


FIGURE 2. Even padded boundary conditions for the hard-core model on a region Λ . The thick black line passes through the interior vertex boundary $\partial^{\text{in}}\Lambda$ of Λ . Vertices required to be occupied or unoccupied by the boundary conditions are drawn as gray or white circles, respectively. Solid black vertices indicate where the configuration is not determined by the boundary conditions; note these vertices may be required to be unoccupied due to sharing an edge with a vertex required to be occupied by the boundary conditions.

(resp. *odd*) if the sum of its coordinates is even (resp. odd). For a finite region Λ , the set of allowed configurations under *even padded boundary conditions* is

$$\mathcal{I}^{\text{even}}(\Lambda) := \{I \in \mathcal{I}(\Lambda) : d_\infty(i, \Lambda^c) \leq 2 \implies \mathbf{1}_{i \in I} = \mathbf{1}_i \text{ even}\},$$

and likewise for $\mathcal{I}^{\text{odd}}(\Lambda)$. See Figure 2. The partition function is

$$Z_\Lambda^{\text{even}}(\lambda) := \sum_{I \in \mathcal{I}^{\text{even}}(\Lambda)} \lambda^{|I|}.$$

Theorem 1.4. *For $d \geq 2$ there exists a $\lambda^* = \lambda^*(d)$ such that for all $\lambda > \lambda^*$, there is an efficient sampling algorithm and an FPTAS for the hard-core model on any finite region Λ of \mathbb{Z}^d with even or odd padded boundary conditions.*

We also establish efficient counting and sampling algorithms on \mathbb{T}_n^d when n is even; this ensures the existence of an independent set that contains half of the vertices of \mathbb{T}_n^d .

Theorem 1.5. *For $d \geq 2$ there exists $\lambda^* = \lambda^*(d)$ and $c = c(d) > 0$ so that for all $\lambda > \lambda^*$ and all $\epsilon \geq e^{-cn}$ there is an algorithm to approximate the partition function to within ϵ -relative error and an ϵ -approximate sampling algorithm both running in time polynomial in n and $1/\epsilon$ for the hard-core model on the torus \mathbb{T}_n^d for even n .*

The value of $\lambda^*(d)$ we obtain is exponentially large in d , as in the results for slow mixing in [15]. We expect Theorem 1.5 to hold for much smaller λ^* , in particular with $\lambda^*(d) \rightarrow 0$ as $d \rightarrow \infty$ as in the proofs of phase coexistence in the hard-core model on \mathbb{Z}^d [26, 47]. See Section 7.

1.2.1. *Related results.* For graphs of maximum degree at most Δ a clear picture has emerged about the existence of an FPTAS for computing $Z_G(\lambda)$. A crucial role is played by the value $\lambda_c(\Delta) := \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^\Delta}$, the uniqueness threshold for the infinite d -regular tree. For $\lambda < \lambda_c(\Delta)$, Weitz [61] gave an FPTAS for approximating $Z_G(\lambda)$ on all graphs of max degree Δ . Conversely, Sly [56], Sly and Sun [57], and Galanis, Štefankovič, and Vigoda [24] showed that for $\lambda > \lambda_c(\Delta)$ there is no FPRAS for approximating $Z_G(\lambda)$ unless $\text{NP} = \text{RP}$, where RP is the class of problems that can be solved in polynomial time by a randomized algorithm.

The problem of counting independent sets on bipartite graphs is called $\#\text{BIS}$, and no such hardness result is known for $\#\text{BIS}$. Several important problems have been shown to be as hard as $\#\text{BIS}$ to approximate, including the problem of approximating the ferromagnetic Potts model partition function on general graphs [29, 19, 25]. The problem $\#\text{BIS}$ may be easier than the problem of approximating the hard-core partition function on general graphs: unlike on general graphs, finding the size of the *largest* independent set is easy on bipartite graphs. It is a major open problem in complexity theory to determine the complexity of $\#\text{BIS}$ [21].

1.3. **Overview of the algorithms.** The preceding theorems will be proven as applications of more general results about *polymer models* and *contour models*. We introduce polymer models in Section 2 below, and contour models in Section 3. In the current section, which gives an informal overview of our algorithms, we elide the distinction between polymers and contours, and for simplicity we will write contour models. The idea behind contour models is introduced in Section 1.3.1, we outline our approximation algorithms in Section 1.3.2, and lastly we describe our sampling algorithms in Section 1.3.3.

1.3.1. *Contour models.* For many discrete statistical mechanics models there are regimes in which the most likely configuration is simple to describe. For example, in the hard-core model the most likely configuration at low fugacities is the empty independent set, while at high fugacities the most likely configurations are the all-even or all-odd occupied independent sets. Contour models are a geometric way to represent spin models in terms of their deviations from these most likely configurations, which we will henceforth call *ground states*.

In the simplest settings such a representation involves re-writing a partition function as a sum over a suitable class of subgraphs. For example, this can be done for the high-temperature Ising model. In more complex situations, Pirogov–Sinai theory provides an appropriate representation. We defer the details of this to Section 3. For the purposes of this introduction it will suffice that the reader has in mind that a contour model expresses the partition function as a sum over collections of disjoint geometric objects.

1.3.2. *Approximation algorithms using contour models.* Our algorithm for approximating the partition function will be based on truncating the Taylor series for $\log Z_G$ after a given number of terms. There are several components to making this work:

- (1) We write the partition function as an abstract contour model as dictated by Pirogov–Sinai theory [49, 50].
- (2) We prove that the partition function, as a function of the inverse temperature, does not vanish outside a disc in the complex plane. We do this by using the

Peierls' condition and a theorem of Borgs and Imbrie [17] implementing Zahradník's version [62] of Pirogov–Sinai theory.

- (3) We use the absence of zeros to write error bounds for the truncated Taylor series for the log partition function, following Barvinok [4, 6].
- (4) We efficiently compute the low-order coefficients of the Taylor series. This is done inductively using the cluster expansion.

None of these components are wholly new – our main contribution is to establish the relevance of Pirogov–Sinai theory to the design of algorithms. In this paper we strive for simplicity and clarity of the main ideas, and so we do not try to pursue optimal bounds or maximal generality in stating theorems. We believe, however, that essentially any application of Pirogov–Sinai theory to prove phase coexistence or to prove slow mixing for discrete lattice spin models can be turned into efficient approximate counting and sampling algorithms with the ideas of this paper.

1.3.3. Sampling algorithms using contour models. Often efficient approximation algorithms lead to efficient sampling algorithms via self-reducibility. The basic idea is that if one can accurately approximate the partition function Z_G for arbitrary G with *arbitrary boundary conditions*, then one can accurately estimate the probability of a configuration by expressing it as a telescoping product of partition functions. The idea is already evident in the expression for the probability that a vertex v is occupied in the hard-core model:

$$\mathbb{P}_{G,\lambda}[v \text{ occupied}] = \lambda \frac{Z_{G \setminus N(v)}(\lambda)}{Z_G(\lambda)},$$

where $N(v)$ is the union of $\{v\}$ and the set of neighbours of v . This expression arises as v being occupied implies that no neighbour of v is occupied. We think of the numerator as being a partition function with a boundary condition that $N(v)$ is unoccupied.

The derivation of contour representations in Pirogov–Sinai theory makes use of particular boundary conditions: the padded boundary conditions introduced in Sections 1.1 and 1.2. This leads to a difficulty in using self-reducibility to define sampling algorithms, as changing the boundary conditions may lead to a situation in which we do not have a contour representation. We circumvent this difficulty by using the idea of self-reducibility on the level of contours: instead of iteratively determining a spin configuration spin by spin, we instead iteratively determine a contour configuration contour by contour. The manner in which contours are defined ensures that we are always able to write the partition functions that arise in terms of contour representations.

Obtaining a spin configuration from a contour configuration is straightforward, and we defer a discussion of this point until after we have defined contour models precisely.

1.4. Organization and Conventions. In Section 2 we define polymer models and present both the cluster expansion and Taylor series for the log partition function. Under the condition of a zero-free region of the partition function in the complex plane, we give an efficient algorithm for approximating the partition function of a polymer model.

In Section 3 we define the more sophisticated contour models from Pirogov–Sinai theory, and show that the algorithm of Section 2 can be applied to approximate the partition function of a contour model under suitable hypotheses. We discuss how to verify the main hypothesis, which is the convergence of the cluster expansion, in Section 4. By

using a theorem of Borgs and Imbrie [17] we verify this condition for the Potts model and the hard-core model.

In Section 5 we prove our main sampling results. Establishing our results for the torus \mathbb{T}_n^d requires some additional work and we carry this out in Section 6. In Section 7 we conclude with some directions for future work.

We end this section with some notation and conventions that will be used throughout. All logarithms are natural logarithms. If G is a graph we write $|G|$ for the size of the vertex set of G .

A finite subset $\Lambda \subset \mathbb{Z}^d$ is *c-connected* if Λ^c is connected under the adjacency relation derived from the distance function $d_\infty(x, y) = \max_{i=1}^d |x_i - y_i|$. We also call c-connected subsets *regions*. The *interior boundary* of a set $A \subset \mathbb{Z}^d$ is $\partial^{\text{in}} A = \{i \in A : d_\infty(i, A^c) = 1\}$. The *exterior boundary* of a set $A \subset \mathbb{Z}^d$ is $\partial^{\text{ex}} A = \{i \in A^c : d_\infty(i, A) = 1\}$. On the torus \mathbb{T}_n^d , with the vertex set viewed as $\{1, \dots, n\}^d$, we define the d_∞ distance in the natural way, with $d_\infty(x, y) = \max_{i=1}^d \min\{(x_i - y_i) \bmod n, (y_i - x_i) \bmod n\}$.

2. CLUSTER EXPANSIONS, TAYLOR SERIES, AND APPROXIMATE COUNTING

In this section we introduce polymer models and the cluster expansion, and describe how they can be used algorithmically. To illustrate the method we recover results of Patel–Regts [46] and Liu–Sinclair–Srivastava [42] on the efficient approximation of the hard-core and Ising models. The method of this section is at the heart of the proofs of our main results for more sophisticated contour models.

2.1. Polymer models. Let $G = (V, E)$ be a finite graph and let Ω be a finite set of *spins*. Define a *polymer* γ in G to be a pair $\gamma = (\bar{\gamma}, \omega_{\bar{\gamma}})$ where $\bar{\gamma}$, the *support* of the polymer, is a connected subgraph of G and $\omega_{\bar{\gamma}}: \bar{\gamma} \rightarrow \Omega$ is an assignment of a spin from Ω to each vertex in $\bar{\gamma}$. The *size* of a polymer is $|\bar{\gamma}|$. A *polymer model* consists of a set $\mathcal{C}(G)$ of polymers along with *weight functions* $w(\gamma, \cdot): \mathbb{C} \rightarrow \mathbb{C}$ for each polymer γ . We need one assumption about the weight functions:

Assumption 1. *The weight functions $w(\gamma, z)$ are analytic functions of z in a neighborhood of the origin of the complex plane, and there is an absolute constant $\rho > 0$ such that for each $\gamma \in \mathcal{C}(G)$ the first non-zero term in the Taylor series expansion of $w(\gamma, z)$ around zero is of order $k \geq |\bar{\gamma}| \rho$.*

Note that Assumption 1 implies $w(\gamma, 0) = 0$ for all γ with non-empty support.

We say two polymers $\gamma, \gamma' \in \mathcal{C}(G)$ are *compatible* if $d(\bar{\gamma}, \bar{\gamma}') > 1$, where $d(\cdot, \cdot)$ is the graph distance in G . Let $\mathcal{G}(G)$ be the collection of all finite sets of polymers from $\mathcal{C}(G)$ that are pairwise compatible, including the empty set of polymers.

The partition function associated to the polymer model defined by $\mathcal{C}(G)$ is

$$(1) \quad Z(G, z) := \sum_{\Gamma \in \mathcal{G}(G)} \prod_{\gamma \in \Gamma} w(\gamma, z)$$

where the term corresponding to the empty set of polymers is 1 by convention. We think of $Z(G, z)$ as a function of one complex variable z .

Example 1 (Hard-core model at low density). The hard-core model is the simplest model to describe as a polymer model. Polymers are single vertices, i.e., $\mathcal{C}(G) = V(G)$.

The spin set, which is superfluous in this simple example, is $\Omega = \{1\}$: every polymer receives the same spin 1, which is interpreted as meaning the vertex is ‘occupied’. The weight function of each polymer is $w(\gamma, z) = z$. Two polymers are compatible if their distance in the graph is more than 1, and so the sets of pairwise compatible polymers are exactly the independent sets of G , and the polymer partition function is the hard-core model partition function at fugacity z :

$$Z(G, z) = \sum_{\Gamma \in \mathcal{G}(G)} \prod_{\gamma \in \Gamma} w(\gamma, z) = \sum_{I \in \mathcal{I}(G)} z^{|I|} = Z_G(z).$$

Example 2 (Ising model with free boundary conditions and an external field). Consider the Ising model with free boundary conditions and an external field z . That is

$$Z_G(\beta, z) := \sum_{\sigma \in \{\pm 1\}^{V(G)}} z^{\sum_{v \in V(G)} \sigma(v)} \prod_{\{u, v\} \in E(G)} e^{\beta \sigma(u) \sigma(v)}.$$

Assume $|z| < 1$, so -1 spins are preferred. To obtain a polymer model representation we can express the partition function in terms of deviations from the all -1 configuration. That is, a polymer γ is a connected induced subgraph $\bar{\gamma}$ of vertices, all labeled $+1$. Then we can write

$$Z_G(\beta, z) = z^{-|G|} e^{\beta|E(G)|} \sum_{\Gamma \in \mathcal{G}(G)} \prod_{\gamma \in \Gamma} w(\gamma, z),$$

where, letting $\partial_e \bar{\gamma} = |\{\{u, v\} \in E(G) : u \in \bar{\gamma}, v \notin \bar{\gamma}\}|$, the weight function is

$$w(\gamma, z) = z^{2|\bar{\gamma}|} e^{-2\beta|\partial_e \bar{\gamma}|}.$$

2.2. The cluster expansion. The *cluster expansion* is the following formal power series representation for $\log Z(G, z)$, see, e.g., [38, 23]. Under suitable conditions, see Section 4 below, it is also an absolutely convergent power series representation.

$$(2) \quad \log Z(G, z) = \sum_{k \geq 1} \frac{1}{k!} \sum_{(\gamma_1, \dots, \gamma_k)} \phi(\gamma_1, \dots, \gamma_k) \prod_{i=1}^k w(\gamma_i, z).$$

The sum in (2) is over ordered k -tuples of polymers from $\mathcal{C}(G)$, and ϕ is the *Ursell function*, which we now define.

Let $H = H(\gamma_1, \dots, \gamma_k)$ be the *incompatibility graph* of polymers $\gamma_1, \dots, \gamma_k$, i.e., the graph on k vertices with an edge between γ_i and γ_j if and only if γ_i and γ_j are not compatible. Then

$$\phi(\gamma_1, \dots, \gamma_k) := \sum_{\substack{E \subseteq E(H) \\ \text{spanning, connected}}} (-1)^{|E|}.$$

The sum is over spanning and connected edge sets of H . Thus $\phi(\gamma_1, \dots, \gamma_k) = 0$ if H is disconnected. By definition, the Ursell function depends only on the graph H induced by the incompatibility relation, and not on the polymers $\gamma_1, \dots, \gamma_k$ themselves.

It will be convenient for us later to rewrite (2) as a sum over unordered multisets of polymers from $\mathcal{C}(G)$. Given a multiset $M = \{\gamma_1^{m_1}, \dots, \gamma_t^{m_t}\}$, there are exactly $\binom{k}{m_1 \dots m_t}$

k -tuples which have M as underlying multiset. Here the exponents m_i denote the multiplicities of the elements in M , and $k = \sum_{i=1}^t m_i$. We can therefore rewrite (2) as

$$\log Z(G, z) = \sum_{k \geq 1} \frac{1}{k!} \sum_{\{\gamma_1^{m_1}, \dots, \gamma_t^{m_t}\}} \binom{k}{m_1 \dots m_t} \phi(\gamma_1^{m_1}, \dots, \gamma_t^{m_t}) \prod_{i=1}^t w(\gamma_i, z)^{m_i},$$

where $\phi(\gamma_1^{m_1}, \dots, \gamma_t^{m_t})$ is the Ursell function applied to the incompatibility graph of the collection of polymers $\underbrace{\gamma_1, \dots, \gamma_1}_{m_1}, \underbrace{\gamma_2, \dots, \gamma_2}_{m_2}, \dots, \underbrace{\gamma_t, \dots, \gamma_t}_{m_t}$.

2.3. The Taylor series. We can also Taylor expand $\log Z(G, z)$ around $z = 0$:

$$\log Z(G, z) = \sum_{k \geq 1} \frac{z^k}{k!} \frac{\partial^k}{\partial z^k} \log Z(G, 0).$$

In fact, as observed by Dobrushin [20], the cluster expansion and Taylor series are the same power series in z , though arranged differently. By our assumptions on the weight functions, for each k only a finite number of terms in the cluster expansion contribute to the coefficient of z^k , and so we can compute the coefficients of the Taylor series via the cluster expansion:

$$(3) \quad \frac{\frac{\partial^k}{\partial z^k} \log Z(G, 0)}{k!} = \sum_{j=1}^k \frac{1}{j!} \sum_{(\gamma_1, \dots, \gamma_j)} \phi(\gamma_1, \dots, \gamma_j) \frac{1}{k!} \frac{\partial^k}{\partial z^k} \left(\prod_{i=1}^j w(\gamma_i, z) \right)_{z=0}.$$

2.4. Approximate counting for polymer models. The partial sums of the Taylor series are

$$T_m(G, z) := \sum_{k=1}^m \frac{z^k}{k!} \frac{\partial^k}{\partial z^k} \log Z(G, 0).$$

If we know $Z(G, z)$ is non-zero in a disc around the origin in the complex plane, then we can control the error of the truncated Taylor series approximation for $\log Z(G, z)$. This is the approach of Barvinok for devising approximation algorithms [2, 3, 6, 4]. The next lemma rephrases [46, Lemma 2.2] and indicates where to truncate the Taylor series to get a good approximation. We use the following notion of relative error for complex numbers.

Definition 4. An ϵ -relative approximation to a complex number $Z \neq 0$ is a complex number $\hat{Z} \neq 0$ so that

$$e^{-\epsilon} \leq \left| \frac{Z}{\hat{Z}} \right| \leq e^\epsilon$$

and the angle between Z and \hat{Z} as vectors in the complex plane is at most ϵ .

Lemma 2.1. Suppose the degree of the polynomial $Z(G, z)$ is at most N and suppose that $Z(G, z) \neq 0$ for all $|z| \leq \delta$. Then for every $\epsilon > 0$ and every $|z| < \delta$, $\exp[T_m(G, z)]$ is an ϵ -relative approximation to $Z(G, z)$ for all

$$m \geq \frac{\log(N/\epsilon)}{1 - |z|/\delta}.$$

Lemma 2.1 implies that if we can compute all of the coefficients $\frac{\partial^k}{\partial z^k} \log Z(G, 0)$ for $k = 1, \dots, m$ in time $\exp(O(m))$, then we obtain an algorithm to produce ϵ -relative approximations of $Z(G, z)$ with a running time polynomial in N and $1/\epsilon$ when $|z| < \delta$.

Definition. *We can compute a function $f(z)$ up to order m if we can compute the coefficients of the Taylor series of $f(z)$ around 0 up to order m .*

Theorem 2.2. *Fix Δ and let \mathfrak{G} be a set of graphs of degree at most Δ . Suppose:*

- *There is a constant C so that $Z(G, z)$ is a polynomial in z of degree at most $C|G|$ for all $G \in \mathfrak{G}$.*
- *The weight functions satisfy Assumption 1, and we can compute $w(\gamma, z)$ up to order m for all $G \in \mathfrak{G}$ and all $\gamma \in \mathcal{C}(G)$ in time $\exp(O(m + \log |G|))$.*
- *For every connected subgraph G' of every $G \in \mathfrak{G}$, we can list all polymers $\gamma \in \mathcal{C}(G)$ with $\bar{\gamma} = G'$ in time $\exp(O(|G'|))$.*
- *There exists $\delta > 0$ so that for all $|z| < \delta$ and all $G \in \mathfrak{G}$, $Z(G, z) \neq 0$.*

Then for every z with $|z| < \delta$, there is an FPTAS for $Z(G, z)$ for all $G \in \mathfrak{G}$.

The proof of Theorem 2.2 requires a few lemmas.

Lemma 2.3 (8). *There is an algorithm to list all rooted, unlabeled trees on at most m vertices that runs in time $\exp(O(m))$.*

Proof. Let $A(m)$ be the number of rooted unlabeled trees on m vertices; $A(m)$ is $\exp(O(m))$ [44], and Beyer and Hedetniemi [8] have given an algorithm that lists all such trees in time $\exp(O(m))$. \square \square

Let $\mathcal{C}_m(G) := \{\gamma \in \mathcal{C}(G) : |\bar{\gamma}| \leq m\}$ be the set of polymers of size at most m . If $|G| = n$ the next lemma shows $\mathcal{C}_m(G)$ can be enumerated in time $\exp(O(m + \log n))$.

Lemma 2.4. *Under the assumptions of Theorem 2.2 we can list all polymers $\gamma \in \mathcal{C}_m(G)$ in time $\exp(O(m + \log |G|))$.*

Proof. There are at most $\exp(O(m + \log |G|))$ such polymers, as (i) the support of a polymer is a connected subgraph of a bounded degree graph, and by [13, Lemma 9] there are $\exp((O(m + \log |G|)))$ of these, and (ii) by assumption we can list all polymers with a given support of size at most m in time $\exp(O(m))$. The list can be created in time $\exp(O(m + \log |G|))$ as in [46, Lemma 3.4]. \square \square

Lemma 2.5. *Under the assumptions of Theorem 2.2, for any polymer γ we can list all polymers γ' such that γ' is incompatible with γ and $|\bar{\gamma}'| \leq m$ in time $\exp(O(m + \log |\bar{\gamma}|))$.*

Proof. For each v such that $d(v, \bar{\gamma}) \leq 1$, we list all polymers of size at most m containing v , then remove duplicates. As in the proof of Lemma 2.4, this can be done in time $\exp(O(m + \log |\bar{\gamma}|))$. \square \square

The computation of the Ursell function of a graph on k vertices by naively summing over all spanning edge sets would take $\exp(O(k^2))$ time. The next lemma does better.

Lemma 2.6. *The Ursell function $\phi(H)$ can be computed in time $\exp(O(|H|))$.*

Proof. Let $\kappa((V, A))$ denote the number of connected components of a graph (V, A) . The *Tutte polynomial* of a connected graph $H = (V, E)$ on k vertices is

$$T_H(x, y) := \sum_{A \subseteq E} (x-1)^{\kappa((V, A))-1} (y-1)^{\kappa((V, A)) + |A| - k}$$

We can express $\phi(H)$ in terms of the Tutte polynomial:

$$\phi(H) = \sum_{A \subseteq E} \mathbf{1}_{\kappa((V, A))=1} \cdot (-1)^{|A|} = (-1)^{k-1} T_H(1, 0).$$

The coefficients of the Tutte polynomial $T_H(x, y)$ can be computed in time $3^k k^{O(1)}$ using an algorithm of Björklund, Husfeldt, Kaski, and Koivisto [10], and hence $T_H(1, 0)$ can be computed in this time. \square

Finally, we give a simple lemma about products of weight functions.

Lemma 2.7. *Let $w_1(z)$ and $w_2(z)$ be two weight functions. If we know $w_1(z)$ and $w_2(z)$ up to order m then we can compute the product $w_1(z)w_2(z)$ up to order m in time $O(m^2)$.*

Proof. It is a simple calculation to express the coefficients of w_1w_2 in terms of those of w_1 and w_2 . This takes $O(m^2)$ time. \square

Proof of Theorem 2.2. Let $n = |G|$ and set $m = \lceil \frac{\log(Cn/\epsilon)}{1-|z|/\delta} \rceil$, where C, δ are the constants from the hypotheses of the theorem. Recall the constant ρ of Assumption 1, and let $m' = \lceil m/\rho \rceil$. Note $m' = \Theta(m)$.

First we create a list of all polymers in $\mathcal{C}_{m'}(G)$, along with the Taylor series coefficients of $w(\gamma, z)$ of order at most m for all $\gamma \in \mathcal{C}_{m'}(G)$. These are the polymers and coefficients that can contribute to the order k coefficients of the Taylor series of $\log Z(G, z)$ for $k \leq m$. The list of polymers can be formed in time $\exp(O(m + \log |G|))$ by Lemma 2.4, and we can compute the coefficients of the weight functions up to order m in time $\exp(O(m) + \log |G|)$ by assumption. Sort this list by $|\bar{\gamma}|$ and call the sorted list \mathcal{L} : $\mathcal{L}(j)$ is the j th polymer in the list.

Next we create a list of all possible rooted unlabeled trees on at most m vertices and call the list \mathcal{T} . The list \mathcal{T} has length at most $\exp(O(m))$ and can be created in time at most $\exp(O(m))$ by Lemma 2.3.

For each rooted unlabeled tree $(T, r) \in \mathcal{T}$, we label its vertices with polymers from $\mathcal{C}_{m'}(G)$ with the condition that two vertices adjacent in the tree must be labeled by incompatible polymers. We do this by first listing all assignments of integers $\ell(v) \in \{1, \dots, m\}$ to the vertices $v \in T$ such that the sum of the $\ell(v)$ is at most m' . By a crude balls-and-boxes argument, there are at most $m' \cdot 2^{m'} = \exp(O(m))$ ways to do this.

Then for each such assignment, we first label the root r with a polymer γ with $|\bar{\gamma}| = \ell(r)$. Crudely, there are at most $|\mathcal{L}| = \exp(O(m + \log |G|))$ ways to do this. Proceeding down the tree we label each vertex $u \in T$ with a polymer γ' from \mathcal{L} so that $|\bar{\gamma}'| = \ell(u)$ and so that the polymer assigned to u is incompatible with the polymer γ assigned to the parent of u in the tree. By Lemma 2.5, there are at most $\exp(O(k + \log |\bar{\gamma}|))$ polymers γ' of size k incompatible with a given polymer γ , and so the total number of choices of polymers for all the nodes of the tree is at most $|\mathcal{L}| \cdot \prod_{(u,v) \in E(T)} \exp(O(\ell(v) + \log \ell(u))) = \exp(O(m + \log |G|))$.

Multiplying over all trees, all assignments $\ell(\cdot)$, and all assignments of polymers respecting ℓ and the incompatibility relation gives a total of at most $\exp(O(m + \log |G|))$ labeled trees, and so we can construct a list of all such possibilities in time at most $\exp(O(m + \log |G|))$.

Now for each labeled tree in this list, we retain only the multiset of polymers assigned to the nodes of the tree, and we sort this list. Call this list of multisets \mathcal{L}' . Each of these multisets is made up of polymers from $\mathcal{C}(G)$ and each has the property that its corresponding incompatibility graph is connected; that is, each corresponds to a *cluster* that contributes to the sum in (2). Moreover, each cluster contributing to (2) whose weight is $z^j(1 + O(z))$ for $j \leq m$ appears at least once in this list. This is because Assumption 1 implies that both a cluster of $\tilde{m} > m$ polymers and a cluster containing a polymer of size $\tilde{m} > m'$ are 0 up to order m . We can obtain a list that contains each possible cluster exactly once by removing all duplicate clusters from \mathcal{L}' . This takes time at most quadratic in the length of the list, which is $\exp(O(m + \log |G|))$.

For each cluster in \mathcal{L}' we can compute its incompatibility graph H in time $O(m^2)$, and we can compute the Ursell function $\phi(H)$ in time at most $\exp(O(m))$ by Lemma 2.6. We can also compute the product of the weight functions of the polymers in the cluster up to order m : since we have already computed the weight functions up to order m we can do this in time $O(m^3)$ by m applications of Lemma 2.7.

We then sum the coefficients of order k over all clusters in \mathcal{L}' to obtain the coefficient of z^k in the Taylor series for $\log Z(G, z)$ by (3). Evaluating $T_m(G, z)$ and exponentiating gives an ϵ -relative approximation to $Z(G, z)$ by Lemma 2.1. The total running time of the algorithm is $\exp(O(m + \log |G|))$. \square \square

Before applying this theorem to our examples, we record a remark that will be needed later.

Remark 1. *In the proof of Theorem 2.2 we only used the fourth hypotheses of the theorem to guarantee the accuracy of the approximation $\exp[T_m(G, z)]$ to $Z(G, z)$. In particular, this hypothesis was not used in the computation of the coefficients of $T_m(G, z)$.*

2.5. Examples. Theorem 2.2 allows us to recover the results of Patel and Regts [46], and independently Harvey, Srivastava and Vondrák [32], for the hard-core model and the results of Liu, Sinclair, and Srivastava [42] for the Ising model with non-zero external field. In both cases we get an FPTAS for these models on graphs with degree at most Δ . Let us briefly justify why Theorem 2.2 applies.

Example 3 (The hard-core model at low density). Recall Example 1. Let \mathfrak{G}_Δ be the set of graphs of maximum degree Δ . The first three conditions of Theorem 2.2 are straightforward to verify. For the fourth condition, Shearer's bound shows that $Z(G, z) \neq 0$ for all $|z| < \frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta}$ and $G \in \mathfrak{G}_\Delta$ [54, 53].

Example 4 (The Ising model with free boundary conditions and an external field). Recall Example 2. It suffices to approximate the polymer model

$$(4) \quad Z(G, z) = z^{|G|} e^{-\beta|E(G)|} Z_G(\beta, z),$$

and by swapping the roles of the +1 and -1 spins, it suffices to consider $|z| < 1$. The first three conditions of Theorem 2.2 are easily verified:

- $Z(G, z)$ is a polynomial of degree $2|G|$ in z .
- Polymers correspond to connected induced subgraphs of G . We can compute the weight functions of all polymers up to order m as follows. First, list all connected induced subgraphs of G of size at most m ; there are at most $\exp(O(m + \log |G|))$ of these and the list can be constructed in this time by Lemma 2.4. For each connected subgraph, the weight function can be computed in time $O(m)$ as it suffices to count $|\bar{\gamma}|$ and $|\partial_e \bar{\gamma}|$.
- For each connected induced subgraph G' of G there is exactly one polymer with support G' .

The fourth condition is provided by the Lee–Yang theorem [40]: for any G and any $\beta > 0$, $Z_G(\beta, z) \neq 0$ if $|z| < 1$. By (4), this implies $Z(G, z) \neq 0$ for $|z| < 1$ as well.

2.6. A generalization. We can generalize the definitions and results above, and this generalization will be useful in what follows. Let $\mathcal{S} \subseteq \mathcal{C}(G)$, and let $\mathcal{G}(\mathcal{S})$ be the collection of all finite sets of polymers from $\mathcal{C}(\mathcal{S})$ that are pairwise compatible, including the empty set of polymers. Abusing notation, we define

$$Z(\mathcal{S}, z) := \sum_{\Gamma \in \mathcal{G}(\mathcal{S})} \prod_{\gamma \in \Gamma} w(\gamma, z).$$

If we know $Z(G, z)$ has a zero-free disk about the origin, then we can efficiently approximate $Z(\mathcal{S}, z)$ for any $\mathcal{S} \subseteq \mathcal{C}(G)$.

Lemma 2.8. *Fix Δ and let \mathfrak{G} be a set of graphs of degree at most Δ . Suppose:*

- *There is a constant C so that $Z(\mathcal{S}, z)$ is a polynomial in z of degree at most $C|\mathcal{S}|$ for all $G \in \mathfrak{G}$ and all $\mathcal{S} \subseteq \mathcal{C}(G)$.*
- *We can compute $w(\gamma, z)$ up to order m for all $\gamma \in \mathcal{C}(G)$ in time $\exp(O(m + \log |G|))$.*
- *For every connected subgraph G' of every $G \in \mathfrak{G}$, we can list all polymers γ with $\bar{\gamma} = G'$ in time $\exp(O(|G'|))$.*
- *There exists $\delta > 0$ so that for all $|z| < \delta$ and all $G \in \mathfrak{G}$, the cluster expansion (2) is absolutely convergent.*

Then for every z with $|z| < \delta$, there is an FPTAS for $Z(\mathcal{S}, z)$ for all $G \in \mathfrak{G}$ and all $\mathcal{S} \subseteq \mathcal{C}(G)$.

The proof is a repetition of the proof of Theorem 2.2 together with one observation: for all $\mathcal{S} \subseteq \mathcal{C}(G)$, $Z(\mathcal{S}, z) \neq 0$ for $|z| < \delta$. This follows since the cluster expansion for $\log Z(G, z)$ is absolutely convergent, and the cluster expansion for $\log Z(\mathcal{S}, z)$ is a subseries so it too must be absolutely convergent.

2.7. Related results. The algorithm of Theorem 2.2 has strong similarity with the algorithms used in [46] and [42]. Both of these results use truncation of the Taylor series for $\log Z$ and the fact that the Taylor series are in some sense supported on connected graphs. Theorem 2.2 makes this notion of connectedness explicit and illustrates the connection to the cluster expansion. As a consequence our result uses analyticity of the weight functions, while the other approaches use more algebraic methods in combination with the Newton identities (see (10) and (11) below).

In the next section we will apply Theorem 2.2 to more sophisticated contour models. It is likely possible to apply the approach of [46] to contour models as well. We have elected to develop the cluster expansion approach as it gives us access to well-developed criteria for verifying the fourth condition of Theorem 2.2, as will be explained in Section 4.

A more careful analysis of our algorithm allows one to recover the result from [45] saying that one can compute the number of independent sets of size m in a bounded degree graph of order n in time $O(nc^m)$.

3. CONTOUR MODELS

A more sophisticated version of a polymer model is a *contour model*, and for this we specialize to \mathbb{Z}^d , $d \geq 2$. Our setup will be an amalgamation of those in [17] and [23, Chapter 7]. The main result is Theorem 3.1. We give examples of contour model representations of spin models in Section 3.5.

3.1. Contour models. Fix a finite set of spins Ω , and let Ξ be a finite set of *ground states*. In spin models ground states correspond to periodic assignments of spins to \mathbb{Z}^d that minimize energy, e.g., monochromatic configurations for the Potts model or the all even/all odd occupied configurations for the hard-core model, but at this level of generality they are just labels.

A contour γ is a pair $(\bar{\gamma}, \omega_{\bar{\gamma}})$; the *support* $\bar{\gamma}$ is a finite subset of \mathbb{Z}^d connected under the d_∞ distance and $\omega_{\bar{\gamma}}: \bar{\gamma} \rightarrow \Omega$ is an assignment of spins to the vertices of $\bar{\gamma}$. The support $\bar{\gamma}$ of a contour partitions $\mathbb{Z}^d \setminus \bar{\gamma}$ into maximal connected components, and in what follows we denote them by A_0, A_1, \dots, A_t , and we assume A_0 is the unique infinite component. Let $\text{ext}\gamma := A_0$ denote the *exterior* of γ and $\text{int}\gamma := \bigcup_{i=1}^t A_i$ denote the *interior* of γ .

A *contour model* is a set of contours \mathcal{C} , a *surface energy* $\|\gamma\| \in \mathbb{N}$ for each contour, and a *labeling function* $\text{lab}_\gamma(\cdot)$ for each contour. The labeling function lab_γ is a map from the collection of connected components $\{A_0, \dots, A_t\}$ to Ξ , the set of ground states. We will assume the labelling function is determined by the contour γ .

We will always make two basic assumptions on contour models. The first is about the computability of contours and their surface energies.

Assumption 2. *For every contour γ we can both determine if $\gamma \in \mathcal{C}$ and compute the labelling function $\text{lab}_\gamma(\cdot)$ in time $\exp(O(|\bar{\gamma}|))$. Moreover, for $\gamma \in \mathcal{C}$ we can compute $\|\gamma\|$ in time $\exp(O(|\bar{\gamma}|))$.*

Our second assumption relates the surface energy to the support of a contour. In applications the upper bound is typically trivial, while the lower bound is non-trivial and is known as the *Peierls' condition*.

Assumption 3. *There are constants $\rho, C > 0$ such that for all $\gamma \in \mathcal{C}$ the surface energy $\|\gamma\|$ is a positive integer satisfying the bound*

$$\rho|\bar{\gamma}| \leq \|\gamma\| \leq C|\bar{\gamma}|.$$

3.2. Partition functions of contour models. There are natural partition functions associated to contour models, and to introduce them we need a few more definitions. Two contours γ and γ' are *compatible* if $d_\infty(\bar{\gamma}, \bar{\gamma}') > 1$. A contour γ is of *type* φ if its

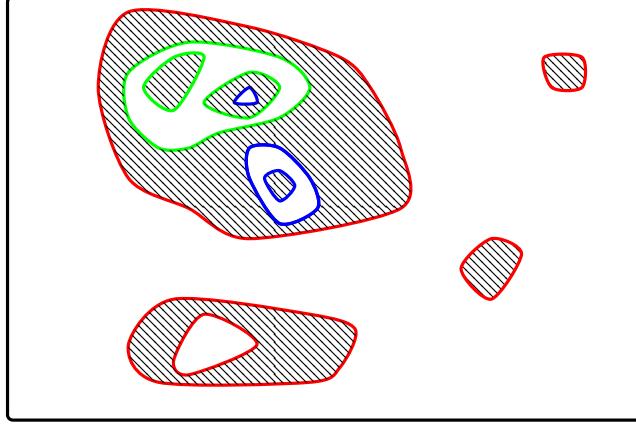


FIGURE 3. A schematic representation of the contour partition function. The boundary condition is red, the shaded set indicates the contours, and colors indicate the labels of the contours. Each connected component of the shaded set is a distinct contour.

exterior is labelled φ . The union of all interior regions of γ with label φ will be denoted

$$\text{int}_\varphi \gamma := \bigcup_{i \geq 1: \text{lab}_\gamma(A_i) = \varphi} A_i.$$

Let Γ be a set of compatible contours.

- (1) We say $\gamma \in \Gamma$ is *external* if $\bar{\gamma} \subset \text{ext} \gamma'$ for all $\gamma' \in \Gamma$, $\gamma' \neq \gamma$,
- (2) We say Γ is *matching and of type φ* if (i) all external contours have type φ , and (ii) either $|\Gamma| = 1$, or for each external contour $\gamma \in \Gamma$ and ground state φ' the subcollection of contours $\Gamma' \subset \Gamma$ whose support is contained in $\text{int}_{\varphi'} \gamma$ is matching and of type φ' .

Let $\mathcal{C}^\varphi \subset \mathcal{C}$ be the set of all contours of type φ , and for a region $\Lambda \subset \mathbb{Z}^d$, let $\mathcal{C}^\varphi(\Lambda)$ be the set of all contours γ of type φ so that $d_\infty(\bar{\gamma}, \Lambda^c) > 1$. We say these contours are *in* Λ . Let $\mathcal{G}_{\text{match}}^\varphi(\Lambda)$ be the collection of all sets of pairwise compatible contours in Λ that are matching and of type φ . Define

$$(5) \quad Z^\varphi(\Lambda, z) := \sum_{\Gamma \in \mathcal{G}_{\text{match}}^\varphi(\Lambda)} \prod_{\gamma \in \Gamma} z^{\|\gamma\|}.$$

We call this the *contour representation* of the partition function. It is clear from (5) that $Z(\Lambda, z)$ is a polynomial in z with constant term 1, and by Assumption 3 it is of degree at most $C|\Lambda|$. See Figure 3 for a schematic representation.

Let $\mathcal{G}_{\text{ext}}^\varphi(\Lambda)$ be the collection of all sets Γ of contours from $\mathcal{C}^\varphi(\Lambda)$ so that every $\gamma \in \Gamma$ is external. By fixing the outer contours in (5) and summing over all possible contours in their interior, we obtain the following inductive representation of Z^φ :

$$(6) \quad Z^\varphi(\Lambda, z) = \sum_{\Gamma \in \mathcal{G}_{\text{ext}}^\varphi(\Lambda)} \prod_{\gamma \in \Gamma} \left(z^{\|\gamma\|} \prod_{\varphi' \in \Xi} Z^{\varphi'}(\text{int}_{\varphi'} \gamma, z) \right),$$

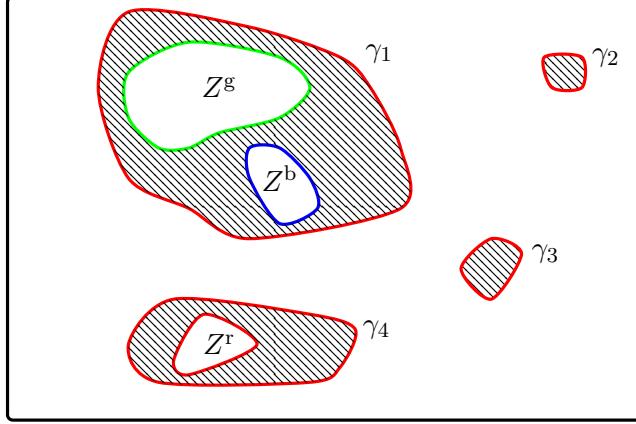


FIGURE 4. A schematic representation of a term of the outer contour partition function. The boundary condition is red, the shaded set indicates the contours, and colors indicate the labels of the contours. The contours γ_1 and γ_4 have interiors with interior partition functions Z^φ for $\varphi \in \{r, b, g\}$, while γ_2 and γ_3 do not have interiors.

which we call the *outer contour representation*. In obtaining (6) we have used that compatibility implies that the distance between contours is at least two, and hence any contour γ of type φ with $\bar{\gamma} \subset \text{int}_\varphi \gamma'$ belongs to $\mathcal{C}^\varphi(\text{int}_\varphi \gamma')$. The base case in (6) is a *thin region* Λ , i.e., one so that $\mathcal{G}_{\text{ext}}^\varphi(\Lambda) = \{\emptyset\}$, in which case $Z^\varphi(\Lambda, z) = 1$. See Figure 4 for a schematic representation.

There are well-known methods to convert discrete statistical physics models into contour representations [23, Chapter 7]. For the convenience of the reader we carry this out in Section 3.5 for the Potts and hard-core models.

3.3. Approximating the contour model partition function. Our main theorem is an algorithm to approximate the contour model partition function.

Theorem 3.1. *Fix $d \geq 2$ and $\varphi \in \Xi$, and suppose that:*

- *The contour model satisfies Assumptions 2 and 3.*
- *There exists $\delta > 0$ so that for $|z| < \delta$ and all regions $\Lambda \subset \mathbb{Z}^d$, $Z^\varphi(\Lambda, z) \neq 0$.*

Then for every z with $|z| < \delta$, there is an FPTAS for $Z^\varphi(\Lambda, z)$ for all regions $\Lambda \subset \mathbb{Z}^d$.

To prove this theorem we will view the outer contour model given by (6) as a polymer model. To make this precise, define the weight function of γ by

$$(7) \quad w^{\text{ext}}(\gamma, z) = z^{\|\gamma\|} \prod_{\varphi' \in \Xi} Z^{\varphi'}(\text{int}_{\varphi'} \gamma, z).$$

The outer contour representation can be rewritten as

$$(8) \quad Z^\varphi(\Lambda, z) = \sum_{\Gamma \in \mathcal{G}_{\text{ext}}^\varphi(\Lambda)} \prod_{\gamma \in \Gamma} w^{\text{ext}}(\gamma, z),$$

which matches the form of (1), except for the fact that the compatibility condition for external contours is not the notion of compatibility that was used for polymer models.

We will address this momentarily. Note that by construction $w^{\text{ext}}(\gamma, z)$ is a polynomial. By Assumption 3 $\|\gamma\| \geq \rho|\bar{\gamma}|$, and hence Assumption 1 is satisfied for these weights.

Two contours γ, γ' are *mutually external* if they are compatible, $\bar{\gamma} \subset \text{ext}\gamma'$, and $\bar{\gamma}' \subset \text{ext}\gamma$. This means neither contour lies in the interior of the other. Let

$$(9) \quad \text{cov}(\gamma) = \bar{\gamma} \cup \bigcup_{\varphi \in \Xi} \text{int}_{\varphi} \gamma.$$

Then two contours γ, γ' of type φ are mutually external if $d_{\infty}(\text{cov}(\gamma), \text{cov}(\gamma')) > 1$. We will use mutual externality as the notion of compatibility for the outer contour model; this replaces the notion of compatibility that was used for polymer models. The cluster expansion (2) holds for this notion of compatibility [23], and the proof of Theorem 2.2 goes through unchanged for this notion of compatibility given the following replacement for Lemma 2.5.

Lemma 3.2. *Suppose it is possible to determine if $\gamma \in \mathcal{C}(G)$ in time $\exp(O(|\bar{\gamma}|))$. Then for any contour γ we can list all contours $\gamma' \in \mathcal{C}(G)$ such that γ, γ' are not mutually external and $|\bar{\gamma}'| \leq m$ in time $\exp(O(m + \log |\bar{\gamma}|))$.*

Proof. We need to list all $\gamma' \in \mathcal{C}(G)$ of size at most m so that $d_{\infty}(\text{cov}(\gamma), \text{cov}(\gamma')) \leq 1$.

For each v such that $d_{\infty}(v, \text{cov}(\gamma)) \leq 1$, and each u such that $d_{\infty}(u, v) \leq m$, we list all d_{∞} -connected subgraphs of size at most m containing u and all assignments of spins from Ω to these subgraphs. This takes time $\exp(O(m))$ by [23, Lemma 9]. By hypothesis we can determine which of these contours are in \mathcal{C} in time $\exp(O(m))$, and hence for each v, u this list can be constructed in time $\exp(O(m))$.

There are at most $2 \cdot 3^d |\bar{\gamma}|^{d/(d-1)}$ [23, Lemma 7.28] such vertices v , and for each v at most $(2m+1)^d$ vertices u , and so the combination of all lists can be constructed in time $\exp(O(m + \log |\bar{\gamma}|))$. Finally for each γ' in the list, we check if $d_{\infty}(\text{cov}(\gamma), \text{cov}(\gamma')) \leq 1$. This can be done in time polynomial in $|\bar{\gamma}| \cdot |\bar{\gamma}'|$. \square \square

Theorem 3.1 will follow directly from Theorem 2.2 if we can verify the second hypothesis, i.e., if we can prove that the weight functions $w^{\text{ext}}(\gamma, z)$ can be computed up to order m for all $\gamma \in \mathcal{C}^{\varphi}(\Lambda)$ in time $\exp(O(m + \log |\Lambda|))$.

Lemma 3.3. *Under the assumptions of Theorem 3.1, we can compute the weight functions $w^{\text{ext}}(\gamma)$ up to order m for all $\varphi \in \Xi$ and all contours $\gamma \in \mathcal{C}_m^{\varphi}(\Lambda)$ in time $\exp(O(m + \log |\Lambda|))$.*

Before we prove Lemma 3.3 we need one useful fact, the Newton identities. Let $Z(z) = 1 + \sum_{k=1}^N e_k z^k$ be a polynomial, and let p_1, p_2, \dots be the coefficients of the Taylor series $\log Z(z) = \sum_{k \geq 1} p_k z^k$ around 0. The Newton identities imply the coefficients p_i can be expressed inductively in terms of the coefficients e_i , and vice-versa (cf. [46]):

$$(10) \quad p_k = -k e_k - \sum_{j=1}^{k-1} e_j p_{k-j},$$

$$(11) \quad e_k = -\frac{1}{k} \sum_{j=0}^{k-1} e_j p_{k-j}.$$

From this it follows that we can compute Z up to order m in time polynomial in m given the Taylor series coefficients of $\log Z$ up to order m and vice versa.

Proof of Lemma 3.3. We compute the weight functions inductively. Let

$$\begin{aligned}\mathcal{C}_m^\varphi(\Lambda) &:= \{\gamma \in \mathcal{C}^\varphi(\Lambda) : |\bar{\gamma}| \leq m\}, \quad \text{and} \\ \mathcal{C}_m(\Lambda) &:= \bigcup_{\varphi' \in \Xi} \mathcal{C}_m^{\varphi'}(\Lambda).\end{aligned}$$

We first give a polynomial-time algorithm to list and order \mathcal{C}_m such that if $\bar{\gamma}$ lies in the interior of γ' then γ comes before γ' in the ordering. In particular, the contours with thin interiors are at the front of the order. To do this we note that by Lemma 2.4 (using Assumption 2 in place of the third hypothesis of Theorem 2.2) we can list $\mathcal{C}_m(\Lambda)$ in time $\exp(O(m + \log |\Lambda|))$. For each γ we can determine the components of $\Lambda \setminus \bar{\gamma}$ in time $|\Lambda|$ by greedily growing the components of the complement. We can then decide how to order a pair $\{\gamma, \gamma'\}$ by checking if each $y \in \bar{\gamma}$ is contained in a single interior component of γ' or not and vice versa; this takes time $O(|\Lambda|m)$. Doing this for each pair of contours can be done in time quadratic in the length of the list, and hence the list can be ordered in time $\exp(O(m + \log |\Lambda|))$.

Given the ordered list, we will compute the weight functions $w^{\text{ext}}(\gamma)$ in order. The base cases are the contours with thin interiors for which $w^{\text{ext}}(\gamma) = z^{\|\gamma\|}$. By Assumption 2 these can each be computed in time in $\exp(O(|\bar{\gamma}|)) = \exp(O(m))$.

Now suppose we have computed the weight functions to order m for every contour γ' that precedes γ in the list. Then we can compute

$$w^{\text{ext}}(\gamma, z) = z^{\|\gamma\|} \prod_{\varphi' \in \Xi} Z^{\varphi'}(\text{int}_{\varphi'}(\gamma), z)$$

as follows. The surface energy can be computed in time $\exp(O(|\bar{\gamma}|))$ by Assumption 2. Each factor $Z^{\varphi'}(\text{int}_{\varphi'}(\gamma), z)$ is a polynomial in z whose first m coefficients can be computed in time $\exp(O(m + \log |\bar{\gamma}|))$ as follows. Recalling Remark 1, the proof of Theorem 2.2 (with Lemma 3.2 taking the place of Lemma 2.5) shows we can compute the first m coefficients of the Taylor series for $\log Z^{\varphi'}(\text{int}_{\varphi'}(\gamma), z)$ in the claimed time. The conditions of the theorem are satisfied since we have already written down to order m the weight function of any contour that can appear in the interior of γ . We can then use the Newton identities (11) to compute the coefficients of $Z^{\varphi'}(\text{int}_{\varphi'}(\gamma), z)$ from the Taylor series coefficients of $\log Z^{\varphi'}(\text{int}_{\varphi'}(\gamma), z)$. Multiplying these factors together, of which there are at most $\exp(O(\log |\bar{\gamma}|))$, and applying Lemma 2.7 shows that we can compute $w^{\text{ext}}(\gamma)$ to order m in time $\exp(O(m + \log |\bar{\gamma}|))$.

The time to compute each weight function to order m is therefore at most $\exp(O(m + \log |\Lambda|))$, and so the total time to compute all weight functions is at most $\exp(O(m + \log |\Lambda|))$ as well. \square \square

Proof of Theorem 3.1. We apply Theorem 2.2 with the class of bounded degree graphs \mathfrak{G} being subgraphs of \mathbb{Z}^d with the d_∞ -distance. The first two hypotheses of the theorem are true by the remarks following (5) and (8) and Lemma 3.3. The third and fourth hypotheses are the first part of Assumption 2 and an assumption of the Theorem, respectively. \square \square

3.4. A slight generalization. As in Section 2.6 we generalize the definitions and results slightly. We will use this generalization in the sampling algorithm of Section 5.

Let $\mathcal{S} \subseteq \mathcal{C}^\varphi(\Lambda)$ for some region Λ . Then define $\mathcal{G}_{\text{ext}}^\varphi(\mathcal{S})$ as the collection of all sets of compatible and mutually external contours from \mathcal{S} . Define

$$Z^\varphi(\mathcal{S}, z) := \sum_{\Gamma \in \mathcal{G}_{\text{ext}}^\varphi(\mathcal{S})} \prod_{\gamma \in \Gamma} w^{\text{ext}}(\gamma, z).$$

Our approximate counting algorithm extends to this generalization.

Lemma 3.4. *Fix $d \geq 2$ and suppose the following:*

- *The contour model satisfies Assumptions 2 and 3.*
- *There exists $\delta > 0$ so that for all $|z| < \delta$, all regions $\Lambda \subset \mathbb{Z}^d$, and all $\varphi \in \Xi$, the cluster expansion for $\log Z^\varphi(\Lambda, z)$ converges absolutely.*

Then for every z with $|z| < \delta$, there is an FPTAS for $Z^\varphi(\mathcal{S}, z)$ for all regions $\Lambda \subset \mathbb{Z}^d$ and all $\mathcal{S} \subseteq \mathcal{C}(\Lambda)$.

As in Section 2.6, it is enough to observe that absolute convergence of the cluster expansion for $\log Z^\varphi(\Lambda, z)$ implies absolute convergence of the cluster expansion for $\log Z^\varphi(\mathcal{S}, z)$.

3.5. Examples. In this section we introduce the contour representations that will be used in the proofs of our main theorems.

Example 5 (The ferromagnetic Potts model). For the ferromagnetic Potts model with no external field the set of ground states is the set of spins (or *colors*) $\Xi = \Omega = [q]$. Recall the padded monochromatic boundary conditions from Section 1.1: for a region Λ and a color $\varphi \in [q]$, the set of allowed configurations is

$$\Omega_\Lambda^\varphi = \{\omega \in [q]^\Lambda : \omega_i = \varphi \forall i \text{ s.t. } d_\infty(i, \Lambda^c) \leq 2\}.$$

We say a vertex $i \in \Lambda$ is *correct* with respect to $\omega \in \Omega_\Lambda^\varphi$ if there exists $\varphi' \in [q]$ so that $\omega_j = \varphi'$ for all $j \in \Lambda$ such that $d_\infty(i, j) \leq 1$; that is, i and its d_∞ neighbors all receive the same color. All other vertices of Λ are *incorrect* with respect to ω . The *boundary* $\Gamma(\omega)$ is the set of all incorrect vertices with respect to ω . See Figure 5. Each connected component (with respect to the d_∞ distance) of $\Gamma(\omega)$ defines the support $\bar{\gamma}$ of a contour γ , and ω_γ is the restriction of ω to γ . By the definition of Ω_Λ^φ we have $d_\infty(\bar{\gamma}, \Lambda^c) > 1$ for all contours.

It is a non-trivial fact that for each contour γ and each connected component A of $\mathbb{Z}^d \setminus \bar{\gamma}$, the set of vertices $i \in A$ such that $d_\infty(i, \bar{\gamma}) = 1$ is connected under the d_∞ distance [23, Appendix B.15] (see also [23, Lemma 7.19]). This implies there exists a φ' such that $\omega_i = \varphi'$ for all such i ; the label of A is φ' . This defines the set of contours and their labelling functions. Note the set of contours $\Gamma(\omega)$ is matching and of type φ .

Conversely, let $\bar{\gamma}$ be a d_∞ -connected subset of Λ so that $d(\bar{\gamma}, \Lambda^c) > 1$. Let $\omega_{\bar{\gamma}}$ be an assignment of spins to $\bar{\gamma}$ so that:

- For every $i \in \bar{\gamma}$, there is a $j \in \bar{\gamma}$, $d_\infty(i, j) = 1$ so that $\omega_{\bar{\gamma}, i} \neq \omega_{\bar{\gamma}, j}$.
- Let A_0, \dots, A_t denote the connected components of $\mathbb{Z}^d \setminus \bar{\gamma}$, with A_0 the unique infinite component. For each i there is a spin $\varphi' := \text{lab}_\gamma(A_i)$ so that $\omega_{\bar{\gamma}, j} = \varphi'$ for all $j \in \bar{\gamma}$, $d_\infty(j, A_i) = 1$. Moreover, $\text{lab}_\gamma(A_0) = \varphi$.

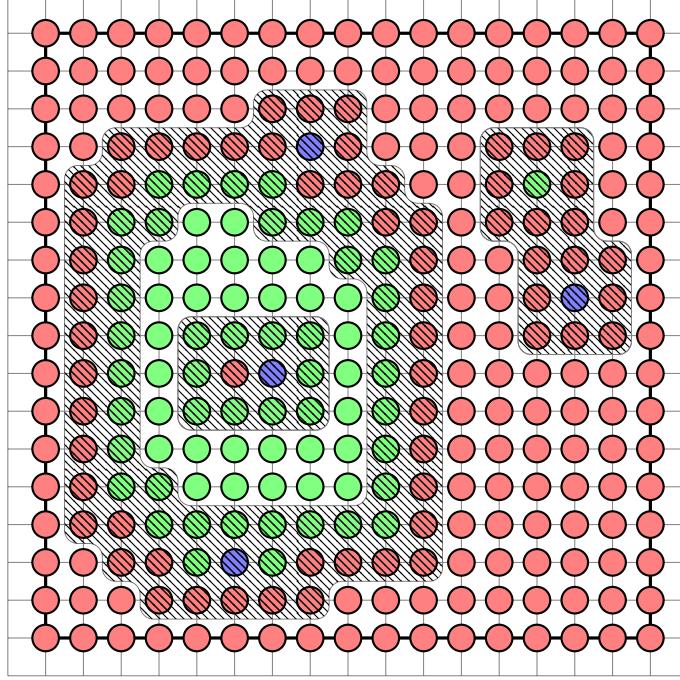


FIGURE 5. A 3-state Potts model configuration with padded red boundary conditions. Incorrect vertices and the contours they define are indicated by shading.

Any contour satisfying these conditions belongs to the set $\mathcal{C}^\varphi(\Lambda)$ and can be realized by a configuration $\omega \in \Omega_\Lambda^\varphi$ by setting $\omega_j = \text{lab}_\gamma(A_i)$ for any $j \in \Lambda \cap A_i$ and $\omega_j = \omega_{\bar{\gamma},j}$ for any $j \in \bar{\gamma}$. Iterating this construction shows that any set of matching contours $\Gamma \in \mathcal{G}_{\text{match}}^\varphi(\Lambda)$ of type φ gives rise to a Potts configuration $\omega \in \Omega_\Lambda^\varphi$.

Recall that we write $E(H) \subset E(\mathbb{Z}^d)$ for the set of edges of a subgraph H of \mathbb{Z}^d . Define the surface energy of a contour γ by

$$(12) \quad \|\gamma\| = \sum_{\{i,j\} \in E(\bar{\gamma})} \mathbf{1}_{\omega_i \neq \omega_j},$$

This is a positive integer by the definition of the boundary. Note also that we can check whether an assignment satisfies the condition of a contour and compute $\|\gamma\|$ in time linear in $|\bar{\gamma}|$, which shows Assumption 2.

Letting $z = e^{-\beta}$, (12) yields an expression for the Potts partition function:

$$\begin{aligned} Z_{q,\Lambda}^\varphi(\beta) &= \sum_{\omega \in \Omega_\Lambda^\varphi} e^{\beta \sum_{\{i,j\} \in E(\Lambda)} \mathbf{1}_{\omega_i = \omega_j}} \\ &= \sum_{\Gamma \in \mathcal{G}_{\text{match}}^\varphi(\Lambda)} z^{-|E(\Lambda)|} \prod_{\gamma \in \Gamma} z^{\|\gamma\|} \\ &= z^{-|E(\Lambda)|} Z^\varphi(\Lambda, z) \end{aligned}$$

where $Z^\varphi(\Lambda, z)$ is the contour model partition function defined in (5).

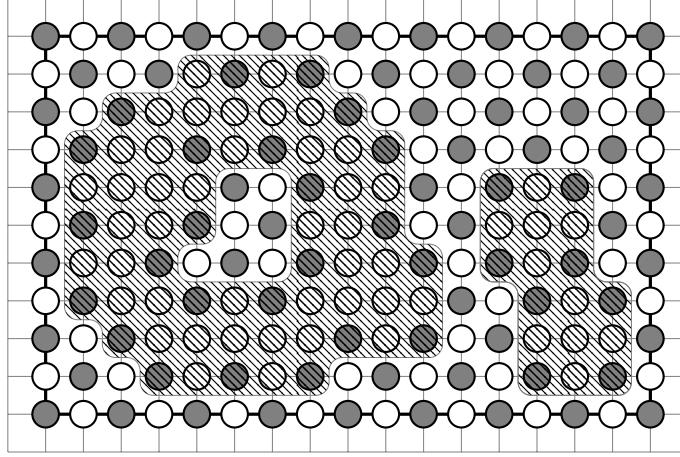


FIGURE 6. A hard-core model configuration with padded even boundary conditions. Incorrect vertices and the contours they define are indicated by shading.

Lastly we must show that Assumption 3 is satisfied. The upper bound is immediate, as each vertex has only $2d$ neighbors. A crude lower bound can be obtained by noting that for $v \in \bar{\gamma}$, there must be a u with $d_\infty(u, v) = 1$ such that $\omega_u \neq \omega_v$. Removing all vertices at d_∞ distance at most 1 from u and v , the same holds true for the remaining vertices of $\bar{\gamma}$. This implies $\|\gamma\| \geq \lceil |\bar{\gamma}|/(2 \cdot 3^d) \rceil$.

Example 6 (The hard-core model). We can express the hard-core model as a contour model in a similar way. We set $\Omega = \{0, 1\}$ and $\Xi = \{\text{even, odd}\}$. It will be convenient to identify independent sets $I \in \mathcal{I}(\Lambda)$ with their characteristic vectors $\omega^I \in \{0, 1\}^\Lambda$. In particular we define $\omega^{\text{even}} \in \mathcal{I}(\mathbb{Z}^d)$ by $\omega^{\text{even}} = \mathbf{1}_i$ is even, and similarly for ω^{odd} . The set of valid configurations for the even padded boundary conditions is

$$\Omega_\Lambda^{\text{even}} = \{\omega \in \{0, 1\}^\Lambda : \omega_i = \omega_i^{\text{even}} \text{ if } d_\infty(i, \Lambda^c) \leq 2\}.$$

We say a vertex $i \in \Lambda$ is *correct* with respect to $\omega \in \Omega_\Lambda^{\text{even}}$ if either $\omega_j = \omega_j^{\text{even}}$ for all $j \in \Lambda$ such that $d_\infty(i, j) \leq 1$ or $\omega_j = \omega_j^{\text{odd}}$ for all $j \in \Lambda$ such that $d_\infty(i, j) \leq 1$. All other vertices of Λ are incorrect. Again $\Gamma(\omega)$ is the set of all incorrect vertices with respect to ω , and each connected component (with respect to the d_∞ distance) of $\Gamma(\omega)$ is the support $\bar{\gamma}$ of a contour γ , and ω_γ is the restriction of ω to γ . See Figure 6 for an illustration. Again we have $d(\bar{\gamma}, \Lambda^c) > 1$ for all contours γ . For each contour γ and each connected component A of $\mathbb{Z}^d \setminus \bar{\gamma}$ either $\omega_i = \omega_i^{\text{even}}$ for all $i \in A$ such that $d_\infty(i, \bar{\gamma}) = 1$ or $\omega_i = \omega_i^{\text{odd}}$ for all $i \in A$ such that $d_\infty(i, \bar{\gamma}) = 1$; this again relies on [23, Appendix B.15] as in Example 5. In the first case, $\text{lab}_\gamma(A) = \text{even}$ and in the second, $\text{lab}_\gamma(A) = \text{odd}$. The set $\mathcal{C}^{\text{even}}(\Lambda)$ consists of all possible contours γ of type even with $d_\infty(\bar{\gamma}, \Lambda^c) > 1$.

Analogously to the Potts model, each configuration $\omega \in \Omega_\Lambda^{\text{even}}$ corresponds to a matching set of contours $\Gamma(\omega)$ of even type and each set of matching contours $\Gamma \in \mathcal{G}_{\text{match}}^{\text{even}}$ corresponds to a configuration $\omega \in \Omega_\Lambda^{\text{even}}$.

Given $A \subset \Lambda$, let A^{even} denote the set of even vertices of A . We define the surface energy of γ to be

$$(13) \quad \|\gamma\| = \frac{1}{4d} \sum_{\substack{i \in \bar{\gamma} \\ \omega_{\bar{\gamma},i}=0}} (2d - \sum_{j \in N(i)} \omega_{\bar{\gamma},j}),$$

where $N(i)$ is the set of neighbors of i in \mathbb{Z}^d . The surface energy is completely determined by $\bar{\gamma}$ and $\omega_{\bar{\gamma}}$. Let $\Gamma(\omega^I)$ denote the set of contours determined by the configuration ω^I . A double counting argument shows that

$$(14) \quad |I| = |\Lambda^{\text{even}}| - \sum_{\gamma \in \Gamma(\omega^I)} \|\gamma\|.$$

Since each contour γ can arise from a hard-core configuration, this formula shows $\|\gamma\|$ is integer valued. We can determine if a given assignment of spins to a d_∞ -connected subgraph $\bar{\gamma}$ satisfies the definition of a contour, and can compute $\|\gamma\|$ in linear time. This shows Assumption 2 holds.

Let $z = 1/\lambda$. Using (14) we obtain

$$\begin{aligned} Z_\Lambda^{\text{even}}(\lambda) &= \sum_{\omega \in \Omega_\Lambda^{\text{even}}} \lambda^{|\Lambda^{\text{even}}|} \prod_{\gamma \in \Gamma(\omega)} \lambda^{-\|\gamma\|} \\ &= z^{-|\Lambda^{\text{even}}|} \sum_{\Gamma \in \mathcal{G}_{\text{match}}^{\text{even}}(\Lambda)} \prod_{\gamma \in \Gamma} z^{\|\gamma\|} \\ &= z^{-|\Lambda^{\text{even}}|} Z^{\text{even}}(\Lambda, z), \end{aligned}$$

where $Z^{\text{even}}(\Lambda, z)$ is the contour model partition function.

Assumption 3 is also satisfied. The upper bound follows as $\|\gamma\| \leq |\bar{\gamma}|$ as each i can contribute at most 1 to the sum in (13). For the lower bound we have $\|\gamma\| \geq \frac{|\bar{\gamma}|}{4d \cdot 2 \cdot 3^d}$; this is a crude bound obtained by using that for every incorrect vertex i there must be a j with $d_\infty(i, j) \leq 1$ so that j is unoccupied and has an unoccupied neighbor.

4. CONVERGENCE OF THE CLUSTER EXPANSION

To apply Theorems 2.2 or 3.1 requires knowing that the partition function is non-zero in a disc around the origin in the complex plane. Occasionally, recall Section 2.5, this is provided by model-specific results. More generally, however, there are criteria for polymer and contour models that guarantee the partition function is non-zero in a disc around the origin.

The following theorem gives a criterion for the convergence of the cluster expansion; it is a special case of a result of Kotecký and Preiss [38]. The theorem says that if the weights decay at fast enough exponential rates, then the partition function is non-vanishing in some disc. For refined criteria, see [22].

Theorem 4.1 (Kotecký and Preiss [38]). *Suppose that for every $\gamma \in \mathcal{C}(G)$,*

$$(15) \quad \sum_{\gamma' \approx \gamma} |w(\gamma', z)| e^{|\bar{\gamma}'|} \leq |\bar{\gamma}|,$$

where the sum is over all polymers γ' incompatible with γ . Then the cluster expansion for $\log Z(G, z)$ converges absolutely and, in particular, $Z(G, z) \neq 0$.

Example 7 (Hard-core model at low density). Recall Example 1. We can apply Theorem 4.1 to the polymer representation of the hard-core model on graphs of maximum degree Δ . Polymers have size 1 and are incompatible with at most $\Delta + 1$ polymers; the $+1$ accounts for incompatibility with itself. Equation (15) becomes $(\Delta + 1)|z|e \leq 1$, or

$$|z| \leq \frac{1}{e(\Delta + 1)}.$$

This radius of convergence is not sharp; recall Example 3. It is, however, asymptotically sharp, since $\frac{(\Delta-1)^{\Delta-1}}{\Delta^\Delta} \sim \frac{1}{e\Delta}$ as $\Delta \rightarrow \infty$. For more on zero-free regions of the hard-core partition function see [53, 48].

We cannot apply a result like Theorem 4.1 to the outer contour model of Section 3 with weights given by (7), as these weight functions generally grow exponentially in the size of a contour and its interior. Instead we use a standard trick in Pirogov–Sinai theory.

Define the weight function

$$(16) \quad w^\varphi(\gamma, z) := z^{\|\gamma\|} \prod_{\varphi' \in \Xi} \frac{Z^{\varphi'}(\text{int}_{\varphi'}(\gamma), z)}{Z^\varphi(\text{int}_{\varphi'}(\gamma), z)}.$$

Then we can rewrite (6) as

$$(17) \quad Z^\varphi(\Lambda, z) = \sum_{\Gamma \in \mathcal{G}_{\text{ext}}^\varphi(\Lambda)} \prod_{\gamma \in \Gamma} \left(w^\varphi(\gamma, z) \prod_{\varphi' \in \Xi} Z^\varphi(\text{int}_{\varphi'}(\gamma), z) \right),$$

and now the partition function $Z^\varphi(\text{int}_{\varphi'}(\gamma), z)$ inside the product can be written using (17) again. Iterating this yields

$$(18) \quad Z^\varphi(\Lambda, z) = \sum_{\Gamma \in \mathcal{G}^\varphi(\Lambda)} \prod_{\gamma \in \Gamma} w^\varphi(\gamma, z),$$

where $\mathcal{G}^\varphi(\Lambda)$ is the collection of all subsets of contours from $\mathcal{C}^\varphi(\Lambda)$ that are pairwise compatible (but are no longer required to be mutually external). We call (18) the *polymer representation* of the partition function. Note however, that unlike the outer contour representation, there is not a mapping from the sets of contours appearing in the sum in (18) to spin configurations.

The polymer representation is of exactly the same form as the polymer partition function (1), but with a different weight function and the restriction that all contours have type φ . Moreover, the weight functions $w^\varphi(\gamma, z)$ satisfy the condition on the weight functions in the polymer model: the first non-zero Taylor series coefficient of $w^\varphi(\gamma, z)$ is of order at least $|\gamma|\rho$.

In the remainder of this section we indicate a method for proving the convergence of the cluster expansion for contour models with weight functions given by (16). The method, due to Borgs and Imbrie [17], is based on Zahradník’s truncation-based approach to Pirogov–Sinai theory [62, 49].

Assumption 4. *The surface energy function $\|\cdot\|$ and the labeling function are translation invariant, i.e., if there is an $a \in \mathbb{Z}^d$ such that $\bar{\gamma}' = \bar{\gamma} + a$ and $\omega_{\bar{\gamma}'}(i) = \omega_{\bar{\gamma}}(i - a)$ for all $i \in \bar{\gamma}'$, then they have the same surface energy and the labelling function respects the translation.*

To state the result of Borgs and Imbrie we must define the notion of a stable contour and a stable ground state. Recall from Section 3 that Ξ denotes the finite set of ground states. A contour γ of type φ is *stable* if

$$\frac{Z^{\varphi'}(\text{int}_{\varphi'}(\gamma), z)}{Z^{\varphi}(\text{int}_{\varphi'}(\gamma), z)} \leq e^{4|\partial^{\text{ex}} \text{int}_{\varphi'}(\gamma)|}$$

for all $\varphi' \in \Xi$. Let $\mathcal{G}_{\text{stab}}^{\varphi}(\Lambda)$ be the collection of all sets of pairwise compatible, stable contours from $\mathcal{C}^{\varphi}(\Lambda)$. The *truncated partition* is

$$Z_{\text{trun}}^{\varphi}(\Lambda, z) := \sum_{\Gamma \in \mathcal{G}_{\text{stab}}^{\varphi}(\Lambda)} \prod_{\gamma \in \Gamma} w^{\varphi}(\gamma, z).$$

If Peierls' condition holds and $|z|$ is small enough then the cluster expansion for the truncated partition function converges, and hence the limiting free energy of the truncated partition functions exists for each ground state $\varphi \in \Xi$, i.e.,

$$f(\varphi) := \lim_{\Lambda \rightarrow \mathbb{Z}^d} \frac{1}{|\Lambda|} \log Z_{\text{trun}}^{\varphi}(\Lambda, z)$$

exists when the limit is taken in the sense of van Hove.¹ A *stable* ground state φ is one for which $\text{Re } f(\varphi) \geq \text{Re } f(\varphi')$ for all $\varphi' \in \Xi$. In particular, at least one stable ground state exists.

Theorem 4.2 (Borgs, Imbrie [17]). *Fix $d \geq 2$. Suppose a contour model satisfies Assumptions 3 and 4. Then there exists a constant $\delta = \delta(d, \rho, \Xi) > 0$ so that for all $z \in \mathbb{C}$ with $|z| < \delta$, all regions Λ , and all stable ground states φ , the weights w^{φ} satisfy (15). In particular, the cluster expansion for $\log Z^{\varphi}(\Lambda, z)$ converges absolutely, and $Z^{\varphi}(\Lambda, z) \neq 0$.*

Proof. We must explain why the analysis of [17] applies when $d \geq 2$ and Assumptions 3 and 4 hold. This is essentially immediate as these assumptions constitute [17] Equation (2.1)], which is the assumption used in [17]. Two further remarks are in order. First, while the setup discussed in the introduction to [17] takes place in \mathbb{R}^d , the analysis applies to partition functions that can be expressed in the algebraic form of [17] Equation (2.6)]. Second, contours in [17] are geometric objects with phase labels, while our contours additionally have spins assigned to vertices. This does not cause any complication as it only modifies the exponential growth rate of the number of contours. \square \square

For this paper we do not need to go into the details of proving the stability of particular ground states: for the Potts and hard-core models symmetry ensures all ground states are stable. Thus by combining Theorem 4.2 with Theorem 3.1 we can prove the FPTAS portions of Theorems 1.4 and 1.2.

¹This means $\frac{|\partial^{\text{in}} \Lambda_n|}{|\Lambda_n|} \rightarrow 0$, see [23, Section 3.2.1]

Proof of Theorem 1.4, FPTAS part. By Example 6 the hard-core model satisfies Assumptions 2, 3, and 4, and hence by Theorem 4.2 there is a zero-free region for the partition function. The result then follows by Theorem 3.1. \square

For arbitrary q and β sufficiently large, the proof of the FPTAS portion of Theorem 1.2 is exactly analogous to that of Theorem 1.4. Example 5 verifies Assumptions 2, 3, and 4, and we obtain a zero-free region from Theorem 4.2.

5. SAMPLING

This section introduces a notion of self-reducibility based on polymers and contours. When combined with the approximate counting algorithms of Theorems 2.2 and 3.1 this yields efficient sampling algorithms.

5.1. Sampling from a polymer model. We will first introduce an algorithm to sample from a polymer model, then use a very similar algorithm to sample from a contour model.

In order to sample from a polymer model we need one further assumption; we note this assumption is simple to verify in the examples of polymer models (Examples 1 and 2) that we have seen so far.

Assumption 5. *For $z > 0$ the weights $w(\gamma, z)$ are non-negative real numbers for all polymers γ . Moreover, we can compute an ϵ -relative approximation to $w(\gamma, z)$ in time polynomial in $|\gamma|$ and $1/\epsilon$.*

Throughout this section we will assume Assumption 5 holds. In this case, given a polymer model on a graph G and a real number $z > 0$ the probability measure μ_G associated to the polymer model is

$$\mu_G(\Gamma) := \frac{\prod_{\gamma \in \Gamma} w(\gamma, z)}{Z(G, z)}, \quad \Gamma \in \mathcal{G}(G),$$

where $Z(G, z)$ is the polymer partition function defined in (1). Under the conditions for which we obtain an FPTAS for $Z(G, z)$ we obtain an efficient sampling algorithm.

Theorem 5.1. *Under the conditions of Lemma 2.8 and Assumption 5, for any positive real number $0 < z < \delta$ there is an efficient sampling algorithm for μ_G for all $G \in \mathfrak{G}$.*

We will begin by describing an idealized sampling algorithm which returns an exact sample from μ_G by sampling a configuration Γ one polymer at a time. We will then describe how to turn this into an efficient approximate sampling algorithm.

For a set of vertices $S \subset V(G)$ and a collection of compatible polymers $\Gamma \in \mathcal{G}(G)$, let $\mathcal{C}_{\Gamma, S}$ be the set of polymers γ given by

$$(19) \quad \mathcal{C}_{\Gamma, S} := \{\gamma \in \mathcal{C}(G) : \bar{\gamma} \cap S = \emptyset, \gamma \cup \Gamma \in \mathcal{G}(G)\}.$$

For a vertex $x \in V(G)$, let $\mathcal{C}_{\Gamma, S}(x) \subset \mathcal{C}_{\Gamma, S}$ be the subset of polymers γ such that $x \in \bar{\gamma}$. Note that if $\gamma, \gamma' \in \mathcal{C}_{\Gamma, S}(x)$ then they are incompatible.

Let γ_\emptyset be the empty polymer, and set $w(\gamma_\emptyset, z) := 1$. Let $\mu_{\Gamma, S, x}$ be the probability measure on $\mathcal{C}_{\Gamma, S}(x) \cup \gamma_\emptyset$ defined by

$$\mu_{\Gamma, S, x}(\gamma) := \frac{w(\gamma, z) Z(\mathcal{C}_{\Gamma \cup \gamma, S \cup x}, z)}{Z(\mathcal{C}_{\Gamma, S}, z)},$$

where we recall the notation $Z(\mathcal{S}, z)$ from Section 2.6. To verify this is a probability measure, note the so-called *fundamental identity* [53]:

$$Z(\mathcal{C}_{\Gamma, S}, z) = \sum_{\gamma \in \mathcal{C}_{\Gamma, S}(x) \cup \{\gamma_\emptyset\}} w(\gamma, z) Z(\mathcal{C}_{\Gamma \cup \gamma, S \cup x}, z).$$

Algorithm 1. Set $\Gamma_0 = \emptyset$, $S_0 = \emptyset$, and order the vertices of G by x_1, \dots, x_n . Repeat the following procedure for $t = 0$ to $n - 1$:

- (1) Sample γ from the measure $\mu_{\Gamma_t, S_t, x_{t+1}}$.
- (2) Set $\Gamma_{t+1} = \Gamma_t \cup \gamma$.
- (3) Set $S_{t+1} = S_t \cup x_{t+1}$.

Return $\Gamma = \Gamma_n$.

Lemma 5.2. *The distribution of Γ returned by Algorithm 1 is exactly $\mu_G(\Gamma)$.*

Proof. By construction the algorithm only outputs collections $\Gamma = \{\gamma_1, \dots, \gamma_k\}$ of polymers that belong to $\mathcal{G}(G)$, so it suffices to compute the probability the algorithm outputs a particular $\Gamma \in \mathcal{G}(G)$.

We first claim that each $\gamma \in \mathcal{C}(G)$ has at most one chance to be added to the collection Γ : at the first step i so that $x_i \in \bar{\gamma}$. For $j < i$, $\gamma \notin \mathcal{C}_{\Gamma_{j-1}, S_{j-1}}(x_j)$ since $x_j \notin \bar{\gamma}$. For $j > i$, $\gamma \notin \mathcal{C}_{\Gamma_{j-1}, S_{j-1}}(x_j)$ since $S_{j-1} \cap \bar{\gamma} \neq \emptyset$. With this in mind, given a collection $\Gamma = \{\gamma_1, \dots, \gamma_k\} \in \mathcal{G}(G)$, let $i(j) = \min\{i : x_i \in \bar{\gamma}_j\}$.

Without loss of generality we may assume the $i(j)$'s are strictly increasing. Set $\Gamma_j = (\gamma_1, \dots, \gamma_j)$ and $X_j = \{x_1, \dots, x_j\}$. Using the convention that $i(0) = 0$, the probability that Γ is returned by the sampling algorithm is

$$\begin{aligned} \mu_{\text{alg}}(\Gamma) &= \prod_{j=1}^k \left(\frac{w(\gamma_j, z) Z(\mathcal{C}_{\Gamma_j, X_{i(j)}}, z)}{Z(\mathcal{C}_{\Gamma_{j-1}, X_{i(j)-1}}, z)} \cdot \prod_{i=i(j-1)+1}^{i(j)-1} \frac{Z(\mathcal{C}_{\Gamma_{j-1}, X_i}, z)}{Z(\mathcal{C}_{\Gamma_{j-1}, X_{i-1}}, z)} \right) \\ &\quad \times \prod_{i=i(k)+1}^n \frac{Z(\mathcal{C}_{\Gamma_k, X_i}, z)}{Z(\mathcal{C}_{\Gamma_k, X_{i-1}}, z)} \\ &= \prod_{j=1}^k \frac{w(\gamma_j, z) Z(\mathcal{C}_{\Gamma_j, X_{i(j)}}, z)}{Z(\mathcal{C}_{\Gamma_{j-1}, X_{i(j)-1}}, z)} \cdot \frac{Z(\mathcal{C}_{\Gamma_k, X_n}, z)}{Z(\mathcal{C}_{\Gamma_k, X_{i(k)}}, z)} \\ &= \frac{\prod_{j=1}^k w(\gamma_j, z)}{Z(G, z)} \end{aligned}$$

which is $\mu_G(\Gamma)$, as desired. In the third equality we have used the fact that $Z(\mathcal{C}_{\Gamma_k, X_n}, z) = 1$ and $Z(\mathcal{C}_{\Gamma_0, X_{i(0)}}, z) = Z(G, z)$. \square \square

To turn Algorithm 1 into an efficient approximate sampling algorithm, we will sample approximately from the measures $\mu_{\Gamma_t, S_t, x_{t+1}}$. To analyze the effect on the output distribution we need a lemma about total variation distance.

Given a family of probability measures $\{\mu_\alpha\}_{\alpha \in \mathcal{A}}$, a μ -sequence of length n is a sequence of random variables $(X_i)_{i=1}^n$, where the conditional distribution of X_i is μ_{α_i} for some $\alpha_i \in \mathcal{A}$ that is a function of the values of the random variables X_j , $j < i$.

Lemma 5.3. *Let $(\mu_\alpha)_{\alpha \in \mathcal{A}}$ and $(\nu_\alpha)_{\alpha \in \mathcal{A}}$ be families of probability measures on a finite set, and suppose $\|\mu_\alpha - \nu_\alpha\|_{TV} < \epsilon'$ for all $\alpha \in \mathcal{A}$. Then if $\epsilon' < \epsilon^2/(9n^2)$ the total variation distance between the distributions of μ - and ν -sequences of length n is at most ϵ .*

Proof. The hypothesis $\|\mu_\alpha - \nu_\alpha\|_{TV} < \epsilon'$ implies the subset $A(\alpha)$ of outcomes such that

$$(1 - \sqrt{\epsilon'})\mu_\alpha(a) \leq \nu_\alpha(a) \leq (1 + \sqrt{\epsilon'})\mu_\alpha(a), \quad a \in A(\alpha),$$

has measure $\mu_\alpha(A(\alpha)) \geq 1 - 2\sqrt{\epsilon'}$ for all $\alpha \in \mathcal{A}$.

Let $(X_i)_{i=1}^n$ and $(Y_i)_{i=1}^n$ be μ - and ν -sequences of length n , respectively. Write μ for the law of the μ -sequence and similarly for ν . Let A be the event that for each $1 \leq j \leq n$ both $X_i \sim \mu_{\alpha_i}$ and $Y_i \sim \nu_{\alpha_i}$ take values in $A(\alpha)$. By a union bound A has μ -measure at least $1 - 2n\sqrt{\epsilon'}$. Moreover,

$$(20) \quad \nu(\bar{a}) \geq (1 - n\sqrt{\epsilon'})\mu(\bar{a})$$

for any $\bar{a} = (a_1, \dots, a_n) \in A$. Recalling the definition of ϵ' , the claim now follows as

$$\|\mu - \nu\|_{TV} = \sum_{\mu(\bar{a}) > \nu(\bar{a})} \mu(\bar{a}) - \nu(\bar{a}) \leq \mu(A^c) + \sum_{\mu(\bar{a}) > \nu(\bar{a})} \mu(\bar{a})n\sqrt{\epsilon'} < 3n\sqrt{\epsilon'},$$

where we have used the estimate (20) to obtain the inequality by splitting the sum into those $\bar{a} \in A$ and those not. \square \square

We need a lemma that tells us we only need to consider polymers of size at most $O(\log(n/\epsilon))$. Recall $\mathcal{C}_m(G) \subset \mathcal{C}(G)$ is the set of all polymers of size at most m , and let $\mathcal{G}_m(G)$ be the collection of all sets of compatible polymers from $\mathcal{C}_m(G)$. Let $Z_m(G, z)$ denote $Z(\mathcal{C}_m(G), z)$, where this partition function is defined according to Section 2.6. Let $\mu_{G,m}$ be the corresponding probability measure. We consider $\mu_{G,m}$ as a measure on $\mathcal{G}(G)$ by setting $\mu_{G,m}(\Gamma) = 0$ for any collection $\Gamma \in \mathcal{G}(G)$ that contains a contour of size larger than m .

Lemma 5.4. *Suppose the polymer model satisfies Assumption 1 with constant ρ , $Z(G, z)$ is a polynomial of degree at most $C|G|$ for all $G \in \mathfrak{G}$, and that the cluster expansion for $\log Z(G, z)$ converges absolutely for all $G \in \mathfrak{G}$ and all $|z| < \delta$. Let*

$$m = \lceil \frac{\log(2C|G|/\epsilon)}{\rho(1 - |z|/\delta)} \rceil.$$

Then

$$\|\mu_G - \mu_{G,m}\|_{TV} < e^{2\epsilon} - 1.$$

Proof. For $\Gamma \in \mathcal{G}_m(G)$ we have

$$\mu_{G,m}(\Gamma) = \mu_G(\Gamma) \frac{Z(G, z)}{Z_m(G, z)}.$$

By Lemma 2.1 and the remark below Lemma 2.8 we have

$$e^{-\epsilon} Z(G, z) \leq \exp[T_m(G, z)] \leq e^\epsilon Z_m(G, z)$$

as the degree of both $Z(G, z)$ and $Z_m(G, z)$ is at most $C|G|$. Thus $\|\mu_G - \mu_{G,m}\|_{TV}$ can be estimated by

$$\sum_{\Gamma: \mu_{G,m}(\Gamma) > \mu_G(\Gamma)} |\mu_{G,m}(\Gamma) - \mu_G(\Gamma)| \leq \sum_{\Gamma} \left| \mu_G(\Gamma) \left(\frac{Z(\Lambda, z)}{Z_m(\Lambda, z)} - 1 \right) \right| \leq e^{2\epsilon} - 1.$$

□ □

Extend the notation given in (19) to polymers of restricted sizes by setting $\mathcal{C}_{\Gamma,S}^m = \mathcal{C}_{\Gamma,S} \cap \mathcal{C}_m$ and $\mathcal{C}_{\Gamma,S}^m(x) = \mathcal{C}_{\Gamma,S}(x) \cap \mathcal{C}_m$. Theorem 5.1 relies on the following algorithm.

Algorithm 2. *Let ϵ' and m be given. Set $\Gamma_0 = \emptyset$, $S_0 = \emptyset$, and order the vertices of G by x_1, \dots, x_n . Repeat the following procedure for $t = 0$ to $n - 1$:*

- (1) *Create the list of polymers $\mathcal{C}_{\Gamma_t, S_t}^m(x_{t+1})$.*
- (2) *For each $\gamma \in \mathcal{C}_{\Gamma_t, S_t}^m(x_{t+1})$, compute $Y(\gamma)$, an ϵ' -relative approximation to $w(\gamma, z)Z(\mathcal{C}_{\Gamma_t \cup \gamma, S_t \cup x_{t+1}}, z)$. Do the same for the empty polymer γ_\emptyset .*
- (3) *Sample γ from the measure $\hat{\mu}_{\Gamma_t, S_t, x_{t+1}}$ defined by*

$$\hat{\mu}_{\Gamma_t, S_t, x_{t+1}}(\gamma) = \frac{Y(\gamma)}{Y(\gamma_\emptyset) + \sum_{\gamma \in \mathcal{C}_{\Gamma_t, S_t}^m(x_{t+1})} Y(\gamma)}.$$

- (4) *Set $\Gamma_{t+1} = \Gamma_t \cup \gamma$.*
- (5) *Set $S_{t+1} = S_t \cup x_{t+1}$.*

Return $\Gamma = \Gamma_n$.

Proof of Theorem 5.1. Let $n = |G|$, and let m be as in the statement of Lemma 5.4. By Lemma 5.4 it is enough to show that Algorithm 2 produces an ϵ -approximate sample from $\mu_{G,m}$ in time polynomial in n and $1/\epsilon$. By Lemma 5.3, Algorithm 2 will output an ϵ -approximate sample from $\mu_{G,m}$ if each approximation in step (2) is an $\epsilon' = O(\epsilon^2/n^2)$ -relative approximation.

Since there are only n steps in Algorithm 2, what remains is to show that each step of the algorithm takes time polynomial in n and $1/\epsilon$. The creation of the list in step (1) can be done in polynomial time by Lemma 2.4: first we list all polymers in \mathcal{C}_m in polynomial time. We can then determine which polymers are in $\mathcal{C}_{\Gamma_t, S_t}(x_{t+1})$ by checking, for each $\gamma' \in \mathcal{C}_m$, (i) if $x_{t+1} \in \gamma'$ and (ii) if there is any $s \in S$ so that $s \in \gamma'$ or any $v \in \bigcup_{\gamma \in \Gamma} \gamma$ so that $d(v, \gamma') \leq 1$. Since there are at most n vertices to check, this last step takes time at most $O(n^2)$.

Computing the approximations in step (2) can be done in polynomial time, as (i) $\mathcal{C}_{\Gamma_t, S_t}^m(x_{t+1})$ has size at most polynomial in n and $1/\epsilon$ by the definition of m , (ii) we can obtain ϵ' -relative approximations to the partition functions in polynomial time by Lemma 2.8, and (iii) we can obtain ϵ' -relative approximations to the weight functions $w(\gamma, z)$ by Assumption 5. □ □

5.2. Sampling from a contour model. Sampling from a contour model is almost the same as sampling from a polymer model, but we must be precise about which probability measure we sample from and the notions of incompatibility used.

For $z > 0$ define the probability measure μ_Λ^φ on $\mathcal{G}_{\text{ext}}^\varphi(\Lambda)$ by

$$(21) \quad \mu_\Lambda^\varphi(\Gamma) := \frac{\prod_{\gamma \in \Gamma} \left(z^{\|\gamma\|} \prod_{\varphi' \in \Xi} Z^{\varphi'}(\text{int}_{\varphi'} \gamma, z) \right)}{Z^\varphi(\Lambda, z)}.$$

Theorem 5.5. *Fix $d \geq 2$ and suppose the conditions of Lemma 3.4 hold for $\varphi \in \Xi$. Then for any $0 < z < \delta$ there is an efficient sampling algorithm for the measure μ_Λ^φ given in (21) for any region $\Lambda \subset \mathbb{Z}^d$.*

The algorithm we use to prove Theorem 5.5 will be a version of Algorithm 2 suited to contour models. The following definitions are analogues of those in Section 5.1. The main difference is that compatibility of polymers now becomes compatibility and mutual externality of contours of the same type, and so instead of attempting to add a polymer such that $x \in \bar{\gamma}$, we attempt to add a contour such that $x \in \text{cov}(\gamma)$, where we recall that $\text{cov}(\gamma)$ was defined in (9).

For $S \subset \Lambda$ and a collection of compatible external contours $\Gamma \in \mathcal{G}^\varphi(\Lambda)$, let

$$\mathcal{C}_{\Gamma, S}^\varphi = \mathcal{C}_{\Gamma, S}^\varphi(\Lambda) := \{\gamma \in \mathcal{C}^\varphi(\Lambda) : \text{cov}(\gamma) \cap S = \emptyset, \gamma \cup \Gamma \in \mathcal{G}^\varphi(\Lambda)\}.$$

For $x \in \Lambda$, set $\mathcal{C}_{\Gamma, S}^\varphi(x)$ denote the subset of contours in $\mathcal{C}_{\Gamma, S}^\varphi$ such that $x \in \text{cov}(\gamma)$. Note that if $\gamma, \gamma' \in \mathcal{C}_{\Gamma, S}^\varphi(x)$ then γ and γ' are not mutually external.

Let γ_\emptyset be the empty contour, and set $w^{\text{ext}}(\gamma_\emptyset, z) := 1$. Let $\mu_{\Gamma, S, x}^\varphi$ be the probability measure on $\mathcal{C}_{\Gamma, S}^\varphi(x) \cup \gamma_\emptyset$ defined by

$$\mu_{\Gamma, S, x}^\varphi(\gamma) = \frac{w^{\text{ext}}(\gamma, z) Z^\varphi(\mathcal{C}_{\Gamma \cup \gamma, S \cup x}^\varphi, z)}{Z^\varphi(\mathcal{C}_{\Gamma, S}^\varphi, z)},$$

where $Z^\varphi(\mathcal{C}_{\Gamma, S}^\varphi)$ is defined as in Section 3.4.

Algorithm 3. *Let $n = |\Lambda|$. Set $\Gamma_0 = \emptyset$, $S_0 = \emptyset$, and order the vertices of Λ by x_1, \dots, x_n .² Repeat the following procedure for $t = 1$ to n :*

- (1) Sample γ from the measure $\mu_{\Gamma_{t-1}, S_{t-1}, x_t}^\varphi$.
- (2) Set $\Gamma_t = \Gamma_{t-1} \cup \gamma$.
- (3) Set $S_t = S_{t-1} \cup x_t$.

Return $\Gamma = \Gamma_n$.

Lemma 5.6. *The output Γ of Algorithm 3 has distribution μ_Λ^φ .*

Proof. The proof is the same as that of Lemma 5.2. □

To turn Algorithm 3 into an efficient approximate sampling algorithm we follow the recipe used in obtaining Algorithm 2: we only consider contours of size $O(\log(n/\epsilon))$ and we approximate the weight functions and partition functions involved in the probability measures $\mu_{\Gamma_{t-1}, S_{t-1}, x_t}^\varphi$. The details follow.

We let $\mathcal{C}_m^\varphi(\Lambda) := \{\gamma \in \mathcal{C}^\varphi(\Lambda) : |\bar{\gamma}| \leq m\}$, $\mathcal{G}_{\text{ext}, m}^\varphi(\Lambda)$ be the collection of all sets of mutually external contours from $\mathcal{C}_m^\varphi(\Lambda)$, $\mathcal{C}_{\Gamma, S}^{\varphi, m} := \mathcal{C}_{\Gamma, S}^\varphi \cap \mathcal{C}_m^\varphi(\Lambda)$, and lastly

²We could consider only vertices x such that $d_\infty(x, \Lambda^c) > 1$, but it does no harm to include the others.

$\mathcal{C}_{\Gamma,S}^{\varphi,m}(x) := \mathcal{C}_{\Gamma,S}^{\varphi}(x) \cap \mathcal{C}_m^{\varphi}(\Lambda)$. With these definitions, we can present the approximate sampling algorithm.

Algorithm 4. Let ϵ' and m be given and let $n = |\Lambda|$. Set $\Gamma_0 = \emptyset$, $S_0 = \emptyset$, and order the vertices of Λ by x_1, \dots, x_n . Repeat the following procedure for $t = 0$ to $n - 1$:

- (1) Create the list of contours $\mathcal{C}_{\Gamma_t, S_t}^m(x_{t+1})$.
- (2) For each $\gamma \in \mathcal{C}_{\Gamma_t, S_t}^m(x_{t+1})$, compute $Y(\gamma)$, an ϵ' -relative approximation to $w^{\text{ext}}(\gamma, z)Z^{\varphi}(\mathcal{C}_{\Gamma_t \cup \gamma, S_t \cup x_{t+1}}^{\varphi}, z)$. Do the same for the empty contour γ_{\emptyset} .
- (3) Sample γ from the measure $\hat{\mu}_{\Gamma_t, S_t, x_{t+1}}$ defined by

$$\hat{\mu}_{\Gamma_t, S_t, x_{t+1}}(\gamma) = \frac{Y(\gamma)}{Y(\gamma_{\emptyset}) + \sum_{\gamma \in \mathcal{C}_{\Gamma_t, S_t}^m(x_{t+1})} Y(\gamma)}.$$

- (4) Set $\Gamma_{t+1} = \Gamma_t \cup \gamma$.
- (5) Set $S_{t+1} = S_t \cup x_{t+1}$.

Return $\Gamma = \Gamma_n$.

We now sketch the proof of Theorem 5.5; it is essentially that of Theorem 5.1.

Proof of Theorem 5.5. By Lemma 5.4, it suffices to sample a configuration of outer contours from $\mathcal{G}_{\text{ext},m}^{\varphi}(\Lambda)$ with $m = O(\log(n/\epsilon))$. We then implement Algorithm 4 with $\epsilon' = O(\epsilon^2/n^2)$, where the ϵ' -relative approximations $Y(\gamma)$ can be computed in time polynomial in n and $1/\epsilon$ by Lemma 3.4. Finally we use Lemma 5.3 to say that the output of the approximate algorithm is a close approximation to the truncated contour probability measure. \square \square

5.3. Applications of Theorem 5.5. The algorithm of Theorem 5.5 returns a collection of contours Γ approximately distributed according to the outer contour measure μ_{Λ}^{φ} . If the contour model arises from a spin system such as the Potts model or hard-core model it is straightforward to recover a spin configuration from inductive calls to this algorithm. We show how to partially determine a configuration $\omega \in \Omega_{\Lambda}^{\varphi}$ given a set of outer contours $\Gamma \in \mathcal{G}^{\varphi}(\Lambda)$.

For each $\gamma \in \Gamma, i \in \bar{\gamma}$, set ω_i to the spin indicated by $\omega_{\bar{\gamma}}$. For each $i \in \Lambda$ so that $i \in \bigcap_{\gamma \in \Gamma} \text{ext} \gamma$, set ω_i to the spin indicated by the ground state φ (e.g., for Potts set $\omega_i = \varphi$, and for hard-core set $\omega_i = \omega_i^{\varphi}$, where we recall ω_i^{even} is the all-even occupied configuration). This leaves ω_i so that $i \in \text{int} \gamma, \gamma \in \Gamma$ unset. To determine these spins, call the algorithm again for $\text{int}'_{\varphi} \gamma$ for each $\gamma \in \Gamma$ and each $\varphi' \in \Xi$.

Using the correspondence between spin configurations and contour configurations given in Examples 5 and 6, this proves the sampling portions of Theorems 1.2 and 1.4.

6. THE TORUS

In this section we give counting and sampling algorithms for contour models on the torus $\mathbb{T}_n^d = \mathbb{Z}^d/(n\mathbb{Z})^d$. We first explain how contour models are defined in this context, and in the subsequent sections we indicate how our previous algorithms can be extended to this setting. For the rest of this section we will fix $d \geq 2$ and write \mathbb{T}_n for \mathbb{T}_n^d to simplify the notation.

6.1. Contour models on \mathbb{T}_n . Contour models on \mathbb{T}_n are defined almost exactly as for regions $\Lambda \subset \mathbb{Z}^d$, but some additional care is needed as the change in topology affects the notion of the exterior of a contour. In the approach below we will largely circumvent topological complications by distinguishing contours that are ‘large’, i.e., those that can detect the change in topology. Large contours make negligible contributions in the cases we are interested in.

A contour γ on the torus \mathbb{T}_n is a pair $(\bar{\gamma}, \omega_{\bar{\gamma}})$ consisting of a subset of vertices $\bar{\gamma} \subset \mathbb{T}_n$ and an assignment $\omega_{\bar{\gamma}}$ of spins from Ω to $\bar{\gamma}$. Letting x_i denote the i th coordinate of $x \in \mathbb{T}_n$, the diameter of a set $A \subset \mathbb{T}_n$ is

$$\text{diam}(A) := \max_{x, y \in A} d_{\infty}(x, y).$$

Following [17], we distinguish between two types of contours, those that are ‘small’ and those that are ‘large’. A *small contour* is a contour γ for which $\text{diam}(\bar{\gamma}) < n/2$ and $\bar{\gamma}$ is d_{∞} -connected. A *large contour* is a contour γ for which each d_{∞} -connected component of $\bar{\gamma}$ has diameter at least $n/2$; note that the support of a large contour need not be connected. Each contour γ partitions $\mathbb{T}_n \setminus \bar{\gamma}$ into d_{∞} -connected components A_0, \dots, A_t . Since each small contour is a subset of a d_{∞} -ball of radius less than $n/2$, we can define the *exterior of a small contour* γ to be the unique region with diameter at least $n/2$, and without loss of generality we can denote this region by A_0 . The regions A_i , $i \geq 1$, are *interior regions*. For a large contour γ we set $A_0 = \emptyset$ and refer to all connected components of $\mathbb{T}_n \setminus \bar{\gamma}$ as interior regions.

A *contour model* on \mathbb{T}_n consists of a set of contours \mathcal{C} , a surface energy $\|\gamma\|$, and labelling function $\text{lab}_{\gamma}(\cdot)$ taking values in Ξ for each contour $\gamma \in \mathcal{C}$. The label of the exterior A_0 of a small contour is called the *type* of the contour. Large contours γ have no exteriors and hence no type, but are still equipped with a surface energy $\|\gamma\|$ and a labelling function $\text{lab}_{\gamma}(\cdot)$ from the connected components of $\mathbb{T}_n \setminus \bar{\gamma}$ to Ξ .

Two contours γ, γ' are *compatible* if $d_{\infty}(\bar{\gamma}, \bar{\gamma}') > 1$, and two compatible small contours γ, γ' of the same type are *mutually external* if $\bar{\gamma} \subset \text{ext} \gamma'$ and $\bar{\gamma}' \subset \text{ext} \gamma$.

Let $\mathcal{C}^{\varphi}(\mathbb{T}_n)$ be the set of all small contours of type φ . Let $\mathcal{G}_{\text{ext}}^{\varphi}(\mathbb{T}_n)$ be the collection of all sets of contours from $\mathcal{C}^{\varphi}(\mathbb{T}_n)$ that are compatible and mutually external. Let $\mathcal{C}_{\text{large}}(\mathbb{T}_n)$ be the set of all large contours.

For $\mathcal{S} \subseteq \mathcal{C}^{\varphi}(\mathbb{T}_n)$, let $\mathcal{G}_{\text{ext}}^{\varphi}(\mathcal{S})$ be the collection of all sets of compatible and mutually external contours of type φ from \mathcal{S} , and define

$$Z^{\varphi}(\mathcal{S}, z) := \sum_{\Gamma \in \mathcal{G}_{\text{ext}}^{\varphi}(\mathcal{S})} \prod_{\gamma \in \Gamma} \left(z^{\|\gamma\|} \prod_{\varphi'} Z^{\varphi'}(\text{int}_{\varphi'} \gamma, z) \right).$$

We now define the *partition function of a contour model on \mathbb{T}_n* by

$$Z(\mathbb{T}_n, z) := Z^{\text{big}}(\mathbb{T}_n, z) + \sum_{\varphi \in \Xi} Z^{\varphi}(\mathbb{T}_n, z),$$

where

$$Z^{\text{big}}(\mathbb{T}_n, z) := \sum_{\gamma \in \mathcal{C}_{\text{large}}(\mathbb{T}_n)} z^{\|\gamma\|} \prod_{\varphi \in \Xi} Z^{\varphi}(\text{int}_{\varphi} \gamma, z),$$

and where $Z^\varphi(\mathbb{T}_n, z)$ is shorthand for $Z^\varphi(\mathcal{S}, z)$ with $\mathcal{S} = \mathcal{C}^\varphi(\mathbb{T}_n)$. Note that each contour configuration contributing to $Z(\mathbb{T}_n, z)$ contains at most one large contour.

We can also write the partition function in an expanded form involving a matching condition. A small contour γ in a collection Γ of compatible small contours is *external* if $\bar{\gamma} \subset \text{ext}\gamma'$ for all $\gamma' \in \Gamma$, $\gamma' \neq \gamma$. As in Section 3, a set Γ of compatible small contours is matching if (i) all external contours have the same type and (ii) for each external contour γ and ground state φ the subcollection of contours Γ' whose support is contained in $\text{int}_\varphi\gamma$ is matching and of type φ . A set of compatible contours Γ containing exactly one large contour γ is matching if for each ground state $\varphi \in \Xi$ the subcollection of contours Γ' whose support is contained in $\text{int}_\varphi\gamma$ is matching and of type φ . Let $\mathcal{G}_{\text{match}}(\mathbb{T}_n)$ be the collection of all sets of matching contours. Then

$$(22) \quad Z(\mathbb{T}_n, z) = (|\Xi| - 1) + \sum_{\Gamma \in \mathcal{G}_{\text{match}}(\mathbb{T}_n)} \prod_{\gamma \in \Gamma} z^{\|\gamma\|}.$$

The term $(|\Xi| - 1)$ is due to the fact that for each $\varphi \in \Xi$ there is a contribution of 1 to $Z^\varphi(\mathbb{T}_n, z)$ from the empty collection of contours.

Moreover, for each φ and $\mathcal{S} \subset \mathcal{C}$, let $\mathcal{G}_{\text{match}}^\varphi(\mathcal{S})$ be the collection of sets of matching small contours from \mathcal{S} whose external contours are all of type φ . Then

$$Z^\varphi(\mathcal{S}, z) = \sum_{\Gamma \in \mathcal{G}_{\text{match}}^\varphi(\mathcal{S})} \prod_{\gamma \in \Gamma} z^{\|\gamma\|}.$$

and again we let $Z^\varphi(\mathbb{T}_n, z) = Z^\varphi(\mathcal{S}, z)$ with $\mathcal{S} = \mathcal{C}(\mathbb{T}_n)$.

Borgs and Imbrie show that under the Peierls condition, for small enough z the relative weight of $Z^{\text{big}}(\mathbb{T}_n, z)$ in $Z(\mathbb{T}_n, z)$ is exponentially small in n . More precisely, and noting that the definition of stable ground states Ξ^{stab} from Section 4 applies equally well to the partition functions of small contours on \mathbb{T}_n , they prove:

Theorem 6.1 (Borgs, Imbrie [17]). *Suppose the contour model satisfies Assumption 3 for some $\rho, C > 0$. Then there exists a constant $\delta = \delta(d, \rho, \Xi) > 0$ and constants $N, c' > 0$ so that for $n > N$, and real $0 < z < \delta$,*

$$\frac{|Z(\mathbb{T}_n, z) - \sum_{\varphi \in \Xi^{\text{stab}}} Z^\varphi(\mathbb{T}_n, z)|}{|Z(\mathbb{T}_n, z)|} \leq e^{-c'n}.$$

Moreover, for all complex $|z| < \delta$ and all $\varphi \in \Xi^{\text{stab}}$, $Z^\varphi(\mathbb{T}_n, z) \neq 0$.

Using this result, we prove our main counting result for the torus. In Section 6.2 below we apply the result to prove Theorems 1.3 and 1.5.

Theorem 6.2. *Fix $d \geq 2$, suppose the contour model satisfies Assumptions 2 and 3 and that all ground states $\varphi \in \Xi$ are stable. Then there exists a constant $\delta = \delta(d, \rho, \Xi) > 0$ and a constant $c = c(d, \rho, \Xi) > 0$ so that for all real $0 < z < \delta$ and all $\epsilon \geq e^{-cn}$, there is an algorithm to obtain an ϵ -relative approximation to $Z(\mathbb{T}_n, z)$ in time polynomial in n and $1/\epsilon$.*

The conclusion of Theorem 6.2 is slightly weaker than that of Theorem 3.1, e.g., it does not allow ϵ to be exponentially small in n^{d-1} . See Section 7 for comments on obtaining a full FPTAS.

Note that we require z to be positive in Theorems 6.1 and 6.2. This is because for complex or negative z there could be cancellations in the sum of partition functions associated to the stable ground states. For models with a symmetric set of ground states, like the Potts and hard-core models, we can take $|z| < \delta$ complex in both theorems as these cancellations cannot occur.

Proof of Theorem 6.2. Let $c = c'/2$ where c' is the constant from Theorem 6.1, and choose $\epsilon = \epsilon(n) \geq e^{-cn}$. By Theorem 6.1, for $0 < z < \delta$ and n large enough we know $\sum_{\varphi \in \Xi} Z^\varphi(\mathbb{T}_n, z)$ is an $\epsilon/2$ -relative approximation to $Z(\mathbb{T}_n, z)$. Hence it suffices to compute $\epsilon/2$ -relative approximations to $Z^\varphi(\mathbb{T}_n, z)$ for each $\varphi \in \Xi$.

We can compute an $\epsilon/2$ approximation to $Z^\varphi(\mathbb{T}_n, z)$ almost exactly as in the proof of Theorem 3.1. Lemma 2.4 applies to $\Lambda = \mathbb{T}_n$ as \mathbb{T}_n is a graph of bounded degree. The proof of Lemma 3.3 carries through as before; we can still order small contours so that γ precedes γ' if γ can appear in the interior of γ' . Moreover, we can inductively compute the weights exactly as before, since $\bar{\gamma}$ has diameter $< n/2$ and so can be embedded in \mathbb{Z}^d . \square \square

6.2. Sampling on the torus. Define the following probability measure associated to the matching contour representation [22]

$$\mu_{\mathbb{T}_n}^{\text{match}}(\Gamma) := \frac{\prod_{\gamma \in \Gamma} z^{\|\gamma\|}}{Z(\mathbb{T}_n, z)}, \quad \Gamma \in \mathcal{G}_{\text{match}}(\mathbb{T}_n).$$

Under the conditions of Theorem 6.2 we obtain an efficient approximate sampling algorithm for $\mu_{\mathbb{T}_n}^{\text{match}}$.

Theorem 6.3. *Fix $d \geq 2$, suppose the contour model satisfies Assumptions 2 and 3 and that all ground states $\varphi \in \Xi$ are stable. Then there exists a constant $\delta = \delta(d, \rho, \Xi) > 0$ and a constant $c = c(d, \rho, \Xi) > 0$ so that for all real $0 < z < \delta$ and all $\epsilon \geq e^{-cn}$, there is an ϵ -approximate sampling algorithm for $\mu_{\mathbb{T}_n}^{\text{match}}$ that runs in time polynomial in n and $1/\epsilon$.*

To prove Theorem 6.3 we need some auxiliary probability measures. The measure $\mu_{\mathbb{T}_n}^{\text{match}}$ conditioned on $\Gamma \in \mathcal{G}_{\text{match}}^\varphi(\mathbb{T}_n)$ is

$$\mu_{\mathbb{T}_n}^{\varphi, \text{match}}(\Gamma) := \frac{\prod_{\gamma \in \Gamma} z^{\|\gamma\|}}{Z^\varphi(\mathbb{T}_n, z)}, \quad \Gamma \in \mathcal{G}_{\text{match}}^\varphi(\mathbb{T}_n).$$

We define the probability measure associated to the outer contour representation of $Z^\varphi(\mathbb{T}_n, z)$ as

$$\mu_{\mathbb{T}_n}^{\varphi, \text{ext}}(\Gamma) := \frac{\prod_{\gamma \in \Gamma} \left(z^{\|\gamma\|} \prod_{\varphi'} Z^{\varphi'}(\text{int}_{\varphi'} \gamma, z) \right)}{Z^\varphi(\mathbb{T}_n, z)}, \quad \Gamma \in \mathcal{G}_{\text{ext}}^\varphi(\mathbb{T}_n).$$

Lemma 6.4. *Under the assumptions of Theorem 6.3, for any $0 < z < \delta$, there is an efficient sampling algorithm for $\mu_{\mathbb{T}_n}^{\varphi, \text{ext}}$.*

The algorithm and proof of Lemma 6.4 are exactly the same as for Theorem 5.5. We now prove Theorem 6.3 using Theorem 6.1, Theorem 6.2, and Lemma 6.4.

Proof of Theorem 6.3. With $c = c'/2$, where c' is the constant from Theorem 6.1 and $\epsilon \geq e^{-cn}$, to obtain an ϵ -approximate sample from $\mu_{\mathbb{T}_n}^{\text{match}}$ it suffices to select $\varphi \in \Xi$ with probability approximately proportional to $Z^\varphi(\mathbb{T}_n, z)$ and return an $\epsilon/2$ approximate sample from $\mu_{\mathbb{T}_n}^{\varphi, \text{match}}$. We can choose the ground state φ within total variation distance $\epsilon/2$ by approximating $Z^{\varphi'}(\mathbb{T}_n, z)$ within relative error $\epsilon/2$ for each φ' using Theorem 6.2.

To obtain an $\epsilon/2$ approximate sample from $\mu_{\mathbb{T}_n}^{\varphi, \text{match}}$, we obtain an $\epsilon/(2n)$ -approximate sample from $\mu_{\mathbb{T}_n}^{\varphi, \text{ext}}$ using Lemma 6.4 and then proceed inductively on the interior regions, as in Section 5.3. To sample approximately from $\mu_{\text{int}_{\varphi'} \gamma}^{\varphi'}$ we can use Theorem 5.5 as $\text{diam}(\text{int}_{\varphi'} \gamma) < n/2$ and so it can be embedded in \mathbb{Z}^d . We return the collection Γ of all contours sampled at each step which is by definition a set of matching contours. $\square \quad \square$

6.3. Applications. Theorem 6.3 immediately implies Theorems 1.3 and 1.5 by the same mapping of a set of matching contours to a spin configuration given in Section 5.3. Note that we must take n even in Theorem 1.5 so that we can properly define the contour models.

7. CONCLUSIONS

We conclude by describing some open problems.

7.1. Extending the region of applicability. It would be interesting to optimize the ranges of parameters for which our algorithms work. The proofs of Theorems 1.2 and 1.3 for the Potts model use techniques from mathematical physics [39, 18] that have also been used to prove slow mixing of the Swendsen–Wang dynamics at β_c when q is sufficiently large [15, 16]. For large q we therefore expect that we can take $\beta^*(d, q) = \beta_c(d, q)$. In fact, for large q the techniques of this paper yields an efficient counting algorithm for $\beta > \beta_c$ and quasi-polynomial-time counting and sampling algorithms for all $\beta \geq \beta_c$.

For the hard-core model, it is known that there is phase coexistence on \mathbb{Z}^d for $\lambda \geq C \frac{\log^2 d}{d^{1/3}}$ [26, 47]. It would be of interest to understand how small λ^* could be taken to obtain an efficient sampling algorithm for the hard-core model on \mathbb{Z}^d .

Open Problem. *Can Theorems 1.4 and 1.5 be extended to $\lambda^*(d) = \tilde{\Theta}(d^{-1/3})$?*

A related direction would be to use more geometrically sophisticated notions of contours to improve the range of parameters for which the condition (15) holds.

Open Problem. *Find an FPTAS and efficient sampling algorithm for the hard-core model on \mathbb{T}_n^2 for $\lambda > 5.3506$, the region of coexistence for the hard-core model on \mathbb{Z}^2 proved in [11].*

With more sophisticated contours, one could hope to find algorithms for models whose ground states consist of collections of configurations, e.g., the q -coloring model.

Open Problem. *Find an FPTAS and efficient sampling algorithm for proper q -colorings of \mathbb{T}_n^d when $d = d(q)$ is sufficiently large.*

7.2. An FPTAS for the torus. The obstacle to applying Theorem 6.2 to obtain a genuine FPTAS for the torus is that if $\epsilon = \exp(-\omega(n))$, then the bound of Theorem 6.1 on the contribution from large contours is not small enough to ignore. However, by using much more sophisticated topological tools, Borgs, Chayes, and Tetali [16] showed a bound of $\exp(-\Theta(n^{d-1}))$ for the contributions to the Potts model partition function due to configurations containing an ‘interface’ of non-zero winding number on the torus. This upper bound is matched by an upper bound of $\exp(\Theta(n^{d-1}))$ on the mixing time of the Glauber dynamics for the Potts model on the torus in the same paper.

Remarkably, these two ingredients together with the techniques of this paper can give a true FPTAS and efficient sampling algorithm on the torus. If $\epsilon = \exp(-o(n^{d-1}))$, then we safely ignore contributions to the partition function from configurations with interfaces and run our counting and sampling algorithms. But if $\epsilon = \exp(-\Omega(n^{d-1}))$ then the Glauber dynamics provide a sampling algorithm that runs in time polynomial in n and $1/\epsilon$. The idea is straightforward, but the topological details are rather complicated, and so we leave this for future work.

7.3. Markov chains. The algorithms we have presented run in time $(n/\epsilon)^{O(\log d)}$, which is polynomial in n and $1/\epsilon$ for fixed d but far from linear time. A more efficient approach would be to use a Markov chain. While the Glauber dynamics is known to mix slowly at low temperature in models of the type we consider here [15], the definition of mixing time is rather strict and slow mixing does not rule out an efficient sampling algorithm based on the Glauber dynamics.

For spin models with finitely many stable and symmetric ground states, like the Potts or hard-core models, we suggest a Markov chain algorithm to sample on the torus \mathbb{T}_n^d .

- (1) Pick a ground state $\varphi \in \Xi$ uniformly at random.
- (2) Run the Glauber dynamics with the ground state configuration corresponding to φ as the initial configuration (i.e. a monochromatic initial configuration for the Potts model; all even or all odd occupied for the hard-core model).

We conjecture that at sufficiently low temperatures (sufficiently high fugacities) in such models the distribution is close to stationary after $O(n \log n)$ steps of the Markov chain; we include the randomness from the choice of the ground state.

Open Problem. *Prove that the above algorithm is an efficient sampling algorithm for the Potts model below the critical temperature or the hard-core model at sufficiently high fugacity.*

For the 2-dimensional Ising model on a box with all plus boundary conditions, Glauber dynamics starting from the all plus configuration does in fact converge rapidly to the stationary distribution for $\beta > \beta_c$ [43].

7.4. Beyond \mathbb{Z}^d and beyond lattices. We have restricted ourselves to the lattices \mathbb{Z}^d for simplicity, and because some geometric lemmas about the connectivity of boundaries in \mathbb{Z}^d have been proved for us (e.g., [23, B.15] and [58]). Similar lemmas can presumably be proved for general lattices of dimension at least 2, but we leave this for future work. In particular, Theorem 1.4 can likely be extended to the entire class of non-sliding models considered by Jauslin and Lebowitz [33].

A related challenge would be to apply these methods to the *hard hexagon model* (i.e., the hard-core model on the triangular lattice) for which it is known that the free energy is analytic for all real non-critical fugacities [7, 37, 59].

Open Problem. *Find efficient counting and sampling algorithms for the hard hexagon model for real $\lambda \neq \lambda_c$.*

The fact that the underlying graph is a lattice does not seem to be entirely necessary. Given the interest in the complexity class $\#\text{BIS}$, it would be interesting to investigate contour representations of the hard-core model on more general families of bipartite graphs. See [41, 19, 25] for more about $\#\text{BIS}$. A cautionary note in this respect is that Bezáková, Galanis, Goldberg, and Štefankovič [9] have shown $\#\text{P}$ -hardness of approximating $Z_G(\lambda)$ on bipartite graphs for *any* complex λ with large real part.

7.5. Approximating the free energy. A computational problem related to the problems considered in this paper is to approximate the limiting free energy $f_d(\lambda) := \lim_{n \rightarrow \infty} \frac{1}{n} \log Z_{\mathbb{T}_n^d}(\lambda)$. The objective is an algorithm which, for any $\epsilon > 0$, outputs a number $\eta \in [f_d(\lambda) - \epsilon, f_d(\lambda) + \epsilon]$, and whose running time grows as slowly as possible as a function of $1/\epsilon$. Gamarnik and Katz [27] gave a polynomial time algorithm for the hard-core model for λ small enough that *strong spatial mixing* holds. This condition implies the hard-core model is in the uniqueness regime. Adams, Briceño, Marcus, and Pavlov [1] gave a polynomial-time algorithm for several models (including the hard-core model) on \mathbb{Z}^2 in a subset of the uniqueness regime. Their results also apply to the hard-core and Widom–Rowlinson models on \mathbb{Z}^2 in a subset of the non-uniqueness regime. This last result is of a similar spirit to the results of this paper, and it would be interesting to understand if our results have any bearing on this problem.

ACKNOWLEDGEMENTS

WP and GR thank Ivona Bezáková, Leslie Goldberg, and Mark Jerrum for organizing the 2017 Dagstuhl seminar on computational counting and Jan Hladký for organizing the 2018 workshop on graph limits in Bohemian Switzerland. Both meetings provided essential inspiration and discussion leading to this work. TH thanks Roman Kotecký for helpful discussions. We thank Eric Vigoda, Matthew Jenssen, and Reza Gheissari for detailed comments on a draft of the paper. TH supported by EPSRC grant EP/P003656/1. WP supported by in part by EPSRC grant EP/P009913/1 and NSF Career award DMS-1847451. GR supported by an NWO Veni grant.

REFERENCES

- [1] Adams, S., Briceño, R., Marcus, B., Pavlov, R.: Representation and poly-time approximation for pressure of \mathbb{Z}^2 lattice models in the non-uniqueness region. *Journal of Statistical Physics* **162**(4), 1031–1067 (2016)
- [2] Barvinok, A.: Computing the partition function for cliques in a graph. *Theory of Computing* **11**(13), 339–355 (2015)
- [3] Barvinok, A.: Computing the permanent of (some) complex matrices. *Foundations of Computational Mathematics* **16**(2), 329–342 (2016)
- [4] Barvinok, A.: Combinatorics and complexity of partition functions. *Algorithms and Combinatorics* **30** (2017)

- [5] Barvinok, A., Regts, G.: Weighted counting of solutions to sparse systems of equations. *Combinatorics, Probability and Computing*, to appear
- [6] Barvinok, A., Soberón, P.: Computing the partition function for graph homomorphisms with multiplicities. *Journal of Combinatorial Theory, Series A* **137**, 1–26 (2016)
- [7] Baxter, R.J.: Hard hexagons: exact solution. *Journal of Physics A: Mathematical and General* **13**(3), L61 (1980)
- [8] Beyer, T., Hedetniemi, S.M.: Constant time generation of rooted trees. *SIAM Journal on Computing* **9**(4), 706–712 (1980)
- [9] Bezáková, I., Galanis, A., Goldberg, L.A., Štefankovič, D.: Inapproximability of the independent set polynomial in the complex plane. In: *Proceedings of the Fiftieth Annual ACM Symposium on Theory of Computing*, STOC 2018, pp. 1234–1240. ACM (2018)
- [10] Björklund, A., Husfeldt, T., Kaski, P., Koivisto, M.: Computing the Tutte polynomial in vertex-exponential time. In: *Proceedings of the Forty-ninth Annual Symposium on Foundations of Computer Science*, FOCS 2008, pp. 677–686. IEEE (2008)
- [11] Blanca, A., Chen, Y., Galvin, D., Randall, D., Tetali, P.: Phase coexistence for the hard-core model on \mathbb{Z}^2 . *Combinatorics, Probability and Computing* **28**(1), 1–22 (2019)
- [12] Blanca, A., Sinclair, A.: Random-cluster dynamics in \mathbb{Z}^2 . *Probability Theory and Related Fields* **168**(3-4), 821–847 (2017)
- [13] Borgs, C., Chayes, J., Kahn, J., Lovász, L.: Left and right convergence of graphs with bounded degree. *Random Structures & Algorithms* **42**(1), 1–28 (2013)
- [14] Borgs, C., Chayes, J.T., Dyer, M., Tetali, P.: On the sampling problem for H-colorings on the hypercubic lattice. In: *Graphs, Morphisms, and Statistical Physics: DIMACS Workshop Graphs, Morphisms and Statistical Physics*, March 19-21, 2001, DIMACS Center, vol. 63, p. 13. American Mathematical Soc. (2004)
- [15] Borgs, C., Chayes, J.T., Frieze, A., Kim, J.H., Tetali, P., Vigoda, E., Vu, V.H.: Torpid mixing of some Monte Carlo Markov chain algorithms in statistical physics. In: *Proceedings of the Fourtieth Annual Symposium on Foundations of Computer Science*, FOCS 1999, pp. 218–229. IEEE (1999)
- [16] Borgs, C., Chayes, J.T., Tetali, P.: Tight bounds for mixing of the Swendsen–Wang algorithm at the Potts transition point. *Probability Theory and Related Fields* **152**(3-4), 509–557 (2012)
- [17] Borgs, C., Imbrie, J.Z.: A unified approach to phase diagrams in field theory and statistical mechanics. *Communications in mathematical physics* **123**(2), 305–328 (1989)
- [18] Borgs, C., Koteky, R., Miracle-Solé, S.: Finite-size scaling for Potts models. *Journal of Statistical Physics* **62**(3-4), 529–551 (1991)
- [19] Cai, J.Y., Galanis, A., Goldberg, L.A., Guo, H., Jerrum, M., Štefankovič, D., Vigoda, E.: # BIS-hardness for 2-spin systems on bipartite bounded degree graphs in the tree non-uniqueness region. *Journal of Computer and System Sciences* **82**(5), 690–711 (2016)
- [20] Dobrushin, R.: Estimates of semi-invariants for the Ising model at low temperatures. *Translations of the American Mathematical Society-Series 2* **177**, 59–82 (1996)
- [21] Dyer, M., Goldberg, L.A., Greenhill, C., Jerrum, M.: The relative complexity of approximate counting problems. *Algorithmica* **38**(3), 471–500 (2004)
- [22] Fernández, R., Procacci, A.: Cluster expansion for abstract polymer models. new bounds from an old approach. *Communications in Mathematical Physics* **274**(1), 123–140 (2007)
- [23] Friedli, S., Velenik, Y.: *Statistical mechanics of lattice systems: a concrete mathematical introduction*. Cambridge University Press (2017)
- [24] Galanis, A., Štefankovič, D., Vigoda, E.: Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *Combinatorics, Probability and Computing* **25**(4), 500–559 (2016)
- [25] Galanis, A., Stefankovic, D., Vigoda, E., Yang, L.: Ferromagnetic Potts model: Refined #BIS-hardness and related results. *SIAM Journal on Computing* **45**(6), 2004–2065 (2016)
- [26] Galvin, D., Kahn, J.: On phase transition in the hard-core model on \mathbb{Z}^d . *Combinatorics, Probability and Computing* **13**(2), 137–164 (2004)
- [27] Gamarnik, D., Katz, D.: Sequential cavity method for computing free energy and surface pressure. *Journal of Statistical Physics* **137**(2), 205–232 (2009)

- [28] Gheissari, R., Lubetzky, E.: Mixing times of critical two-dimensional Potts models. *Communications on Pure and Applied Mathematics* **71**(5), 994–1046 (2018)
- [29] Goldberg, L.A., Jerrum, M.: Approximating the partition function of the ferromagnetic Potts model. *Journal of the ACM* **59**(5), 25 (2012)
- [30] Grimmett, G.R.: The random-cluster model, second edn. Springer-Verlag (2006)
- [31] Guo, H., Jerrum, M.: Random cluster dynamics for the Ising model is rapidly mixing. *Ann. Appl. Probab.* **28**(2), 1292–1313 (2018). DOI 10.1214/17-AAP1335
- [32] Harvey, N.J., Srivastava, P., Vondrák, J.: Computing the independence polynomial: from the tree threshold down to the roots. In: *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1557–1576. SIAM (2018)
- [33] Jauslin, I., Lebowitz, J.L.: High-fugacity expansion and crystalline ordering for non-sliding hard-core lattice particle systems. *arXiv preprint arXiv:1708.01912* (2017)
- [34] Jenssen, M., Keevash, P., Perkins, W.: Algorithms for #BIS-hard problems on expander graphs. In: *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 2235–2247 (2019) Full version available at *arXiv:1807.04804*
- [35] Jerrum, M., Sinclair, A.: Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on computing* **22**(5), 1087–1116 (1993)
- [36] Jerrum, M.R., Valiant, L.G., Vazirani, V.V.: Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science* **43**, 169–188 (1986)
- [37] Joyce, G.: On the hard-hexagon model and the theory of modular functions. *Phil. Trans. R. Soc. Lond. A* **325**(1588), 643–702 (1988)
- [38] Kotecký, R., Preiss, D.: Cluster expansion for abstract polymer models. *Communications in Mathematical Physics* **103**(3), 491–498 (1986)
- [39] Laanait, L., Messager, A., Miracle-Solé, S., Ruiz, J., Shlosman, S.: Interfaces in the Potts model I: Pirogov-Sinai theory of the Fortuin-Kasteleyn representation. *Communications in Mathematical Physics* **140**(1), 81–91 (1991)
- [40] Lee, T.D., Yang, C.N.: Statistical theory of equations of state and phase transitions. II. Lattice gas and Ising model. *Physical Review* **87**(3), 410 (1952)
- [41] Liu, J., Lu, P.: FPTAS for #BIS with degree bounds on one side. In: *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing, STOC 2015*, pp. 549–556. ACM (2015)
- [42] Liu, J., Sinclair, A., Srivastava, P.: The Ising partition function: Zeros and deterministic approximation. *Journal of Statistical Physics* **174**, 287–315 (2019)
- [43] Lubetzky, E., Martinelli, F., Sly, A., Toninelli, F.L.: Quasi-polynomial mixing of the 2D stochastic Ising model with “plus” boundary up to criticality. *Journal of the European Mathematical Society* **15**, 339 (2013)
- [44] Otter, R.: The number of trees. *Annals of Mathematics* pp. 583–599 (1948)
- [45] Patel, V., Regts, G.: Computing the number of induced copies of a fixed graph in a bounded degree graph. *Algorithmica*, **81**(5), 1844–1858 (2019)
- [46] Patel, V., Regts, G.: Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials. *SIAM Journal on Computing* **46**(6), 1893–1919 (2017)
- [47] Peled, R., Samotij, W.: Odd cutsets and the hard-core model on \mathbb{Z}^d . *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques* **50**(3), 975–998 (2014)
- [48] Peters, H., Regts, G.: On a conjecture of Sokal concerning roots of the independence polynomial. *Michigan Math. J.* **68**(1), 33–55 (2019)
- [49] Pirogov, S.A., Sinai, Y.G.: Phase diagrams of classical lattice systems. *Theoretical and Mathematical Physics* **25**(3), 1185–1192 (1975)
- [50] Pirogov, S.A., Sinai, Y.G.: Phase diagrams of classical lattice systems continuation. *Theoretical and Mathematical Physics* **26**(1), 39–49 (1976)
- [51] Randall, D.: Slow mixing of Glauber dynamics via topological obstructions. In: *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006*, pp. 870–879. Society for Industrial and Applied Mathematics (2006)
- [52] Randall, D., Wilson, D.: Sampling spin configurations of an ising system. In: *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 1999*, pp. 959–960. Society for Industrial and Applied Mathematics (1999)

- [53] Scott, A.D., Sokal, A.D.: The repulsive lattice gas, the independent-set polynomial, and the Lovász local lemma. *Journal of Statistical Physics* **118**(5-6), 1151–1261 (2005)
- [54] Shearer, J.B.: On a problem of Spencer. *Combinatorica* **5**(3), 241–245 (1985)
- [55] Sinclair, A., Jerrum, M.: Approximate counting, uniform generation and rapidly mixing Markov chains. *Information and Computation* **82**(1), 93–133 (1989)
- [56] Sly, A.: Computational transition at the uniqueness threshold. In: *Proceedings of the Fifty-first Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pp. 287–296. IEEE (2010)
- [57] Sly, A., Sun, N.: Counting in two-spin models on d-regular graphs. *The Annals of Probability* **42**(6), 2383–2416 (2014)
- [58] Timár, Á.: Boundary-connectivity via graph theory. *Proceedings of the American Mathematical Society* **141**(2), 475–480 (2013)
- [59] Tracy, C.A., Grove, L., Newman, M.: Modular properties of the hard hexagon model. *Journal of Statistical Physics* **48**(3-4), 477–502 (1987)
- [60] Ullrich, M.: Rapid mixing of Swendsen–Wang dynamics in two dimensions. *Dissertationes Mathematicae* **502**, 1–64 (2014)
- [61] Weitz, D.: Counting independent sets up to the tree threshold. In: *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing, STOC 2006*, pp. 140–149. ACM (2006)
- [62] Zahradník, M.: An alternate version of Pirogov-Sinai theory. *Communications in Mathematical Physics* **93**(4), 559–581 (1984)

UNIVERSITY OF BRISTOL

E-mail address: th17948@bristol.ac.uk

UNIVERSITY OF ILLINOIS AT CHICAGO

E-mail address: math@willperkins.org

UNIVERSITY OF AMSTERDAM

E-mail address: guusregts@gmail.com