


ARTICLE

Algorithms for the ferromagnetic Potts model on expanders

Charlie Carlson¹, Ewan Davies², Nicolas Fraiman³, Alexandra Kolla⁴, Aditya Potukuchi⁵ and Corrine Yap⁶ 

¹Computer Science Department, University of California, Santa Barbara, Santa Barbara, CA, USA, ²Department of Computer Science, Colorado State University, Fort Collins, CO, USA, ³Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA, ⁴Computer Science and Engineering Department, University of California, Santa Cruz, Santa Cruz, CA, USA, ⁵Department of Electrical Engineering & Computer Science, York University, Toronto, ON, Canada, and ⁶School of Mathematics, Georgia Institute of Technology, Atlanta, GA, USA
Corresponding author: Corrine Yap; Email: math@corrineyap.com

(Received 10 October 2023; revised 16 February 2024; accepted 20 February 2024; first published online 5 April 2024)

Abstract

We give algorithms for approximating the partition function of the ferromagnetic q -color Potts model on graphs of maximum degree d . Our primary contribution is a fully polynomial-time approximation scheme for d -regular graphs with an expansion condition at low temperatures (that is, bounded away from the order-disorder threshold). The expansion condition is much weaker than in previous works; for example, the expansion exhibited by the hypercube suffices. The main improvements come from a significantly sharper analysis of standard polymer models; we use extremal graph theory and applications of Karger's algorithm to count cuts that may be of independent interest. It is #BIS-hard to approximate the partition function at low temperatures on bounded-degree graphs, so our algorithm can be seen as evidence that hard instances of #BIS are rare. We also obtain efficient algorithms in the Gibbs uniqueness region for bounded-degree graphs. While our high-temperature proof follows more standard polymer model analysis, our result holds in the largest-known range of parameters d and q .

Keywords: Potts model; approximation algorithm; expander graphs; cluster expansion

2020 MSC Codes: Primary: 82B20; Secondary: 68W25, 05C30

1. Introduction

The q -state Potts model on a graph $G = (V, E)$ at inverse temperature β is given by the partition function

$$Z_G(q, \beta) = \sum_{\sigma: V \rightarrow [q]} e^{\beta m(G, \sigma)}, \quad (1)$$

where the sum is over all assignments of the spins $[q] = \{1, 2, \dots, q\}$ to the vertices of G , and $m(G, \sigma)$ is the number of edges such that both endpoints receive the same spin under σ . The Potts model has been of continued interest in combinatorics and computer science, most notably because of its direct relations to notions such as q -cuts and graph colourings. The Potts model also arises in physics and other areas as a generalisation of the Ising model. The main computational question associated to the Potts model is to approximate the partition function. The model is a canonical example of a Markov random field, and hence an excellent testbed for algorithmic techniques.

The Potts model as above is *ferromagnetic* when $\beta > 0$ and *antiferromagnetic* when $\beta < 0$. We are interested in the ferromagnetic model because for $q \geq 3$, at low temperatures approximating the partition function is #BIS-hard [16, 19]. The complexity class known as #BIS consists of problems that are equivalent to counting the number of independent sets in bipartite graphs [11], and understanding the complexity of approximating problems in #BIS is a longstanding problem. More precisely, approximating the partition function of the ferromagnetic Potts model for fixed β is #BIS-hard [16], and restricted to bounded-degree graphs it is #BIS-hard at low temperatures [19]. That is, for any $d \geq 3$ and $q \geq 3$, for any $\beta > \beta_o(q, d)$ it is #BIS-hard to approximate $Z_G(q, \beta)$ on graphs of maximum degree d , where β_o is given by

$$\beta_o(q, d) = \ln \left(\frac{q-2}{(q-1)^{1-2/d} - 1} \right). \quad (2)$$

The parameter β_o has a precise definition as the *order-disorder* threshold of the Potts model on the random d -regular graph, and we note that it is not the same as the Gibbs uniqueness phase transition on the d -regular tree. See e.g. [19, 10] for a discussion of these thresholds. The Gibbs measure $\mu_{G,q,\beta}$ for the Potts model on the spin assignments of G is given by

$$\mu_{G,q,\beta}(\sigma) = e^{\beta m(G,\sigma)} / Z_G(q, \beta)$$

and, alongside approximating Z , algorithms that approximately sample from $\mu_{G,q,\beta}$ are of interest in theoretical computer science. By self-reducibility, the standard problems of approximating the partition function and approximately sampling from the Gibbs measure are equivalent.

Several recent works have given algorithms for #BIS-hard problems on random graphs or graphs with strong expansion [24, 21, 14, 9]. One consequence of such research is a restriction on possible hard instances of #BIS, since under somewhat general conditions these works give efficient algorithms. In the absence of strong evidence that #BIS admits a general polynomial-time approximation algorithm, a weaker but intriguing prospect is that of a subexponential-time algorithm. The recent algorithm of Jenssen *et al.* [26] counts independent sets in d -regular bipartite graphs in subexponential time when d grows with the number of vertices of the instance. Our work is motivated in part by the question of whether a subexponential-time approximation algorithm for the Potts partition function can be found (at temperatures known to be #BIS-hard).

In general, the interaction between probabilistic phase transitions or thresholds such as β_o and the computational complexity of approximating partition functions has received substantial attention. For example, while the (Gibbs uniqueness) phase transition on the infinite d -regular tree corresponds to an NP-hardness threshold in graphs of maximum degree d for antiferromagnetic 2-spin models such as the hard-core model and the Ising model [37, 39, 18], for the ferromagnetic Potts model on the random d -regular graph neither the Gibbs uniqueness phase transition nor the order-disorder threshold correspond to a computational threshold [21] when q is large enough in terms of d . The methods of [21] rely on large q , and our work addresses the question of how one can work with smaller q .

While our primary objective is to give algorithms for approximating $Z_G(q, \beta)$ in the #BIS-hardness region (i.e. low temperatures), we also investigate phase transitions subject to much weaker conditions than those satisfied by the random regular graph. To this end, we develop a high-temperature algorithm for general d -regular instances that works up to some $\beta_1(q, d)$ such that $\beta_1 \sim \beta_o$ as $d \rightarrow \infty$. Since the Gibbs uniqueness threshold is less than β_o , our work gives a polynomial-time approximation algorithm in the Gibbs uniqueness region for regular graphs and does so for the largest-known range of q .

1.1 Our results

Let $G = (V, E)$ be a d -regular graph on n vertices. For a set of vertices $A \subset V$, let the *edge-boundary* $\nabla(A)$ of A be the set of edges of G with precisely one endpoint in A . For $\eta > 0$, we say that a graph

G is an η -expander if

$$\text{every nonempty set } A \subset V \text{ with } |A| \leq n/2 \text{ satisfies } |\nabla(A)| \geq \eta|A|. \quad (3)$$

We note here that the d -dimensional hypercube Q_d is a 1-expander, and that η -expansion differs by a factor d from the usual definition of edge expansion ϕ in an n -vertex d -regular graph, $\phi(G) = \min_{0 < |A| \leq n/2} |\nabla(A)|/(d|A|)$.

For $z > 0$ and $\delta \in [0, 1)$, we say that \hat{z} is a δ -relative approximation to z if $1 - \delta \leq z/\hat{z} \leq 1 + \delta$. For a given $q \in \mathbb{N}$ and $\beta > 0$, a *fully polynomial-time approximation scheme* (FPTAS) for $Z_G(q, \beta)$ is an algorithm that for every $\delta > 0$ outputs a δ -relative approximation to $Z_G(q, \beta)$ and runs in time polynomial in n and $1/\delta$. A *polynomial-time sampling scheme* for $\mu_{G,q,\beta}$ outputs a random spin assignment with distribution $\hat{\mu}$ that lies within δ total variation distance of $\mu_{G,q,\beta}$ and runs in time polynomial in n and $1/\delta$. Our main theorem is the following.

Theorem 1. *For every $\epsilon > 0$, there exists $d_0(\epsilon)$ such that for $d \geq d_0(\epsilon)$, $q \geq d^c$ where c is an absolute constant, and positive $\beta \notin ((1 - \epsilon)\beta_o, (1 + \epsilon)\beta_o)$, there exist*

- (1) *an FPTAS for $Z_G(q, \beta)$, and*
- (2) *a polynomial-time sampling scheme for $\mu_{G,q,\beta}$*

for G in the class of d -regular 2-expander graphs, and for G in the class of triangle-free d -regular 1-expander graphs.

While the constants 2 and 1 marking the limit of our expansion range may seem arbitrary, they represent a subtle barrier in our proofs. All n -vertex d -regular graphs that are 2-expanders or triangle-free 1-expanders must have min-cut d and the property that a set A with $2 \leq |A| \leq n/2$ vertices on one side has edge-boundary at least $2d - 2$ (see (7)). We choose not to pursue any additional case analysis and parameter trade-offs that might overcome this barrier. One option, as in [21], is to require stronger (e.g. $\Omega(d)$) expansion for small sets but we avoided this assumption for Theorem 1. It is also worth noting that we require $d_0(\epsilon) \geq e^{\Omega(1/\epsilon)}$.

In the low-temperature regime, $\beta > (1 + \epsilon)\beta_o$, this is a significant improvement over a previous result of Jenssen et al. [24] which required that $q \geq d^{\Omega(d)}$, as well as the graph being an $\Omega(d)$ -expander when considering the same range of temperatures. Briefly, the statement in [24] permits weaker η -expansion than $\eta = 2$, but at the cost of larger β . We focus¹ on obtaining an algorithm that still applies when β is close to β_o . For a δ -relative approximation, the high-temperature algorithm presented in Theorem 1 runs in time $(n/\delta)^{O_\epsilon(\ln d)}$ and the low-temperature algorithm runs in time $(n/\delta)^{O_\epsilon(d/\eta)}$.

For the high-temperature regime where $\beta \leq (1 - \epsilon)\beta_o$, an expansion assumption is not required and the d -regular condition can be relaxed to maximum degree d . Moreover, there is a function $q_0(\epsilon)$ such that our method only requires $q \geq q_0(\epsilon)$. In particular, this gives an FPTAS to sample in the uniqueness regime (and a bit beyond) for large enough q and d . This improves upon ([2], Theorem 2.4) which gives an FPTAS in the range $\beta \leq \frac{3}{2} \ln(q)/d$ for all $d \geq 2$ when $q \geq \exp(\Omega(d \ln d))$ (see also a slightly different version in [8] that applies up to $(1 - \epsilon)\beta_o$ but requires $q \geq \exp(\Omega(d^{3/2} \ln d))$).

The ‘gap’ in allowed values of β in Theorem 1 is partly due to our focus on reducing the necessary lower bound on q . Such a gap should not be necessary (see [21]) for large enough q , and it would be interesting to close the gap without strengthening the lower bound on q in the statement.

Guided by results for other models in statistical physics, one might expect efficient approximation algorithms for $Z_G(q, \beta)$ to exist on graphs of maximum degree d when (q, β) lie in the Gibbs uniqueness region of the infinite d -regular tree, for any $d, q \geq 3$ (For $d = 2$ the model is exactly

¹Our work could also be used to prove a version of the low-temperature statement in Theorem 1 for a strictly larger range of parameters than those handled in [24] since we extend their method with a sharper analysis.

solvable and $q = 2$ gives the Ising model, which is rather well-understood by comparison). Our methods give an FPTAS in this region for the largest-known range of d and q .

Theorem 2. *There exist absolute constants d' and q' such that for all $d \geq d'$ and $q \geq q'$, there is an FPTAS for $Z_G(q, \beta)$ on graphs of maximum degree d when $\beta > 0$ lies in the Gibbs uniqueness region of the q -color Potts model on the infinite d -regular tree.*

While our high-temperature result does not require the same novel combinatorial techniques (described below) as the low-temperature improvement, our work as a whole emphasises some interesting aspects of the polymer models that we use to prove the above results. We discuss in a concluding section the ‘critical window’ where β is close to β_0 and the difficulties associated with lowering q .

One of our key innovations for the low-temperature result, which is of independent interest, is an upper bound on the number of sets in a d -regular η -expander with a prescribed edge-boundary size b . This is reminiscent of container theorems from extremal graph theory that bound the number of independent sets in some graph or hypergraph. The origin of container theorems can be traced back to Kleitman and Winston [30, 31] but the ideas were developed significantly by Sapozhenko [35, 36], who specifically studied independent sets in expander graphs and the hypercube. Algorithmic applications of containers for counting independent sets in bipartite graphs were recently given by Jenssen and Perkins [25], also with Potukuchi [26]. The type of bound we need does not seem to relate closely to an existing container theorem, and we prove the following novel result using ideas from randomised algorithms. We say that a subset A of vertices in a graph G is connected if the induced subgraph $G[A]$ is connected.

Theorem 3. *Let $G = (V, E)$ be a d -regular η -expander and let u be a vertex in G . Then the number of connected sets $A \subset V$ of size at most $|V|/2$ such that $u \in A$ and $|\nabla(A)| = b$ is at most $d^{O((1+1/\eta)b/d)}$.*

One way of interpreting the theorem is that we show that a large-girth graph (or, if you prefer, an infinite regular tree) is extremal for the number of connected sets with a given boundary b , we explain this heuristic in the overview below.

Our methods also give an idea of the typical structure of the Potts model on η -expanders below and above the order-disorder threshold β_0 . More precisely, we show the following.

Theorem 4. *Let ϵ , d , q , and G be as in Theorem 1, with $n = |V(G)|$. As $n \rightarrow \infty$, for a colouring sampled from the ferromagnetic Potts model,*

- (1) *for $\beta < (1 - \epsilon)\beta_0$, each colour class has size $(1 + o(1))n/q$ with high probability, and*
- (2) *for $\beta > (1 + \epsilon)\beta_0$, there is a colour class that contains at least $(1 - o_d(1))n$ vertices with high probability.*

The main result of [21] gives a much more complete description of Potts model (in fact, the more general random cluster model that we describe below) on d -regular locally tree-like graphs subject to an η -expansion condition of the form $\eta = \Omega(d)$ and a stronger small-set expansion condition. It would be interesting to increase the precision of our approach to give a similar description subject to weaker conditions, e.g. for the Potts model on the hypercube. We give an example of what our methods currently reveal about the Potts model on the hypercube below.

An equivalent definition of Z is given by the Fortuin–Kasteleyn representation [13]

$$Z_G(q, \beta) = \sum_{A \subseteq E} q^{c(A)} (e^\beta - 1)^{|A|}, \quad (4)$$

where $c(A)$ is the number of connected components of the graph (V, A) on the same vertex set as G but with edge set A . In this form, Z is known as the partition function of the random cluster model, and it is easy to see that Z is a reparametrization of the Tutte polynomial of G . The random cluster

model is a distribution on subsets of edges where the probability mass of $A \subseteq E$ is proportional to $q^{c(A)}(e^\beta - 1)^{|A|}$. It is standard to change parameters via $p = e^\beta - 1$. As an illustration of the power of our new techniques, we give the following structural result for the components induced by an edge set drawn from the random cluster model on the discrete d -dimensional hypercube Q_d .

Theorem 5. *Let $\epsilon > 0$ and $d \geq d_0(\epsilon)$ for the d_0 in Theorem 1. For a constant c at least the minimum c of Theorem 1, let $q = d^c$ be a positive integer. There is a sequence $p_o = p_o(d, q)$ such that in the random cluster model on the d -regular discrete hypercube Q_d , for every $\epsilon > 0$,*

- (1) *when $p \leq (1 - \epsilon)p_o$, every component has size at most d with high probability, and*
- (2) *when $p \geq (1 + \epsilon)p_o$, there is a unique connected component of size $2^d(1 - o(1))$ with high probability.*

Some remarks are in order here. First, the condition $q = d^c$ means that the number of colours grows (poly-logarithmically) with the number of vertices in the graph. This is admittedly a slightly different regime than the more common one where q and d are both fixed. This difference turns out to be inconsequential, and the behaviour predicted by Theorem 4 ultimately holds.

Second, we note that for any constant $c > 0$ and $q = d^c$, $\ln(1 + p_o) \sim p_o \sim \frac{2 \ln q}{d}$ as $d \rightarrow \infty$. When considering the random cluster model one should not need to restrict q to the positive integers, though our methods currently require the condition that q is an integer. It would be interesting to remove this restriction, and we discuss the barriers in the concluding section.

1.2 An overview of techniques

Our algorithms follow from the cluster expansion for abstract polymer models, which was first used to design sampling algorithms for the Potts model by Helmuth et al. [23]. The method has been studied intensively since, including further works concerning the ferromagnetic Potts model [24, 2, 21] and many other works tackling various Markov random fields. At a high level, the method proceeds by transforming the partition function into a representation that corresponds to an abstract but particularly well-studied *polymer model* from statistical mechanics. The strength of this approach lies in the fact that general conditions on these models give a wealth of probabilistic and algorithmic information. The specific polymer models we rely on have appeared in several earlier works [24, 2, 21, 8, 7], and our key innovations are improved analyses of the models to obtain algorithms under significantly weaker assumptions. Let us split the rest of the discussion into the *low-temperature* regime ($\beta \geq (1 + \epsilon)\beta_o$) and the *high-temperature* regime ($\beta \leq (1 - \epsilon)\beta_o$).

Low temperature. This part contains our main contributions. Some of the difficulties we face boil down to understanding certain combinatorial questions that are of independent interest. Informally, two problems that arise are as follows. Let G be a d -regular η -expander with n vertices.

- (a) What is the number of connected, induced subgraphs of G containing a given vertex and with an edge-boundary of size b ?
- (b) What is the number of q -colorings of G that have exactly k non-monochromatic edges?

Let us first quantitatively motivate the types of answers that can be expected for these questions. For (a), suppose that G is weakly ‘locally tree-like’ in the sense that every connected set of $a \ll n$ vertices spans $O(a)$ edges and therefore has an edge-boundary of size $\Omega(da)$. Working backwards, a naive heuristic is that a typical connected subset of $\Theta(b/d)$ vertices has an edge-boundary of size b . The locally tree-like property also means that the number of connected, induced subgraphs of G containing a given vertex is roughly the number of trees with degree at most d rooted at the vertex, and hence for subgraphs on $\Theta(b/d)$ vertices this number is roughly $d^{O(b/d)}$, we can answer (a) for locally tree-like graphs. The same argument also works for graphs satisfying condition (3)

(by noting that a set of boundary size b may have at most b/η vertices) to give a bound of $d^{O(b/\eta)}$. Theorem 3 gives a stronger upper bound for graphs satisfying our expansion condition (3) by finding a small ‘certificate’ for each connected, induced subgraph with a prescribed edge-boundary size. Since the certificate is small, there cannot be too many valid certificates and an upper bound on the number of sets we are interested in follows. One of the key tools for the proof is a standard but elegant adaptation of Karger’s randomised algorithm for finding minimum cuts [27] to the problem of counting cuts; see Theorem 13. To the best of our knowledge, our application of this celebrated result in combinatorial optimisation to analyse a counting algorithm is novel, and the technique may be of independent interest. Some additional ideas used in the proof can be found in [34].

For (b), the heuristic is as follows: the simplest way that one can obtain a q -coloring of G with k non-monochromatic edges is by colouring all but k/d randomly chosen vertices, which for small k are likely to form an independent set, in the same colour. The independent set on k/d vertices can be coloured arbitrarily with the other colors, giving the required k non-monochromatic edges. This gives a lower bound of roughly $\binom{n}{k/d} q^{k/d} = (ndq/k)^{\Omega(k/d)}$. We provide another enumeration result (Lemma 16) that again justifies this heuristic for graphs satisfying the expansion condition (3). Note that the emergence of problem (b) involving colourings means that it is not straightforward to generalise our low-temperature algorithm to the random cluster model (with non-integer q).

In a calculation establishing that the polymer models we use provide efficient algorithms, we require upper bounds on $Z_\gamma(q-1, \beta)$ where γ is a connected, induced subgraph of G . Previous works used crude bounds here, and we use better bounds from [40] that are tight when γ is isomorphic to K_{d+1} , or in the triangle-free case, complete bipartite graphs $K_{d,d}$. This is still not ideal as an η -expander cannot contain such subgraphs, but it is generally difficult to prove improved bounds, or even identify graph properties that would allow an improved bound in such problems [34]. This is another place in which our methods require the Potts model, as such bounds are not known for the random cluster model. In this way, our algorithmic work motivates new questions in extremal graph theory; see Conjectures 25 and 26.

High temperature. The main idea to handle this regime is to switch to the random cluster model. The primary reason for this is to access a convenient representation as a polymer model. Using the Edwards–Sokal coupling [12], one can view both of these models on the same probability space. The coupling gives an algorithm for converting a random cluster sampling algorithm to a Potts model sampling algorithm and vice versa. The polymer model we use was also generalised to a partition function related to Unique Games in [8]. The analysis given in this paper improves upon the analyses in [2] and [8]. Moreover, this part of the proof works for *any* graph of maximum degree d (i.e. we can dispense with d -regularity and with an expansion assumption), a fact present in these earlier works which does not seem to be well-known in general.

1.3 Related work

Algorithmic applications of the cluster expansion originate in [23], and for the ferromagnetic Potts model in particular there are two works of note that establish algorithms at all temperatures on subgraphs of the integer lattice \mathbb{Z}^d [2] and on regular graphs with strong expansion [21]. These results rely on q being at least exponentially large in d (and even larger if the expansion guarantee is as weak as that of Theorem 1), which is a substantial drawback but in return the approaches work for all temperatures. The focus of our work is slightly different, and essentially we pay for working with smaller q that grows only polynomially with d (along with much weaker expansion) by accepting a small gap in our techniques around the threshold β_o .

Our main tool is the cluster expansion for abstract polymer models, as in [23, 24, 2, 21] (and many more works), where deterministic approximation algorithms for partition functions follow

from convergent cluster expansions by approximating a truncated series. See also Barvinok's work (e.g. the monograph [1]) which follows a similar series approximation approach. A twist on the cluster expansion method originating in [9] uses a Markov chain to sample from an abstract polymer model. The approaches of [24, 9] were combined and generalised in [14] to give algorithms for general spin systems on bipartite expanders. We do not pursue Markov-chain-based sampling from polymer models here as the conditions required in [9, 14, 3] (i.e. upper bounds on polymer weights) are stronger than what we can guarantee with our current methods, and we do not want to restrict our attention to bipartite graphs.

A popular approach to approximate counting and sampling is to simulate a Markov chain whose stationary distribution is the Gibbs measure of interest. For the Potts model there are two main Markov chains, the *Glauber dynamics* and the *Swendsen–Wang dynamics*. The main result of [6] shows for graphs of maximum degree d that Glauber dynamics is rapidly mixing for β up to a certain threshold and slow mixing for β at least some slightly larger threshold. For Glauber dynamics on the random regular graph there is an improved understanding: the Gibbs uniqueness threshold marks the change from rapid mixing to torpid mixing [4, 10]. For comparison with our results, we note that each of the three thresholds discussed above lie below β_o and are asymptotic to $\ln(q)/d$ as $d \rightarrow \infty$ (which is asymptotically half of $\beta_o \sim 2 \ln(q)/d$ as $d \rightarrow \infty$). It should be noted though, that these results hold for *every* q , while the results in this paper are restricted to large q . The picture for Swendsen–Wang dynamics on the random regular graph is less well-understood; see [10] and the references therein for details. An alternative approach to sampling from the ferromagnetic Potts model (and random cluster model) on the random regular graph in the Gibbs uniqueness regime of the ferromagnetic Potts model was given in [5].

1.4 organization

Section 2 is an introduction to abstract polymer models and their use in obtaining approximate counting and sampling algorithms. In Section 3 we prove our algorithmic result – Theorem 1 – for the low-temperature case ($\beta > (1 + \epsilon)\beta_o$) using the solutions to problems (a), (b), and extremal results on partition functions (Theorem 3, Lemma 8, and Lemmas 11 and 12 respectively). Section 4.1 is dedicated to the proof of our container-like result Theorem 3, solving problem (a); and Section 4.2 to the proof of Lemma 8, solving problem (b). In Section 5, we prove the extremal results, Lemmas 11 and 12. Section 6 consists of the proof of the high-temperature case ($\beta < (1 - \epsilon)\beta_o$) in Theorem 1. In Section 7 we use our results to characterise the structure of the Potts model, proving Theorem 4. Finally, Section 8 is dedicated to the structure of the random cluster model on the hypercube, namely the proof of Theorem 5. We conclude with a discussion of further directions to pursue.

2. Preliminaries

2.1 Abstract polymer models

An *abstract polymer model* [17] is given by a set \mathcal{P} of polymers, a compatibility relation \sim on polymers, and a weight w_γ associated to each polymer $\gamma \in \mathcal{P}$. In applications, polymers might be combinatorial objects such as vertex subsets or connected subgraphs of a graph. Let Ω denote the collection of all sets of pairwise compatible polymers, including the empty set (i.e. the set of independent sets in the graph on \mathcal{P} with edges between incompatible polymers). Then the polymer model partition function is

$$\Xi = \sum_{\Lambda \in \Omega} \prod_{\gamma \in \Lambda} w_\gamma.$$

The associated Gibbs measure ν on Ω is given by $\nu(\Lambda) = \prod_{\gamma \in \Lambda} w_\gamma / \Xi$, in much the same way as the Gibbs measures associated to the partition function of the Potts and random cluster models. It may be helpful to observe that Ξ is the multivariate independence polynomial of the graph on \mathcal{P} with edges between incompatible polymers, evaluated at the polymer weights given by w , and hence ν is the associated hard-core measure.

The *cluster expansion* of Ξ is a formal power series for $\ln \Xi$ with terms given by *clusters*. A cluster Γ is an ordered tuple of polymers, and to each Γ we associate an *incompatibility graph* $H(\Gamma)$. This graph has vertex set Γ (i.e. a vertex for each polymer present in Γ with multiplicity) and each pair of vertices representing incompatible polymers is an edge of $H(\Gamma)$. By convention, a polymer is incompatible with itself. We write \mathcal{C} for the set of all clusters. The formal power series corresponding to the cluster expansion is then

$$\ln \Xi = \sum_{\Gamma \in \mathcal{C}} \phi(\Gamma) \prod_{\gamma \in \Gamma} w_\gamma,$$

where ϕ is the Ursell function

$$\phi(\Gamma) = \sum_{\substack{F \subseteq E(H(\Gamma)) \\ (V(H(\Gamma)), F) \text{ connected}}} (-1)^{|F|}.$$

We remark that for any polymer γ , the tuples (γ) , (γ, γ) , \dots are valid clusters, and so the cluster expansion is an infinite sum. We use the following theorem to establish convergence and an important tail bound.

Theorem 6 (Kotecký–Preiss [29]). *If there exist functions $f : \mathcal{P} \rightarrow [0, \infty)$ and $g : \mathcal{P} \rightarrow [0, \infty)$ such that for all $\gamma \in \mathcal{P}$,*

$$\sum_{\gamma' \not\sim \gamma} w_{\gamma'} e^{f(\gamma') + g(\gamma')} \leq f(\gamma), \quad (5)$$

then the cluster expansion converges absolutely. Moreover, if we let $g(\Gamma) = \sum_{\gamma \in \Gamma} g(\gamma)$, then for every polymer γ ,

$$\sum_{\substack{\Gamma \in \mathcal{C} \\ \Gamma \not\sim \gamma}} \left| \phi(\Gamma) \prod_{\gamma' \in \Gamma} w_{\gamma'} \right| e^{g(\Gamma)} \leq f(\gamma), \quad (6)$$

where we write $\Gamma \not\sim \gamma$ to mean that the cluster Γ contains a polymer that is incompatible with γ .

2.2 Algorithms from convergent cluster expansions

Approximate counting and sampling algorithms for polymer models follow from Theorem 6 in a rather standard way due to [23] and subsequent works. In this paper we consider two polymer models that we define for an n -vertex, d -regular graph G , and then we consider approximating Ξ or approximately sampling from ν with time complexities measured in terms of n and a desired approximation error δ .

For a polymer model as above, let $g : \mathcal{P} \rightarrow [0, \infty)$ be as in Theorem 6, and extend g to clusters $\Gamma \in \mathcal{C}$ by $g(\Gamma) = \sum_{\gamma \in \Gamma} g(\gamma)$. An FPTAS for Ξ follows from a few steps that one must check can be done efficiently:

- (1) Use (6) to show that there exists some $L = L(\epsilon)$ such that

$$\Xi(L) := \exp \left(\sum_{\substack{\Gamma \in \mathcal{C} \\ g(\Gamma) \leq L}} \phi(\Gamma) \prod_{\gamma \in \Gamma} w_{\gamma} \right)$$

is a close enough relative approximation of Ξ .

- (2) List all clusters $\Gamma \in \mathcal{C}$ such that $g(\Gamma) \leq L$.
 (3) For each such cluster Γ , compute $\phi(\Gamma)$ and $\prod_{\gamma \in \Gamma} w_{\gamma}$.
 (4) Compute $\Xi(L)$ directly from the above quantities.

The running time of this algorithm depends on L , the definition of ‘close enough’, the time complexity of listing clusters and of computing polymer weights, ϕ and g . It has been established [23, 24, 2] that for the polymer models we consider, the above steps can be carried out sufficiently fast for the definition of an FPTAS.

Provided that the above algorithm yields an FPTAS for Ξ , there is a generic polynomial-time approximate sampling algorithm for ν that follows from a *self-reducibility* property of the polymer model; see ([23], Theorem 10). We omit the details and summarise these facts as the following theorem which applies when \mathcal{P} is (a subset of) the connected subgraphs of some bounded-degree graph G . For a subgraph γ , we use ν_{γ} to denote the number of vertices in γ . The important fact is that both the high-temperature and low-temperature polymer models we study satisfy the hypotheses of the result below.

Theorem 7 ([24], Theorem 8 and [23], Theorem 10). *Fix d and let \mathcal{G} be some class of graphs of maximum degree at most d . Suppose the following hold for a polymer model with partition function $\Xi(G)$ where the polymers \mathcal{P} are connected subgraphs of some $G \in \mathcal{G}$, and with decay function g .*

- (1) *There exist constants $c_1, c_2 > 0$ such that given a connected subgraph γ , determining whether γ is a polymer in \mathcal{P} , and computing w_{γ} and $g(\gamma)$ can be done in time $O(\nu_{\gamma}^{c_1} e^{c_2 \nu_{\gamma}})$.*
 (2) *There exists $\rho = \rho(d) > 0$ so that for every $G \in \mathcal{G}$ and every $\gamma \in \mathcal{P}$, $g(\gamma) \geq \rho \nu_{\gamma}$.*
 (3) *The Kotecký–Preiss condition (5) holds with the given function g .*

Then there is an FPTAS for the partition function $\Xi(G)$ of the polymer model for $G \in \mathcal{G}$ that with approximation error δ and $n = |V(G)|$ has running time $n \cdot (n/\delta)^{O((\ln d + c_2)/\rho)}$. There is also a polynomial-time sampling scheme for the associated polymer measure ν .

3. Low-temperature algorithms

3.1 Sketch of low-temperature argument

For the low-temperature regime, we decompose the partition function into colourings that have a majority colour, and the remaining colourings. Our polymer model can only represent colourings with a majority colour, so it is vital to prove that the remaining colourings have a small contribution to the partition function, which is the content of our Lemma 8. Working with fewer colours and smaller expansion makes this step more challenging than in [24]. The core of the proof involves obtaining upper bounds on the number of vertex-colorings of a graph that induce few monochromatic edges. In light of the heuristic (b), which relates colourings with few monochromatic edges to cuts, this is done by analysing a Karger-like random sequence of edge contractions and the probability that this yields a particular cut.

We then handle each majority colour with a polymer model where the polymers are connected, induced subgraphs of G whose vertices do not receive the majority colour. The weight of a polymer is related to the $(q - 1)$ -spin partition function on that subgraph. Since we are working

with colourings with a majority colour, the ‘defects’ (vertices that don’t have the majority colour) occupy at most half the vertices, which is a global constraint on the polymers that we wish to avoid. We pass to a different polymer model which requires only that each polymer has at most half the vertices. This partition function is clearly an overestimate; however, in Lemma 9 we show that this overestimation is negligible.

All that remains is to verify the Kotecký–Preiss condition for this polymer model. There is, however, another challenge here: the aforementioned $(q - 1)$ -spin partition function plays a significant role in this computation, and one needs to bound it in order to proceed. Fortunately, our understanding of extremal bounds on such functions has progressed recently, most notably due to Sah *et al.* [40]. We use their results to derive bounds on the partition functions of the subgraphs that make up our polymers. For triangle-free graphs the available bounds are stronger, which lets us work with weaker expansion in this case, and ultimately with the hypercube. The extremal results we derive from [40] are the content of Lemmas 11 and 12.

The final ingredient in verifying the Kotecký–Preiss condition is a bound on the number of connected subsets of vertices in G with a given edge-boundary size. At first glance, this looks similar to a lemma used in [25, 26] that gives a bound on the number of connected subsets of vertices with a given *vertex-boundary* size (in bipartite graphs), which followed from mild adaptations of existing container methods [35, 36].

While the case of edge boundaries does not seem to follow directly from these methods, we have developed Theorem 3 inspired by these previous works. At a high-level, container methods work by finding an ‘encoding’ or ‘certificate’ for each object of the type we wish to enumerate. If these certificates are small then the number of certificates, and thus objects, must be relatively small. In our case, we group sets of vertices A with a fixed edge-boundary B of size b by identifying a ‘core’ subset $A_0 \subset A$ of size $O(b/d)$. We show these core sets must be connected in G^7 , and hence their number can be controlled. In this sense, each core set A_0 certifies a container which contains sets A (of the type whose number we wish to bound) obtained by growing A_0 appropriately, see (19). To complete the bound on the number of sets A of interest, we must bound for each core A_0 the number of possible A , or equivalently edge boundaries B , that can be associated with A_0 . We achieve this using Karger’s algorithm involving randomised edge contractions to count cuts.

3.2 Polymer model

In Theorem 1, we deal with the class of d -regular 2-expanders and triangle-free d -regular 1-expanders. Both these families of graphs satisfy

$$\text{every set } A \subset V \text{ such that } 2 \leq |A| \leq |V|/2 \text{ satisfies } |\nabla(A)| \geq 2d - 2. \quad (7)$$

In particular, this also holds for the hypercube Q_d . Property (7) is immediate in d -regular 2-expanders and follows from Mantel’s theorem [33] in the triangle-free 1-expanding case. To see this, note that the expansion condition suffices for $|A| \geq 2d - 2$, and if $|A| < 2d - 2$ then we use the bound

$$\nabla(A) = d|A| - 2|E(G[A])| \geq d|A| - \frac{|A|^2}{2},$$

which follows from Mantel’s theorem. For $2 \leq |A| \leq 2d - 2$, this lower bound is at least $2d - 2$ and property (7) follows.

The min-cut of a graph is the minimum number of edges that need to be deleted in order to disconnect the remaining graph. Any d -regular graph G that satisfies (7) also satisfies

$$\text{the min-cut of } G \text{ has size at least } d. \quad (8)$$

In other words, the smallest set of edges that can be deleted to disconnect the graph is simply the set of all edges incident to any given vertex. Throughout this section we will use that G satisfies (7), and therefore (8).

We now define a polymer model for the low-temperature regime, following [24]. Here we take $\epsilon \in (0, 1)$; then, for some $d_0(\epsilon)$ and an absolute constant c , we have $d \geq d_0$, $q \geq d^c$, and $\beta \geq (1 + \epsilon)\beta_o(q, d)$. We also assume that q is large enough (by ensuring that d_0 and c are large enough) that this implies $\beta \geq (2 + \epsilon) \ln(q)/d$.

Let a polymer γ be a connected, induced subgraph of G with at most $n/2$ vertices. For each polymer γ , let v_γ and e_γ be the number of vertices and edges respectively. Let ∇_γ denote the number of boundary edges (i.e. edges of G with one endpoint in the vertex set of γ), and let $\eta_\gamma := \nabla_\gamma / v_\gamma$, be a measure of the expansion of the vertex set of γ in G . The polymer weights w_γ are given by

$$w_\gamma = e^{-\beta(\nabla_\gamma + e_\gamma)} \cdot Z_\gamma(q - 1, \beta),$$

where we interpret γ as a graph and hence $Z_\gamma(q - 1, \beta)$ is the Potts partition function of γ with $q - 1$ colours. We say that two polymers γ and γ' are incompatible if $\text{dist}_G(\gamma, \gamma') \leq 1$. That is, $\gamma \not\sim \gamma'$ if and only if they share a vertex or one contains a neighbour of the other. As before, we write Ω for the collection of pairwise compatible sets of polymers and $\Xi = \Xi_G(q, \beta)$ for the partition function of this polymer model.

For our low-temperature algorithm, there is not a precise correspondence between $Z_G(q, \beta)$ and Ξ , and we must establish that approximating Ξ does in fact yield an approximation of $Z_G(q, \beta)$. The set $[q]^{V(G)}$ of q -colourings of $V(G)$ admits a partition into S_0, S_1, \dots, S_q where S_0 is the set of colourings such that each colour occupies at most $n/2$ vertices and for $j \in [q]$, S_j is the set of colourings in which j is the majority colour (i.e. $S_j = \{\sigma \in [q]^{V(G)} : |\sigma^{-1}(j)| > n/2\}$). We have the following lemma whose proof is postponed to Section 4.2.

Lemma 8. *There exists an absolute constant c such that the following holds. For every $\epsilon \in (0, 1)$ there exists $d_0(\epsilon)$ such that for every $d \geq d_0$, $q \geq d^c$ and $\beta \geq (1 + \epsilon)\beta_o(q, d)$, we have*

$$\frac{\sum_{\sigma \in S_0} e^{\beta m(G, \sigma)}}{Z_G(q, \beta)} \leq q^{-\Omega(\frac{\epsilon n}{d})}.$$

Let us define the partition function-like quantity

$$\tilde{\Xi} = \tilde{\Xi}_G(q, \beta) := \sum_{\substack{\Lambda \in \Omega \\ \|\Lambda\| < n/2}} \prod_{\gamma \in \Lambda} w_\gamma,$$

where $\|\Lambda\| = \sum_{\gamma \in \Lambda} v_\gamma$ is the size of a set Λ of pairwise compatible polymers. We write $\tilde{\Omega} = \{\Lambda \in \Omega : \|\Lambda\| < n/2\}$ and note that $\tilde{\Xi}$ is not the true partition function of the polymer model because of the global constraint $\|\Lambda\| < n/2$. In terms of the Potts partition function $Z_G(q, \beta)$, the quantity $\tilde{\Xi}$ faithfully represents the colourings of G which have a majority colour because

$$\sum_{j \in [q]} \sum_{\sigma \in S_j} e^{\beta m(G, \sigma)} = e^{\beta dn/2} \sum_{j \in [q]} \sum_{\substack{\Lambda \in \Omega \\ \|\Lambda\| < n/2}} \prod_{\gamma \in \Lambda} w_\gamma = q e^{\beta dn/2} \tilde{\Xi}.$$

To see this, consider a colouring $\sigma \in [q]^V$ with majority colour j and decompose $\sigma^{-1}([q] \setminus \{j\}) \subset V$ into sets such that their pairwise distance in G is at least two. These sets induce subgraphs which form pairwise compatible polymers. Hence for any $j \in [q]$ the colourings with majority colour j are in bijection with $\{\Lambda \in \Omega : \|\Lambda\| < n/2\}$. The weight function and incompatibility criterion have been carefully defined to make the above calculation work on the level of partition function contribution as well.

Rewriting Lemma 8 in this terminology gives us that

$$\frac{qe^{\beta dn/2}}{Z_G(q, \beta)} \tilde{\Xi} = 1 - q^{-\Omega(\frac{\epsilon n}{d})}.$$

Thus, approximating $\tilde{\Xi}$ is a viable avenue for approximating $Z_G(q, \beta)$. There is one final complication, however, related to the fact that $\tilde{\Xi}$ is not precisely the partition function Ξ of the low-temperature polymer model defined above. Since the standard algorithmic setup (Theorem 7) allows us to approximate Ξ and approximately sample from the associated polymer measure ν , we must additionally show that restricting to Λ such that $\|\Lambda\| < n/2$ does not harm the approximation too much. The proof of the low-temperature case of Theorem 1 therefore follows from the following two lemmas and the standard algorithmic setup of Theorem 7.

Lemma 9. *For Ξ and $\tilde{\Xi}$ as defined above, we have*

$$\tilde{\Xi} \leq \Xi \leq \tilde{\Xi} \left(1 + e^{-\Omega(n/d)}\right).$$

Lemma 10. *For every $\epsilon \in (0, 1)$, there exist $d_0(\epsilon)$ and an absolute constant c such that for every $d \geq d_0$, $q \geq d^c$ and $\beta \geq (1 + \epsilon)\beta_0(q, d)$, the Kotecký–Preiss condition (see (5) in Theorem 6) holds for the low-temperature polymer model with $f(\gamma) = \frac{\epsilon v_\gamma \ln q}{4d}$ and $g(\gamma) = \frac{\epsilon \nabla_\gamma \ln q}{4d} + \frac{v_\gamma}{d}$.*

Note that the left-hand side of condition (5) decreases if g decreases, and hence the above lemma also implies that Theorem 6 holds for any smaller choice of g . We prove the above results, starting with Lemma 10, in the following subsections.

3.3 Proof of Lemma 10

We first give the proof in the case that G is a 2-expander, and then show the modifications necessary for triangle-free 1-expanders.

We use $N(\gamma) = N_G(\gamma) = \{v : \exists u \in \gamma \text{ such that } \{u, v\} \in E\}$ for the *neighbourhood* of γ in G and $N_G[\gamma] = N_G(\gamma) \cup V(\gamma)$ for the *closed neighbourhood* of γ . We upper bound the left-hand side of condition (5), here noting that the incompatibility relation on polymers yields

$$\sum_{\gamma' \not\sim \gamma} w_{\gamma'} e^{f(\gamma') + g(\gamma')} \leq \sum_{u \in N_G[\gamma]} \sum_{\gamma' \ni u} w_{\gamma'} e^{f(\gamma') + g(\gamma')} = \sum_{u \in N_G[\gamma]} \sum_{b \geq d} \sum_{\substack{\gamma' \ni u \\ \text{s.t. } \nabla_{\gamma'} = b}} w_{\gamma'} e^{f(\gamma') + g(\gamma')},$$

Since $|N_G[\gamma]| \leq v_\gamma + \nabla_\gamma$, it suffices to show that for all $u \in V$ we have

$$\sum_{b \geq d} \sum_{\substack{\gamma' \ni u \\ \text{s.t. } \nabla_{\gamma'} = b}} w_{\gamma'} e^{f(\gamma') + g(\gamma')} \leq \min_{\gamma \in \mathcal{P}} \frac{f(\gamma)}{v_\gamma + \nabla_\gamma}. \quad (9)$$

We use the following lemma, whose proof is postponed to Section 5, to bound the weights $w_{\gamma'}$ that appear in (9).

Lemma 11. *For any polymer $\gamma \in \mathcal{P}$ we have*

$$Z_\gamma(q - 1, \beta) \leq e^{\beta e_\gamma} q^{v_\gamma - \frac{2e_\gamma}{d} + \frac{e_\gamma}{\binom{d+1}{2}}} 2^{\frac{e_\gamma}{\binom{d+1}{2}}}.$$

To bound the exponent of q in Lemma 11, observe that $dv_\gamma = 2e_\gamma + \nabla_\gamma \geq 2e_\gamma$ by a double-counting argument, and recall that $\eta_\gamma = \nabla_\gamma / v_\gamma$ to obtain

$$v_\gamma - \frac{2e_\gamma}{d} + \frac{e_\gamma}{\binom{d+1}{2}} \leq \frac{\nabla_\gamma}{d} \left(1 + \frac{1}{\eta_\gamma}\right).$$

Plugging these facts and Lemma 11 into the definition of w_γ gives

$$w_\gamma = e^{-\beta(e_\gamma + \nabla_\gamma)} Z_\gamma(q-1, \beta) \leq e^{-\beta \nabla_\gamma} q^{\frac{\nabla_\gamma}{d}(1 + \frac{1}{\eta_\gamma})} 2^{\frac{\nabla_\gamma}{d\eta_\gamma}},$$

and using $\beta \geq (2 + \epsilon) \ln(q)/d$ we have

$$e^{-\beta \nabla_\gamma} \leq e^{-(2+\epsilon)\nabla_\gamma \frac{\ln q}{d}} = q^{-(2+\epsilon)\frac{\nabla_\gamma}{d}}.$$

We assume q is large enough (by assuming d_0 and c are large enough) so that $2^{\frac{\nabla_\gamma}{d\eta_\gamma}} \leq q^{\frac{\nabla_\gamma}{d} \cdot \frac{\epsilon}{4}}$, and thus

$$w_\gamma \leq q^{-\frac{\nabla_\gamma}{d} \left(1 + \frac{3\epsilon}{4} - \frac{1}{\eta_\gamma}\right)}.$$

Using our choice of f and g , we also have

$$e^{f(\gamma)+g(\gamma)} = q^{\frac{\epsilon}{4d}(v_\gamma + \nabla_\gamma)} \cdot e^{v_\gamma/d} = q^{\frac{\nabla_\gamma}{d} \frac{\epsilon}{4} \left(1 + \frac{1}{\eta_\gamma}\right)} \cdot e^{\frac{v_\gamma}{d}},$$

which combined with the bound $\eta_\gamma \geq 2$ gives

$$w_\gamma e^{f(\gamma)+g(\gamma)} \leq q^{-\frac{\nabla_\gamma}{d} \left(1 + \frac{3\epsilon}{4} - \frac{1}{\eta_\gamma} - \frac{\epsilon}{4\eta_\gamma} - \frac{\epsilon}{4}\right)} e^{\frac{v_\gamma}{d}} \leq q^{-\frac{\nabla_\gamma}{d} \left(\frac{1}{2} + \frac{3\epsilon}{8}\right)} e^{\frac{v_\gamma}{d}} \leq q^{-\frac{\nabla_\gamma}{2d}}, \quad (10)$$

where the final inequality holds again when q is large enough in terms of ϵ .

We now use our container-type result. Applying Theorem 3 and (10) to bound the left-hand side of (9), we have (for q large enough that the series converges)

$$\sum_{b \geq d} \sum_{\substack{\gamma' \ni u \\ \text{s.t. } \nabla_{\gamma'} = b}} w_{\gamma'} e^{f(\gamma')+g(\gamma')} \leq \sum_{b \geq d} d^{O((1+1/\eta)b/d)} q^{-\frac{b}{2d}} \leq \frac{d^{c'}/\sqrt{q}}{1 - (d^{c'}/\sqrt{q})^{1/d}},$$

where the absolute constant $c' \geq 4$ is large enough that the term $O((1+1/\eta)b/d)$ is at most $c'b/d$ for all $b \geq d$ and all $\eta \geq 1$. Note that the right-hand side above is a decreasing function of q . To satisfy (9) we want this to be at most

$$\min_{\gamma \in \mathcal{P}} \frac{f(\gamma)}{v_\gamma + \nabla_\gamma} = \frac{\epsilon \ln q}{4d} \min_{\gamma \in \mathcal{P}} \frac{v_\gamma}{v_\gamma + \nabla_\gamma} = \frac{\epsilon \ln q}{4d(d+1)},$$

which is an increasing function of q . Thus, it suffices to take q larger than some known value for the desired inequality holds. It is straightforward to check that $q = d^{3c'}$ suffices for all d large enough.

3.4 The triangle-free case

If G is triangle free then a stronger upper bound on partition functions of the form $Z_\gamma(q-1, \beta)$ holds.

Lemma 12. *If G is triangle free, then*

$$Z_\gamma(q-1, \beta) \leq e^{\beta e_\gamma} \cdot q^{v_\gamma - \frac{2e_\gamma}{d} + \frac{e_\gamma}{d^2}} \cdot 2^{\frac{e_\gamma}{d^2}}.$$

Plugging this inequality in the computation in the previous section, we have

$$v_\gamma - \frac{2e_\gamma}{d} + \frac{e_\gamma}{d^2} \leq \frac{\nabla_\gamma}{d} \left(1 + \frac{1}{2\eta_\gamma} \right),$$

and hence only assuming $\eta \geq 1$ we have the bound

$$w_\gamma e^{f(\gamma)+g(\gamma)} \leq q^{-\frac{\nabla_\gamma}{d} \left(\frac{1}{2} + \frac{\epsilon}{4} \right)} e^{\frac{v_\gamma}{d}} \leq q^{-\frac{\nabla_\gamma}{2d}}.$$

We may then apply Theorem 3 to reach the same conclusion as before.

3.5 Proof of Lemma 9

Let ν and $\tilde{\nu}$ denote the probability distributions on Ω given by the partition functions Ξ and $\tilde{\Xi}$ respectively, meaning that $\nu(\Lambda) = \prod_{\gamma \in \Lambda} w_\gamma / \Xi$ for any $\Lambda \in \Omega$ and $\tilde{\nu}(\Lambda) = 0$ if $\|\Lambda\| \geq n/2$ but $\tilde{\nu}(\Lambda) = \prod_{\gamma \in \Lambda} w_\gamma / \tilde{\Xi}$ otherwise. We briefly prove a large deviation inequality on the size of the defects, similar to [25, 26]. Here, we carefully make use of the term $e^{v_\gamma/d}$ in the definition of $g(\gamma)$ and the monotonicity of condition (5) in g . This immediately implies that Theorem 6 holds for an alternative polymer model on the same polymers but with weights

$$w'_\gamma := w_\gamma \cdot e^{\frac{v_\gamma}{d}},$$

because for the same choice $f(\gamma) = \frac{\epsilon v_\gamma \ln q}{4d}$ and now taking $g(\gamma) = \frac{\epsilon \nabla_\gamma \ln q}{4d}$, condition (5) is the same for both models.

Write

$$\Xi' := \sum_{\Lambda \in \Omega} \prod_{\gamma \in \Lambda} w'_\gamma = \sum_{\Lambda \in \Omega} \prod_{\gamma \in \Lambda} w_\gamma e^{v_\gamma/d},$$

and let Λ denote a random set of pairwise compatible polymers drawn from ν . Observe that Ξ' naturally appears in an expression for $\mathbb{E}e^{\|\Lambda\|/d}$ as follows:

$$\mathbb{E}e^{\|\Lambda\|/d} = \sum_{\Lambda \in \Omega} \nu(\Lambda) e^{\|\Lambda\|/d} = \frac{1}{\Xi} \sum_{\Lambda \in \Omega} \prod_{\gamma \in \Lambda} (w_\gamma e^{v_\gamma/d}) = \frac{\Xi'}{\Xi}.$$

Then Theorem 6 and the fact that $\Xi \geq 1$ gives us

$$\ln \mathbb{E}e^{\|\Lambda\|/d} = \ln \frac{\Xi'}{\Xi} \leq \ln \Xi' = \sum_{\Gamma \in \mathcal{C}} \phi(\Gamma) \prod_{\gamma \in \Gamma} w'_\gamma,$$

where the final equality is the cluster expansion for the abstract polymer model defined by Ξ' , and hence the sum is over clusters $\Gamma \in \mathcal{C}$ as defined in Section 2. We can crudely upper bound this cluster expansion of by summing over all polymers of the form $\gamma = (\{v\}, \emptyset)$ induced by a vertex and applying the tail bound (6). This gives

$$\ln \Xi' \leq \sum_{v \in V(G)} \sum_{\substack{\Gamma \in \mathcal{C} \\ \Gamma \not\supset (\{v\}, \emptyset)}} \left| \phi(\Gamma) \prod_{\gamma \in \Gamma} w'_\gamma \right| \quad (11)$$

$$\leq \sum_{v \in V(G)} \sum_{\substack{\Gamma \in \mathcal{C} \\ \Gamma \not\supset (\{v\}, \emptyset)}} q^{-\epsilon/4} \left| \phi(\Gamma) \prod_{\gamma \in \Gamma} w'_\gamma \right| e^{g(\Gamma)} \quad (12)$$

$$\leq n \cdot \frac{\epsilon \ln q}{4q^{\epsilon/4}d}. \quad (13)$$

Line (12) holds because every cluster $\Gamma \not\sim (\{v\}, \emptyset)$ contains at least one polymer (which has at least one vertex) and hence using (8) we have

$$g(\Gamma) = \frac{\epsilon \ln q}{4d} \sum_{\gamma \in \Gamma} \nabla_{\gamma} \geq \frac{\epsilon \ln q}{4}.$$

To obtain line (13) we use (6) with $f(\{v\}, \emptyset) = \frac{\epsilon \ln q}{4d}$.

Combining the above calculations, we have $\mathbb{E}e^{\|\Lambda\|/d} \leq e^{\frac{n}{4d}}$, and hence Markov's inequality gives

$$\mathbb{P}(\|\Lambda\| \geq n/2) = \mathbb{P}(e^{\|\Lambda\|/d} \geq e^{n/(2d)}) \leq \exp\left(\frac{\epsilon n \ln q}{4q^{\epsilon/4}d} - \frac{n}{2d}\right) = \exp\left(-\frac{n}{4d}\right), \quad (14)$$

which implies $\nu(\tilde{\Omega}) \geq 1 - e^{-n/(4d)}$. Here we have used the fact that q is large in terms of ϵ so $\frac{\epsilon \ln q}{4q^{\epsilon/4}} \leq \frac{1}{4}$. We would like to point out that similar computations are required for the proofs of Theorems 4 and 5.

Since the probabilities for $\tilde{\nu}$ on outcomes in $\tilde{\Omega}$ are made by redistributing the probability mass $\nu(\Omega \setminus \tilde{\Omega}) \leq e^{-n/(4d)}$ onto outcomes in $\tilde{\Omega}$, we immediately have a bound on the total variation distance

$$\|\tilde{\nu} - \nu\|_{TV} \leq e^{-n/(4d)}.$$

The conclusion follows because by the definitions of $\tilde{\nu}$ and total variation distance we have

$$0 \leq \tilde{\nu}(\tilde{\Omega}) - \nu(\tilde{\Omega}) = 1 - \nu(\tilde{\Omega}) \leq \|\tilde{\nu} - \nu\|_{TV},$$

and hence dividing by $\nu(\tilde{\Omega})$ gives

$$0 \leq \frac{1}{\nu(\tilde{\Omega})} - 1 = \frac{\mathbb{E}}{\tilde{\mathbb{E}}} - 1 \leq \frac{\|\tilde{\nu} - \nu\|_{TV}}{\nu(\tilde{\Omega})} \leq \frac{e^{-n/(4d)}}{1 - e^{-n/(4d)}}.$$

4. Enumerative results

In this section, we prove Theorem 3 and Lemma 8.

4.1 Connected sets with small edge boundaries: Proof of Theorem 3

The main combinatorial result underlying our container-like lemma is the following standard adaptation of Karger's algorithm to count α -min-cuts.

Theorem 13 (Karger [27]). *Let G be a graph whose min-cut has t edges, and let $\alpha \geq 1$ be a real number. The number of cuts in G with at most αt edges is at most $\binom{n}{2\alpha} \cdot 2^{2\alpha}$.*

We also require a well-known result giving an upper bound to the number of connected induced subgraphs containing a fixed vertex.

Proposition 14 ([28], Vol. 3 p. 396, Ex.11). *The number of rooted labelled trees of maximum degree Δ on $n \geq 1$ vertices is*

$$\frac{\binom{\Delta n}{n}}{(\Delta - 1)n + 1} \leq (e\Delta)^{n-1},$$

and hence, for a graph of maximum degree Δ and a fixed vertex v , the number of connected induced subgraphs on n vertices containing v is also at most this number.

With this, we now prove Theorem 3.

Proof of Theorem 3. We first introduce some notation. Recall V and E are the sets of vertices and edges of G respectively. For $u, v \in V$, we use $\text{dist}_G(u, v)$ to denote the distance in G from u to v , that is, the number of edges on a shortest path between u and v . For every $k \in \mathbb{N}$, we define the k th neighbourhood of u in G as $N_G^k(u) = \{v \in V : 1 \leq \text{dist}_G(u, v) \leq k\}$. Then $N_G(u) = N_G^1(u)$. We may also define the k th neighbourhood of a set of vertices $A \subseteq V$ as $N_G^k(A) = \cup_{u \in A} N_G^k(u)$. The k th power of the graph G is the graph G^k on vertex set V where $\{u, v\} \in E(G^k)$ if and only if $1 \leq \text{dist}_G(u, v) \leq k$. Observe that $N_{G^k}(u) = N_G^k(u)$.

Let I be the vertex-edge incidence graph of G , which is a bipartite graph with vertex set $V \sqcup E$ and edges $\{(u, e) \in V \times E \mid u \in e\}$. Let J be the bipartite graph with vertex set $V \sqcup E$ and edges given by $\{(u, e) \in V \times E \mid \exists v \in N_G(u) \text{ s.t. } v \in e\}$. That is, (u, e) forms an edge of J if and only if e is an edge incident in G to a neighbour of u (this includes all edges incident to u itself). Note that I is a subgraph of J .

Observe that for every $u \in V$, the degree of u in J is at least $\binom{d+1}{2} \geq d^2/2$ and at most d^2 . To see the lower bound, consider that each of the d neighbours of u is itself incident to d edges, and the number of distinct edges incident to $N_G(u)$ is minimised when $N_G(u)$ forms a clique on $d+1$ vertices. The upper bound holds because the maximum number of edges of G incident to $N_G(u)$ is d^2 (which occurs if and only if u is contained in no cycles of length 3 or 4).

Now consider a set of vertices $A \subseteq V$ of size at most $n/2$ such that $G[A]$ (the subgraph of G induced by A) is connected and $x \in A$. Let $B = \nabla(A)$ and $|B| = |\nabla(A)| = b$.

We then have $N_I(A) = B \sqcup W$ where $W = E(G[A])$. Note that by summing degrees, $d|A| = b + 2|W|$, and by the expansion assumption (3) we have $b \geq \eta|A|$. Hence,

$$|W| \leq \frac{b}{2} \left(\frac{d}{\eta} - 1 \right) \leq \frac{bd}{2\eta}. \quad (15)$$

Observe that I is a subgraph of J , so $N_I(A) \subseteq N_J(A)$. Let $B' = N_J(A) \setminus N_I(A)$. If $e \in B'$, then e must contain the external endpoint of some boundary edge of A . There are at most b such external endpoints, and each is incident to at most $d-1$ such non-boundary edges. Thus,

$$|B'| \leq b(d-1). \quad (16)$$

Let $A_0 \subseteq A$ be a maximal subset of vertices with pairwise-disjoint neighbourhoods in J . We will first determine how many choices there are for A_0 . We claim that

$$|A_0| \leq \frac{2}{d^2} (|B'| + |B| + |W|) \leq \frac{b}{d} \left(2 + \frac{1}{\eta} \right). \quad (17)$$

The first inequality holds because the above definitions give $|N_J(A)| = |B'| + |B| + |W|$, and each vertex in A has degree at least $d^2/2$ in J . The second inequality follows from (15) and (16).

The maximality of A_0 gives us that

$$\text{for any } u \in A \setminus A_0, \text{ there is a } v \in A_0 \text{ such that } N_J(u) \cap N_J(v) \neq \emptyset. \quad (18)$$

As a result, we have,

$$A_0 \cup N_J^2(A_0) \supseteq A. \quad (19)$$

Moreover, we make the following claim (where we recall $x \in A$ is some specified vertex).

Claim 15. *The set of vertices $A_0 \cup \{x\}$ is connected in G^7 .*




Figure 1. The three different cases for an edge e to be in $N_J(u_i) \cap N_J(w_i)$.




Figure 2. An example of part of the path P' in red arrows.

Proof. We start by proving that A_0 is connected in G^7 . Since A is connected in G , and $A_0 \subseteq A$, for any pair of distinct vertices $u, v \in A_0$, there must be a path P from u to v along edges of G using only vertices in A . Let the vertices of P be $w_0 = u, w_1, \dots, w_k = v$.

Observe that every vertex in P is distance at most three in G from A_0 . This follows from (18); indeed, either $w_i \in A_0$ or there is some $u_i \in A_0$ and $e \in G$ such that $e \in N_J(w_i) \cap N_J(u_i)$. This means e contains a vertex in $N_G(w_i)$ and a vertex in $N_G(u_i)$. There are three possibilities for e : either $e = u_i v_i$, e is incident to exactly one of u_i or w_i , or e is incident to neither u_i nor w_i . These cases give $\text{dist}_G(w_i, u_i)$ as 1, 2, and 3, respectively. See Fig. 1 for an illustration.

We can now construct a walk P' along edges of G from u to v that visits a vertex in A_0 at most every seven steps. Let Q_i be a path from w_i to A_0 such that $|Q_i| \leq 3$, and let Q_i^{-1} be the reverse path from A_0 to w_i . If $w_i \in A_0$, set Q_i and Q_i^{-1} to be empty. Then the edges of P' are

$$P' = (uw_1, Q_1, Q_1^{-1}, w_1w_2, Q_2, Q_2^{-1}, \dots, w_{k-2}w_{k-1}, Q_{k-1}, Q_{k-1}^{-1}, w_{k-1}w_k)$$

See Fig. 2 for an example. The walk P' visits vertices of A_0 at most every seven steps, so there is a walk in G^7 from u to v using only vertices in A_0 . As u and v were arbitrary, it follows that A_0 is connected in G^7 .

This completes the proof of the claim in the case $x \in A_0$. If $x \notin A_0$, then it is distance at most three in G from the nearest vertex in A_0 and hence is adjacent to a vertex of A_0 in G^7 . Thus, $A_0 \cup \{x\}$ is connected in G^7 as required. \square

As a result, we can specify A_0 by specifying a tree of degree at most d^7 rooted at v (since G^7 has maximum degree $d(d-1)^6 \leq d^7$). Using Proposition 14 and the bound on $|A_0|$ given by (17), there are at most $(ed^7)^{\frac{b}{d}(2+\frac{1}{\eta})} = d^{O((1+1/\eta)b/d)}$ possibilities for A_0 .

We now want to count the number of choices for B . Let E_0 be the edges of G in $N_J^3(A_0)$, meaning the edges which have an endpoint at distance at most 5 from A_0 , and let G' be the graph G with

every edge in $E \setminus E_0$ contracted. By (19) we have that

$$B \subseteq N_I(A) \subseteq N_J(A) \subseteq N_J^3(A_0),$$

so $B \subset E(G')$. We also have $v \in V(G')$ if and only if $\text{dist}_G(u, v) \leq 5$ for some $u \in A_0$. Since $|N_G^{k-1}(A_0)| \leq (d+1)|N_G^k(A_0)|$ for $k \geq 1$, this implies $|V(G')| \leq (d+1)^4|A_0| = O(bd^3)$.

By (8), the min-cut size of G is at least d (and in fact exactly d since G is d -regular). Contracting edges does not decrease the min-cut size of a graph, so the min-cut size of G' is also at least d .

Thus, the number of possible boundaries B is upper bounded by the number of cuts of size at most b in G' , which by Theorem 13 is at most

$$\binom{O(bd^3)}{\leq 2b/d} \cdot 2^{2b/d} = d^{O(b/d)}.$$

We are done, as we can uniquely determine A from its boundary. □

4.2 Colourings with small colour classes: Proof of Lemma 8

Let us now turn to Lemma 8. For a colouring σ , let us use $\text{nm}(\sigma) = \text{nm}(G, \sigma)$ to denote the number of non-monochromatic edges in σ ; observe that $m(G, \sigma) = |E(G)| - \text{nm}(\sigma)$. Recall that S_0 is the set of colourings in which each colour occupies at most $n/2$ vertices. The following lemma is the main point of this section.

Lemma 16. *Let G be as in Theorem 1. For $k \geq d$, let \mathcal{L}_k denote the number of q -colorings from S which give exactly k non-monochromatic edges. Then for $\delta > \frac{12}{\ln d}$,*

$$\mathcal{L}_k \leq n^4 \cdot q^{\frac{(2+\delta)k}{d} - \Omega\left(\frac{\delta k}{d}\right)}.$$

We would like to briefly comment on parameter δ . If one is interested in a combinatorial bound, one could simply plug in the ‘best possible’ value (i.e., $\delta = 12/\ln d$). We allow for a choice in the parameter δ solely to make the subsequent calculations slightly easier.

Before using this to prove Lemma 8, we will need the following.

Lemma 17. *Let G be an η -expander on n vertices. For every $\sigma \in S_0$ we have $\text{nm}(\sigma) \geq \eta n/2$.*

Proof. Let $A_i = \sigma^{-1}(i)$ be the set of vertices which get colour i under σ , which by the fact that $\sigma \in S_0$ satisfies $|A_i| \leq n/2$. The non-monochromatic edges of G under σ are precisely those that appear as $\nabla(A_i)$ for two distinct choices of $i \in [q]$. This gives

$$2 \text{ nm}(\sigma) = \sum_{i \in [q]} |\nabla(A_i)| \geq \eta \sum_{i \in [q]} |A_i| = \eta n.$$

□

We are now ready to prove Lemma 8.

Proof of Lemma 8. We start off with the easy lower bound $Z_G(q, \beta) \geq q \cdot e^{\beta|E(G)|}$ obtained by considering only the colourings that give every vertex the same colour. So for $\beta \geq (2 + \epsilon) \frac{\ln q}{d}$, we have

$$\begin{aligned}
 \frac{\sum_{\sigma \in S_0} e^{\beta(|E(G)| - \text{nm}(\sigma))}}{Z_G(q, \beta)} &\leq \frac{\sum_{k \geq \eta n/2} \sum_{\substack{\sigma \in S_0 \\ \text{nm}(\sigma) = k}} e^{\beta(|E(G)| - k)}}{q \cdot e^{\beta|E(G)|}} \\
 &\leq \frac{1}{q} \sum_{k \geq \eta n/2} \sum_{\substack{\sigma \in S_0 \\ \text{nm}(\sigma) = k}} q^{-(2+\epsilon)\frac{k}{d}} \\
 &\leq \frac{n^4}{q} \sum_{k \geq \eta n/2} q^{-\Omega\left(\frac{\epsilon k}{d}\right)}
 \end{aligned}$$

where the first inequality uses Lemma 17 and the last inequality uses Lemma 16 with $\delta = \epsilon$ (note that we already assume that $d \geq \exp(\Omega(1/\epsilon))$). If $d \geq \sqrt{n}$, then $\frac{n^4}{q} < 1$. Otherwise, $n^4 = q^{o\left(\frac{\epsilon k}{d}\right)}$. In either case,

$$\frac{n^4}{q} \sum_{k \geq \eta n/2} q^{-\Omega\left(\frac{\epsilon k}{d}\right)} \leq q^{-\Omega\left(\frac{\epsilon n}{d}\right)}.$$

□

One ingredient in the proof of Lemma 16 is the following, which requires the same techniques as in the proof of Theorem 13.

Lemma 18. *Let G be as in Theorem 1. Let $\mathcal{C}(\ell, s)$ be the number of q -colorings of the vertices of G such that there are exactly ℓ non-monochromatic edges, where each monochromatic component has at least s vertices. Then we have*

- (1) $\mathcal{C}(\ell, 1) \leq \binom{n}{2\ell/d} q^{\frac{2\ell}{d}}$, and
- (2) $\mathcal{C}(\ell, 2) \leq \binom{n}{2\ell/d} q^{\frac{2\ell}{2d-2}} \cdot (2d)^{\frac{2\ell}{d}}$.

Proof. For part 1: Fix any q -coloring σ with exactly ℓ non-monochromatic edges and run the following algorithm on G .

- I. While there are more than $\frac{2\ell}{d}$ vertices, choose a uniformly random edge and contract it. Delete self-loops, if any, after each contraction.
- II. Colour every vertex uniformly and independently from $[q]$. Output the final graph G' and colouring σ' .

Observe that σ' naturally corresponds to a q -coloring σ in the original graph defined by $\sigma(v) = \sigma'(v')$ where v' is the vertex in G' into which v was contracted.

Recall that G has a min-cut of size d by (8), and it does not decrease with the contraction operation during step I. So at any point in this step, if the current graph has m vertices, then it has at least $\frac{md}{2}$ edges.

Moreover, each contraction reduces the number of vertices by 1, and all deleted self-loops must be monochromatic. So the probability that the ℓ non-monochromatic edges remain uncontracted in G' is at least

$$\left(1 - \frac{2\ell}{nd}\right) \left(1 - \frac{2\ell}{(n-1)d}\right) \cdots \frac{d}{2\ell} = \binom{n}{2\ell/d}^{-1}.$$

During step II, the probability that every contracted vertex gets the correct colour is $q^{-\frac{2\ell}{d}}$. Therefore, the probability that σ was recovered by this procedure is at least $\binom{n}{2\ell/d}^{-1} q^{-\frac{2\ell}{d}}$. Since

this holds for any σ , there are at most $\binom{n}{2\ell/d} q^{\frac{2\ell}{d}}$ many q -colorings with ℓ non-monochromatic edges.

For part 2: Fix any q -coloring σ with at least ℓ non-monochromatic edges where each monochromatic component has at least 2 vertices, and consider the following algorithm.

- I. While there are more than $\frac{2\ell}{d}$ vertices, choose a uniformly random edge and contract it. Delete self-loops, if any, after each contraction.
- II. While a vertex in this contracted graph has degree d , contract it with one of its neighbours uniformly at random, deleting self-loops at each stage.
- III. While there are more than $\frac{2\ell}{2d-2}$ vertices, choose a uniformly random edge and contract it. Delete self-loops, if any, after each contraction.
- IV. Colour every vertex uniformly and independently at random from $[q]$.

Again, the final colouring naturally corresponds to a q -coloring in G . Step I is analysed in an identical manner as (1). The probability that the ℓ non-monochromatic edges remain uncontracted at the end of this step is at least $\binom{n}{2\ell/d}^{-1}$.

Since every monochromatic component in σ has at least two vertices, every uncontracted vertex at the end of step I is incident to at least one monochromatic edge. So during step II, the probability that the ℓ non-monochromatic edges remain uncontracted is at most $d^{-\frac{2\ell}{d}}$.

After the end of step II, (7) guarantees that the min-cut is at least $2d - 2$, and this does not decrease with the contraction operation in step III. So, using a calculation similar to that in the analysis of step I, the probability that the ℓ non-monochromatic edges remain uncontracted at the end of this step is at least $2^{-\frac{2\ell}{d}}$.

During step IV, the probability that every contracted vertex gets the correct colour is $q^{-\frac{2\ell}{2d-2}}$. Thus, σ is recovered with probability at least $\binom{n}{2\ell/d}^{-1} q^{-\frac{2\ell}{(2d-2)}} (2d)^{-\frac{2\ell}{d}}$. Since this holds for any σ , we have that there are at most $\binom{n}{2\ell/d} q^{\frac{2\ell}{(2d-2)}} (2d)^{\frac{2\ell}{d}}$ many q -colorings with ℓ non-monochromatic edges and each monochromatic component having size at least 2. \square

One can check by hand that our bound is not tight in several cases, such as when $s = 1$ and $\ell = d$. Indeed, the original algorithm by Karger for min-cuts gives a reasonable tight lower bound on the number of near-minimum cuts when G is a cycle – a case that is not included in the class of graphs under consideration for Theorem 1. However, recent work obtaining a more fine-grained understanding of the Karger process may be adaptable to our setting of q -colorings; we refer the interested reader to [15].

We are now ready to prove Lemma 16. Recall that \mathcal{L}_k is the number of q -colorings from S_0 which induce exactly k non-monochromatic edges, and we want to show that

$$\mathcal{L}_k \leq n^4 \cdot q^{\frac{(2+\delta)k}{d} - \Omega\left(\frac{\delta k}{d}\right)}.$$

Proof of Lemma 16. Case 1: $k \leq nd^{1-\frac{\delta}{4}}$. For a subset $A \subset V$, let $E(A)$ denote the set of edges induced by A . Consider a partition of the vertices of G into L monochromatic components. Let T be the set of components containing exactly one vertex. Let N be the set of non-monochromatic edges. We then have the partition

$$N = E(T) \sqcup \nabla(T) \sqcup \left(N \cap E(\overline{T})\right).$$

Let $e_T := |E(T)|$, $b_T := |\nabla(T)|$, and $\ell := |N \cap E(\overline{T})| = k - e_T - b_T$. We have that $|T| \leq \frac{2k}{d} \leq nd^{-\delta} < n/3$ and each monochromatic component has size at most $n/2$. So, one can choose some

$S \subseteq V$ such that $n/3 \leq |S| \leq 2n/3$ by greedily including all vertices from the smallest remaining monochromatic component. In particular, $S \supseteq T$. The expansion condition (3) then implies that $|\nabla(S)| \geq n/3$. However, $\nabla(S) \subseteq \nabla(T) \cup (N \cap E(\overline{T}))$, and so $|\nabla(S)| \leq b_T + (k - e_T - b_T) = k - e_T = \ell + b_T$. As a result, we have

$$b_T + \ell = \Omega(n). \quad (20)$$

A q -coloring counted by \mathcal{L}_k can be chosen by (I) first choosing the cut (T, \overline{T}) , (II) giving each element on the T -side of the cut a colour, and (III) choosing the monochromatic components in $V \setminus T$ so that each component has at least two vertices.

Theorem 13 gives us that (I) can be done in at most $\binom{n}{2b_T/d} \cdot 2^{2b_T/d}$ ways. We also have that (II) can be done in at most $q^{|T|}$ ways. Part 2 of Lemma 18 gives us that (III) can be done in at most $\binom{n}{2\ell/d} \cdot q^{\frac{2\ell}{2d-2}} \cdot (2d)^{\frac{2\ell}{d}}$ ways. So, the number of q -colorings of V with the above parameters is at most

$$\begin{aligned} Q &:= \binom{n}{2b_T/d} \cdot 2^{2b_T/d} \cdot q^{|T|} \cdot \binom{n}{2\ell/d} \cdot q^{\frac{2\ell}{2d-2}} \cdot (2d)^{\frac{2\ell}{d}} \\ &= \binom{n}{2b_T/d} \cdot \binom{n}{2\ell/d} \cdot q^{\frac{\ell+2e_T+b_T}{d}} \cdot q^{\frac{\ell}{d^2}} \cdot (2d)^{\frac{2\ell}{d}}, \end{aligned}$$

where we have used $d|T| = 2e_T + b_T$. So we have

$$\begin{aligned} Q \cdot q^{-\frac{2k}{d}} &= \binom{n}{2\ell/d} \binom{n}{2b_T/d} \cdot q^{\frac{-b_T-\ell}{d}} \cdot q^{\frac{\ell}{d^2}} \cdot (2d)^{\frac{2\ell}{d}} \\ &\leq \binom{n}{2(\ell+b_T)/d} \cdot 2^{\frac{4(b_T+\ell)}{d}} q^{\frac{-b_T-\ell}{d}} \cdot q^{\frac{\ell}{d^2}} \cdot (2d)^{\frac{2\ell}{d}} \\ &\leq d^{O(\frac{\ell+b_T}{d})} q^{-\frac{\ell+b_T}{d}} \cdot 2^{\frac{4(b_T+\ell)}{d}} q^{\frac{\ell}{d^2}} \\ &\leq q^{-\Omega(\frac{n}{d})}, \end{aligned}$$

where the first inequality used the fact that for $a + b \leq x$, we have $\binom{x}{a} \binom{x}{b} \leq \binom{x}{a+b} 4^{a+b}$, and the next two inequalities use (20). So summing over all values for ℓ and b_T gives us that for $k \leq nd^{1-\frac{\delta}{4}}$,

$$\mathcal{L}_k \leq n^4 \cdot Q \leq n^4 \cdot q^{\frac{2k}{d} - \Omega(\frac{n}{d})}.$$

Case 2: $k \geq nd^{1-\frac{\delta}{4}}$. This case is relatively straightforward. Using part 1 of Lemma 18, we have

$$\mathcal{L}_k \leq \binom{n}{2k/d} q^{\frac{2k}{d}} \leq (ed)^{\frac{2\delta k}{3d}} \cdot q^{\frac{2k}{d}} = q^{-\frac{(2+\delta)k}{d} - \Omega(\frac{\delta k}{d})}.$$

□

5. Extremal results for partition functions

An important step in upper bounding the weight $w_\gamma = e^{-\beta(\nabla_\gamma + e_\gamma)} Z_\gamma(q-1, \beta)$ is to upper bound the partition function $Z_\gamma(q-1, \beta)$. Bounds on partition functions in this setting are a well-studied topic in extremal combinatorics.

5.1 General graphs and cliques

Theorem 19 (Sah et al. [40], Theorem 1.14 restated). *Let G be a graph, $q \geq 2$ be an integer and $\beta \geq 0$. Then*

$$Z_G(q, \beta) \leq \prod_{v \in V(G)} Z_{K_{d_v+1}}(q, \beta)^{\frac{1}{d_v+1}},$$

where d_v is the degree of a vertex $v \in V(G)$.

We briefly remark that their result was much more general; [40] shows that in fact every ferromagnetic model (using a definition introduced in [18] as a generalisation of the ferromagnetic Potts model) is ‘clique-maximizing’, which translates to the above statement on the level of partition functions.

Corollary 20. *Let G be a graph of maximum degree Δ with n vertices and m edges. Then for any integer $q \geq 2$ and $\beta \geq 0$,*

$$Z_G(q, \beta) \leq q^{n - \frac{2m}{\Delta}} Z_{K_{\Delta+1}}(q, \beta)^{\frac{2m}{\Delta(\Delta+1)}}.$$

Proof. This follows from Theorem 19 and the fact that $Z_{K_{d+1}}(q, \beta)^{\frac{1}{d+1}}$ forms a log-convex sequence for $d \geq 0$, which is proved in ([40], Lemma 5.3). In particular, consider the graph with partition function $Z_G(q, \beta)^\Delta$ formed from the disjoint union of Δ copies of G . Applying Theorem 19, we obtain an upper bound on $Z_G(q, \beta)^\Delta$ in terms of a degree sequence. We can further increase the upper bound because the aforementioned log-convexity implies that when $a \leq b$, replacing two degrees a and b by degrees $a - 1$ and $b + 1$ can only increase the bound. Then over graphs of maximum degree Δ with Δn vertices and Δm edges, the maximum upper bound Theorem 19 can give is from a degree sequence where $2m$ vertices have degree Δ and the rest have degree zero. The result follows upon taking the power $1/\Delta$ of this inequality. \square

5.2 Triangle-free graphs and bicliques

The upper bound in Theorem 19 can be improved in the case that G is triangle-free.

Theorem 21 (Sah et al. [40], Theorem 1.9 and the subsequent remark). *Let G be a triangle-free graph with no isolated vertices, $q \geq 2$ be an integer and $\beta \geq 0$. Then*

$$Z_G(q, \beta) \leq \prod_{uv \in E(G)} Z_{K_{d_u, d_v}}(q, \beta)^{\frac{1}{d_u d_v}},$$

where d_u is the degree of a vertex $u \in V(G)$.

Corollary 22. *Let G be a triangle-free graph of maximum degree Δ with n vertices and m edges. Then for any integer $q \geq 2$ and $\beta \geq 0$,*

$$Z_G(q, \beta) \leq q^{n - \frac{2m}{\Delta}} Z_{K_{\Delta, \Delta}}(q, \beta)^{\frac{m}{\Delta^2}}.$$

Proof. Let G have t isolated vertices. Then by Theorem 21 we have

$$Z_G(q, \beta) \leq q^t \prod_{uv \in E(G)} Z_{K_{d_u, d_v}}(q, \beta)^{\frac{1}{d_u d_v}}.$$

The desired result is now a consequence of the fact that for $1 \leq a \leq c$ and $1 \leq b \leq d$ we have

$$q^{-\left(\frac{1}{a} + \frac{1}{b}\right)} Z_{K_{a,b}}(q, \beta)^{\frac{1}{ab}} \leq q^{-\left(\frac{1}{c} + \frac{1}{d}\right)} Z_{K_{c,d}}(q, \beta)^{\frac{1}{cd}}. \quad (21)$$

To see this, apply the inequality to each term of the product and collect the factors of q to obtain

$$Z_G(q, \beta) \leq q^{t + \sum_{uv \in E(G)} \left(\frac{1}{d_u} + \frac{1}{d_v} \right) - \frac{2m}{\Delta}} Z_{K_{\Delta, \Delta}}(q, \beta)^{\frac{m}{\Delta^2}}.$$

Some simple counting gives

$$\sum_{uv \in E(G)} \left(\frac{1}{d_u} + \frac{1}{d_v} \right) = n - t$$

since each vertex u appears as the endpoint of precisely d_u edges, and the result follows.

It remains to prove inequality (21), which follows from Hölder's inequality and Jensen's inequality as noted² in [40]. For concreteness, we observe that inequality (21) is a direct consequence of a generalised Hölder inequality stated as Theorem 3.1 in [32]. We can write $\Omega = [q]$, and let μ be the uniform probability measure on Ω and μ^a be the uniform probability measure on Ω^a such that for the non-negative functions $f: \Omega^2 \rightarrow \mathbb{R}$ and $g: \Omega^a \rightarrow \mathbb{R}$ given by

$$f(x, y) = \begin{cases} e^\beta & x = y \\ 1 & \text{otherwise,} \end{cases}$$

$$g(x_1, \dots, x_a) = \int_{\Omega} \prod_{i=1}^a f(x_i, y) \, d\mu.$$

Then to prove (21) in this notation, we have

$$\text{LHS} := q^{-\left(\frac{1}{a} + \frac{1}{b}\right)} Z_{K_{a,b}}(q, \beta)^{\frac{1}{ab}} = \left(\int_{[q]^a} |g|^b \, d\mu^a \right)^{\frac{1}{ab}}.$$

When $b \leq d$ we can apply the generalised Hölder inequality (or in this somewhat special case iterated applications of the usual Hölder inequality followed by Jensen's inequality) to the integral over $[q]^a$ above to obtain the upper bound

$$\text{LHS} \leq \left(\int_{\Omega^a} g^d \, d\mu^a \right)^{\frac{1}{ad}} = q^{-\left(\frac{1}{a} + \frac{1}{d}\right)} Z_{K_{a,d}}(q, \beta)^{\frac{1}{ad}}.$$

A symmetric application of this argument for any integer c such that $a \leq c$ gives inequality (21). \square

Lemma 23. Let $q \geq d^4$ and for some fixed $\epsilon > 0$ let $\beta \geq (1 + \epsilon)\beta_0$. Then,

- (1) $Z_{K_{d+1}}(q, \beta) \leq (1 + q^{-\Omega(\epsilon)}) q e^{\beta \binom{d+1}{2}}$, and
- (2) $Z_{K_{d,d}}(q, \beta) \leq (1 + q^{-\Omega(\epsilon)}) q e^{\beta d^2}$.

Proof. We will use the fact that both K_{d+1} and $K_{d,d}$ are η -expanders with $\eta = d/2$. We will state the proof for K_{d+1} , and proof for $K_{d,d}$ follows similarly. Let σ be any colouring of the vertices of K_d , and let $c(\sigma)$ denote the number of colours that appear in σ .

We first observe that the proof of Lemma 16 also extends to all σ such that $c(\sigma) \geq 2$. Indeed, the only place where S was used was to establish Lemma 17, and (20). Both of these also hold for

²See Definition 1.11 and the subsequent discussion of monotonicity in [40]. Note that for this monotonicity to hold one must work with probability measures on the sets of spins, which is not necessarily the case for other 'scale-free' results proved in [40] and related works.

K_{d+1} whenever there are at least two colours. So we have

$$\begin{aligned} \frac{\sum_{\sigma: c(\sigma) \geq 2} e^{\beta \left(\binom{d+1}{2} - \text{nm}(\sigma) \right)}}{q \cdot e^{\beta \binom{d+1}{2}}} &\leq \frac{1}{q} \sum_{k \geq d} \sum_{\substack{\sigma: c(\sigma) \geq 2 \\ \text{nm}(\sigma) = k}} q^{-(2+\epsilon) \frac{k}{d}} \\ &\leq \frac{d^4}{q} \sum_{k \geq d} q^{-\Omega\left(\frac{\epsilon k}{d}\right)} \\ &\leq q^{-\Omega(\epsilon)}. \end{aligned}$$

□

The proofs of Lemmas 11 and 12 are now straightforward calculations.

Proof of Lemma 11. By Corollary 20 and Lemma 23,

$$\begin{aligned} Z_{G[\gamma]}(q-1, \beta) &\leq (q-1)^{v_\gamma - \frac{2e_\gamma}{d}} Z_{K_{d+1}}(q-1, \beta)^{\frac{2e_\gamma}{d(d+1)}} \\ &\leq (q-1)^{v_\gamma - \frac{2e_\gamma}{d}} \left(2(q-1) e^{\beta \binom{d+1}{2}} \right)^{\frac{2e_\gamma}{d(d+1)}} \\ &< q^{v_\gamma - \frac{2e_\gamma}{d} + \frac{e_\gamma}{\binom{d+1}{2}}} e^{\beta e_\gamma 2^{\frac{e_\gamma}{\binom{d+1}{2}}}}. \end{aligned}$$

□

Proof of Lemma 12. By Corollary 22 and Lemma 23,

$$\begin{aligned} Z_{G[\gamma]}(q-1, \beta) &\leq (q-1)^{v_\gamma - \frac{2e_\gamma}{d}} Z_{K_{d,d}}(q-1, \beta)^{\frac{e_\gamma}{d^2}} \\ &\leq (q-1)^{v_\gamma - \frac{2e_\gamma}{d}} \cdot (2(q-1) e^{\beta d^2})^{\frac{e_\gamma}{d^2}} \\ &\leq q^{v_\gamma - \frac{2e_\gamma}{d} + \frac{e_\gamma}{d^2}} \cdot 2^{\frac{e_\gamma}{d^2}} \cdot e^{\beta e_\gamma}. \end{aligned}$$

□

6. High-temperature algorithms

The polymer model we use in the high-temperature region is the same as the one considered in [2, 8, 21], which we now define. Given a graph $G = (V, E)$, the set \mathcal{P} of polymers is the set of connected subgraphs (not necessarily induced) of G on at least two vertices. The polymer weights are given by $w_\gamma = q^{1-v_\gamma} p^{e_\gamma}$, where we write $p = e^\beta - 1$. Two polymers γ and γ' are incompatible if and only if their union is not connected in G .

The key observation is that edge subsets $A \subset E$ are in bijection with the collection Ω of sets of pairwise compatible polymers by the map taking A to the set of components of (V, A) on at least two vertices. In this setup, the polymer model partition function Ξ is given by

$$\begin{aligned}
\Xi &= \sum_{\Lambda \in \Omega} \prod_{\gamma \in \Lambda} w_{\gamma} \\
&= \sum_{A \subseteq E} \prod_{\substack{\text{components } \gamma \\ \text{of } (V, A)}} q^{1-v_{\gamma}} p^{e_{\gamma}} \\
&= \sum_{A \subseteq E} q^{c(A)-n} (e^{\beta} - 1)^{|A|} = q^{-n} Z_G(q, \beta),
\end{aligned}$$

where $c(A)$ is the number of components of the graph (V, A) . This shows that $q^n \Xi$ is precisely the function which we wish to approximate, and hence any multiplicative approximation to Ξ corresponds to a multiplicative approximation to Z with the same relative error. Given the standard approach to approximating polymer model partition functions outlined in Section 2.2, to complete our high-temperature algorithm it suffices to verify the conditions of Theorem 6 which imply that the cluster expansion of this polymer model is convergent.

Lemma 24. *For every $\epsilon \in (0, 1)$ there exist $d_0(\epsilon)$ and $q_0(\epsilon)$ such that for every $d \geq d_0$, $q \geq q_0$ and $\beta \leq (1 - \epsilon)\beta_o(q, d)$, the Kotecký–Preiss condition (see (5) in Theorem 6) holds for the high-temperature polymer model with $f(\gamma) = g(\gamma) = v_{\gamma}$.*

Proof. We assume that $d_0(\epsilon)$ and $q_0(\epsilon)$ are large enough that $\beta \leq (1 - \epsilon)\beta_o$ implies that $\beta \leq (2 - \epsilon) \ln(q)/d$ and hence $q \geq (1 + p)^{d/(2-\epsilon)}$.

Write $\mathcal{P}_{u,k,t}$ for the subset of polymers $\gamma \in \mathcal{P}$ such that γ contains a fixed vertex u , $v_{\gamma} = k$, and $e_{\gamma} = t$. Given an arbitrary polymer γ , we bound from above the sum in the left-hand side of (5) as follows

$$\sum_{\gamma' \not\supseteq \gamma} w_{\gamma'} e^{f(\gamma') + g(\gamma')} \leq \sum_{u \in \gamma} \sum_{\gamma' \supseteq u} w_{\gamma'} e^{f(\gamma') + g(\gamma')}.$$

Given this and our choice $f(\gamma) = v_{\gamma}$, it suffices to show that for every vertex u in G ,

$$F_u := \sum_{\gamma' \supseteq u} w_{\gamma'} e^{f(\gamma') + g(\gamma')} \leq 1.$$

Observe that

$$F_u = q \sum_{\gamma' \supseteq u} q^{-v_{\gamma'}} p^{e_{\gamma'}} e^{f(\gamma') + g(\gamma')} = q \sum_{k \geq 2} \sum_{t \geq k-1} |\mathcal{P}_{u,k,t}| q^{-k} p^t e^{2k}.$$

For $t < k - 1$ we have $|\mathcal{P}_{u,k,t}| = 0$ since polymers must be connected, hence the above sums over t start at $k - 1$. We bound $|\mathcal{P}_{u,k,t}|$ from above to continue, noting that

$$|\mathcal{P}_{u,k,t}| \leq \min \left\{ (ed)^{k-1} \cdot \binom{k}{t-k+1}, (ed)^{k-1} \cdot \binom{dk/2-k+1}{t-k+1} \right\}.$$

To see this, observe that there are at most $(ed)^{k-1}$ choices of a spanning tree of $\gamma \in \mathcal{P}_{u,k,t}$ by Proposition 14, and the binomial coefficient gives an upper bound on the number of ways of adding edges of G to this spanning tree to form γ . In the case $k \leq d + 1$, we can use $\binom{k}{2}$ as an upper bound on the number of edges of G incident to a vertex of this spanning tree; else, we use $dk/2$. Then

$$F_u \leq \frac{q}{ed} \sum_{k=2}^d (e^3 d/q)^k p^{k-1} \sum_{t=k-1}^{\binom{k}{2}} p^{t-k+1} \binom{k}{t-k+1} \\ + \frac{q}{ed} \sum_{k=d+1}^{\infty} (e^3 d/q)^k p^{k-1} \sum_{t=k-1}^{dk/2} p^{t-k+1} \binom{dk/2-k+1}{t-k+1}.$$

The inner sums over t are precisely $(1+p)^{\binom{k}{2}-k+1}$ and $(1+p)^{dk/2-k+1}$ respectively, so that

$$F_u \leq \frac{1+p}{p} \frac{q}{ed} \left(\sum_{k=2}^d (e^3 dp/q)^k (1+p)^{\binom{k}{2}-k} + \sum_{k=d+1}^{\infty} (e^3 dp/q)^k (1+p)^{dk/2-k} \right).$$

Let $r = e^3 dp(1+p)^{d/2-1}/q$, and observe that for $k \in [2, d]$ we have $\binom{k}{2} - k \leq \frac{dk}{2} - k - \frac{d}{2}$. Then

$$F_u \leq \frac{1+p}{p} \frac{q}{ed} \left(\sum_{k=2}^d \frac{r^k}{(1+p)^{d/2}} + \sum_{k=d+1}^{\infty} r^k \right) = \frac{e^2}{1-r} \left(r - r^d + (1+p)^{d/2} r^d \right).$$

To finish the proof, observe that in order to show $F_u \leq 1$ it suffices to show that $(1+p)^{1/2} r \leq 1/100$.

Substituting the definition of r and $q \geq (1+p)^{d/(2-\epsilon)}$ into $(1+p)^{1/2} r$ we have

$$(1+p)^{1/2} r \leq e^3 dp(1+p)^{\frac{d}{2} - \frac{d}{2-\epsilon} - \frac{1}{2}}.$$

This has a unique stationary point on $p \in [0, \infty)$ at $p^* = \frac{4-2\epsilon}{(d+1)\epsilon-2}$ provided that $d \geq 10/\epsilon$ (which also ensures that $p^* > 0$), and this stationary point is a maximum. Assuming that $d \geq 10/\epsilon$, one can verify that $(1+p)^{1/2} r < 1/100$ whenever $p \geq 100p^*/\epsilon^2$. In terms of q , to ensure that $p \geq 100p^*/\epsilon^2$, it suffices to take $q > \exp(400/\epsilon^3)$. This establishes Lemma 24. \square

We briefly note that the above proof requires $\epsilon > 0$ so we cannot, for instance, simply allow $\epsilon < 0$ in order to extend Theorem 1 to all temperatures. However, we can extend our results if we modify our high-temperature polymer model to consider only a subset of polymers – namely, those at most a certain size. We discuss this approach in slightly more detail in Section 9.

We also note that Theorem 2 follows from Lemma 24 and the standard method sketched in Subsection 2.2 which uses Theorem 7.

Proof of Theorem 2. As both q and d tend to infinity, there exists $\beta_u(q, d) \sim \ln(q)/d$ such that the Gibbs uniqueness region of the infinite d -regular tree corresponds to $\beta < \beta_u$ (see [20, 6]). Since $\beta_o \sim 2 \ln(q)/d$ we can pick (say) $\epsilon = 1/4$ in Lemma 24 and so long as d and q are large enough (i.e. at least some absolute constants) we will have the conditions of that lemma, as well as $\beta_u \leq (1-\epsilon)\beta_o$. The algorithm is given in Subsection 2.2. \square

7. Typical structure of the Potts model

In this section, we give a proof of Theorem 4. A useful perspective on the random cluster model comes from tilted bond percolation, providing a distribution on edge subsets of a graph that with $q = 1$ is precisely the usual notion of (independent) bond percolation. For integer q , there is a coupling between the distributions given by the Potts and random cluster models due to Edwards and Sokal [12] which we now describe. Given a subset of edges from the random cluster model, colour each component with a colour uniformly from $[q]$ to obtain a vertex colouring distributed according to the Potts model. In reverse, one performs bond percolation (with a particular probability) on the monochromatic components of a colouring from the Potts model. See e.g. [38] for

a thorough treatment of the function Z and the various combinatorial quantities that it encodes, including the equivalence of (1) and (4).

Proof of Theorem 4. For part 1, we use the Kotecký–Preiss theorem for the high-temperature cluster expansion to argue that all polymers have size $O(\ln n)$ with high probability. In order to do this, let $L = (1 + \delta) \ln n$ and consider the truncated cluster expansion for Ξ . Recall that

$$\Xi(L) = \exp \left(\sum_{\substack{\Gamma \in \mathcal{C} \\ g(\Gamma) \leq L}} \phi(\Gamma) \prod_{\gamma \in \Gamma} w_{\gamma} \right).$$

We have that $|\ln \Xi(L) - \ln \Xi| \leq n^{-\delta}$. Therefore, if we let \tilde{Z}_G be the sum over configurations with components of size at most L , we have

$$\frac{\tilde{Z}_G}{Z_G} = \frac{q^n \Xi(L)}{q^n \Xi} \geq e^{-n^{-\delta}} = 1 - \Omega(n^{-\delta}),$$

and so, with high probability all components are on at most $(1 + \delta) \ln n$ vertices. To get a configuration of the Potts model we make use of the Edwards–Sokal coupling [12] by colouring each component at random with one of the q colours. Note that if C_1, \dots, C_K are the components we have

$$\sum_{i=1}^K |C_i|^2 \leq \max_{i \in \{1, \dots, K\}} |C_i| \cdot \sum_{i=1}^K |C_i| \leq (1 + \delta) n \ln n.$$

Let N_j be the number of vertices with colour j . Clearly, $\mathbb{E}N_j = n/q$. By Hoeffding's inequality ([22], Theorem 2), we have that

$$\mathbb{P} \left(\left| N_j - \frac{n}{q} \right| > tn \right) \leq 2 \exp \left(\frac{-2t^2 n^2}{\sum_{i=1}^K |C_i|^2} \right) \leq 2 \exp \left(-\frac{2t^2 n}{(1 + \delta) \ln n} \right).$$

Taking $t = o\left(\frac{1}{q}\right) \cap \omega\left(\sqrt{\frac{\ln n \ln q}{n}}\right)$ and a union bound over all colours q , we have that with high probability all the N_j satisfy $N_j = (1 + o(1))n/q$.

The strengthening of Part 2 now follows. Replacing $n/2$ with δn in (14), we have that

$$\begin{aligned} \mathbb{P}(\|\mathbf{A}\| \geq (1 - \delta)n) &= \mathbb{P}(e^{\|\mathbf{A}\|/d} \geq e^{(1-\delta)n/d}) \\ &\leq \exp \left(\frac{\epsilon n \ln q}{4q^{\epsilon/4} d} - (1 - \delta) \frac{n}{d} \right) \end{aligned}$$

This gives us that for some $\delta = O\left(\frac{\epsilon \ln q}{q^{\epsilon/4}}\right)$ we have

$$\text{the largest color class has at least } n(1 - \delta) \text{ vertices w.h.p.} \quad (22)$$

□

Note that there is some subtlety in the final argument. We did not start out by defining the ‘majority colorings’ (those not in S_0) to be the ones with at least $(1 - o(1))n$ vertices of the same colour; if we had, the combinatorial arguments necessary to prove the Kotecký–Preiss condition would have been far more difficult (or in some cases false) due to the weakness of our expansion assumptions. Instead, we take a two-layered approach. In the first layer, we use combinatorial arguments to show that $Z_G(q, \beta)$ is exponentially well-approximated by the majority colourings, where ‘majority’ here means at least $\frac{n}{2}$ vertices of the same colour (Lemma 8). In the second layer, we show that this approximation is itself well-approximated by a polymer model with convergent cluster expansion (Lemma 9). Due to the convergence (Lemma 10), we can discern an even finer

level of detail within the polymer model – that the polymer model partition function is itself well-approximated by majority colourings, where we now take ‘majority’ to mean at least $(1 - o(1))n$ vertices of the same colour. Our approximations at each step are good enough that we can chain them together to achieve an approximation of our original $Z_G(q, \beta)$.

8. The random cluster model on the discrete hypercube

In this section, we prove Theorem 5 about the existence of a structural phase transition for the random cluster model on the discrete hypercube.

Proof of Theorem 5. The proof of part 1 follows from the proof of part 1 of Theorem 4, which shows that with high probability, each component in the random cluster model has size at most $(1 + o(1)) \ln n \leq \ln_2 n \leq d$.

For part 2, suppose there are K components with vertex sets C_1, \dots, C_K , where $|C_1| \geq \dots \geq |C_K|$. We make use of the Edwards–Sokal coupling [12]. For a subset of edges distributed according to the random cluster model, colour each component at random with one of the q colours to get a configuration of the Potts model.

By (22), every component with more than δn vertices should get the same colour. If there are at least two such components, then the probability of this occurring is at most $1/q \ll 1$, thus contradicting (22).

Henceforth, let us assume that there is at most one component of size at least δn , that is, $\max_{i \in \{2, \dots, K\}} |C_i| \leq \delta n$. Let $n_1 := n - |C_1|$. We have

$$\sum_{i=2}^K |C_i|^2 \leq \max_{i \in \{2, \dots, K\}} |C_i| \cdot \sum_{i=2}^K |C_i| = \delta n \cdot n_1.$$

Let N_j be the number of vertices in components C_2, \dots, C_K with colour j . Clearly, $\mathbb{E}N_j = \frac{n_1}{q}$. Set $t = \sqrt{\delta} \cdot \ln q$. By Hoeffding’s inequality, we have that

$$\mathbb{P} \left(N_j - \frac{n_1}{q} > tn_1 \right) \leq \exp \left(\frac{-2t^2 n_1^2}{\sum_{i=2}^K |C_i|^2} \right) \leq \exp \left(-\frac{2t^2 n_1}{\delta n} \right).$$

If $n_1 \leq n/\ln q$, then $|C_1| \geq n(1 - 1/\ln q)$ and we are done. Else, we have that $\frac{2t^2 n_1}{\delta n} \geq 2 \ln q$, and so $\mathbb{P}(N_j > (t + 1/q)n_1) \leq e^{-2 \ln q} = 1/q^2$. By taking a union bound over all colours q , we have that with high probability all of the N_j are smaller than $(t + 1/q)n_1$, and so with high probability, the largest colour class has at most

$$|C_1| + (t + 1/q)n_1 \leq n \left(1 - \Omega \left(\frac{1}{\ln q} \right) \right)$$

vertices, contradicting (22). □

9. Conclusion

We conclude with a discussion on future directions and potential extensions of our arguments.

9.1 Extension to all temperatures

A deceptively simple-sounding task would be to extend our algorithms to work at all temperatures, meaning the range $(1 - \epsilon)\beta_0 \leq \beta \leq (1 + \epsilon)\beta_0$. This requires convergence of both our low- and high-temperature cluster expansions beyond the threshold β_0 . The main results of [21] show that

this must be possible for large enough q and strong enough expansion conditions, so the main interest here is to close the gap without significantly strengthening our assumptions.

For the high-temperature regime, we can take the same polymer model restricted to polymers on at most $\frac{n}{2}$ vertices and repeat the argument from Section 6. Subject to the condition of η -expansion, the constraint in the polymer size allows for better bounds on the size of $\mathcal{P}_{u,k,t}$ (the k -vertex t -edge polymers contained a fixed vertex u), and ultimately lets us establish the Kotecký–Preiss condition for this ‘smaller’ polymer model at temperatures slightly beyond β_0 . There is a delicate interaction between the expansion and improved range of β , however, and one must handle the colourings that are no longer represented due to the restriction on polymer size.

More significant obstacles appear in extending our low-temperature results. In this regime, our current methods do not extend in a straightforward way beyond β_0 unless we allow q to be exponential in d , thus losing some novelty in our results. That is, some extension of what we prove here is plausible, but our conditions seem to degrade and become comparable to the setup of [21], while still requiring q to be an integer. One possible approach to maintaining q polynomial in d is to improve Lemmas 11 and 12 by providing a better analysis of the partition functions of K_{d+1} and $K_{d,d}$. Even a simple prerequisite such as more a comprehensive version of Lemma 23 that operates in the entire temperature range, appears not to be known. We leave the pursuit of such bounds to future work.

9.2 Weaker expansion, fewer colours

A natural further direction is to extend our results to a weaker notion of expansion that only guarantees large edge-boundary for small sets, e.g. for some $k \geq 2$ we have

$$\text{Every set } A \text{ of at most } n/k \text{ vertices satisfies } |\nabla(A)| \geq 2|A|.$$

Such a condition is related to the gap between eigenvalues k and $k+1$ of the graph, and it would be interesting to determine whether techniques from combinatorial optimisation that handle small-set expansion and such higher-order eigenvalue gaps can be used for approximating partition functions. This was initiated in [7] for the low-temperature polymer model that we study here.

We believe that the lower bounds on q (and to some extent d) in our results are artefacts of our techniques. Extending our results to the case of all d , $q \geq 3$ presents an interesting challenge, as raised previously in [21]. The slow mixing results of [10] provide some evidence that this may be possible, at least at low temperatures.

9.3 Conjectured extension to the random cluster model

It would also be interesting to remove the restriction that q is an integer and handle the random cluster model directly. That is, are there aspects of our approach that can be extended to work with the more general polymer models of [21]?

Our high-temperature algorithm already works in the extra generality of the random cluster model, but some of the tools we rely on for the low-temperature algorithm only apply to the Potts model. We conjecture that the results we rely on in Section 5 hold in the required generality.

Conjecture 25. *For all real $q > 1$ and $\beta > 0$, for any graph G we have*

$$Z_G(q, \beta) \leq \prod_{v \in V(G)} Z_{K_{d_v+1}}(q, \beta)^{1/(d_v+1)}.$$

Conjecture 26. *For all real $q > 1$ and $\beta > 0$, for any triangle-free graph G we have*

$$Z_G(q, \beta) \leq \prod_{uv \in E(G)} Z_{K_{d_u, d_v}}(q, \beta)^{1/(d_u d_v)}.$$

Extending our low-temperature algorithm to the random cluster model may not require results as strong as these conjectures, and instead one could work with analogues of Corollaries 20 and 22 that offer bounds in terms of the total number of edges and maximum degree parameter. So pursuing such slightly weaker results is also of interest. Note that $q \geq 1$ is required in these conjectures; they do not hold for $q < 1$ and at $q = 1$ we have $Z_G(q, \beta) = e^{\beta e(G)}$ so in this degenerate case the conjectures are trivial.

Financial support

NF is supported in part by NSF grant CCF-1934964, AK is supported by NSF CAREER grant 1452923, AP is supported in part by NSF grant CCF-1934915, CY is supported in part by NSF grant CCF-1814409 and NSF grant DMS-1800521. An extended abstract of this work appeared at FOCS 2022.

Acknowledgments

This work was undertaken as part of the Phase Transitions and Algorithms working group in the SAMSI Spring 2021 semester programme on Combinatorial Probability. We thank the semester organisers for bringing us together. We also thank Will Perkins and the anonymous referees for helpful feedback on this paper.

References

- [1] Barvinok, A. (2016) *Combinatorics and Complexity of Partition Functions*, Vol. 30 of *Algorithms and Combinatorics*. Springer International Publishing, Cham.
- [2] Borgs, C., Chayes, J., Helmuth, T., Perkins, W. and Tetali, P. (2020) Efficient sampling and counting algorithms for the Potts model on \mathbb{Z}^d at all temperatures. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, pages 738–751, Chicago, IL, USA, June 2020. Association for Computing Machinery.
- [3] Blanca, A., Cannon, S. and Perkins, W. (2022) Fast and perfect sampling of subgraphs and polymer systems.
- [4] Blanca, A. and Gheissari, R. (2021) Random-cluster dynamics on random regular graphs in tree uniqueness. *Commun. Math. Phys.* **386**(2) 1243–1287.
- [5] Blanca, A., Galanis, A., Goldberg, L. A., Štefankovič, D., Vigoda, E. and Yang, K. (2020) Sampling in uniqueness from the Potts and random-cluster models on random regular graphs. *SIAM J. Discrete Math.* **34**(1) 742–793.
- [6] Bordewich, M., Greenhill, C. and Patel, V. (2016) Mixing of the Glauber dynamics for the ferromagnetic Potts model. *Random Struct. Algor.* **48**(1) 21–52.
- [7] Carlson, C., Davies, E. and Kolla, A. (2020) Efficient algorithms for the Potts model on small-set expanders.
- [8] Coulson, M., Davies, E., Kolla, A., Patel, V. and Regts, G. (2020) Statistical physics approaches to Unique Games. In *35th Computational Complexity Conference (CCC 2020)*, Vol. 169 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Dagstuhl, Germany. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, pp. 13:1–13:27.
- [9] Chen, Z., Galanis, A., Goldberg, L. A., Perkins, W., Stewart, J. and Vigoda, E. (2021) Fast algorithms at low temperatures via Markov chains. *Random Struct. Algor.* **58**(2) 294–321.
- [10] Coja-Oghlan, A., Galanis, A., Goldberg, L. A., Ravelomanana, J. B., Stefankovic, D. and Vigoda, E. (2022) Metastability of the Potts ferromagnet on random regular graphs.
- [11] Dyer, M., Goldberg, L. A., Greenhill, C. and Jerrum, M. (2004) The relative complexity of approximate counting problems. *Algorithmica* **38**(3) 471–500.
- [12] Edwards, R. G. and Sokal, A. D. (1988) Generalization of the Fortuin-Kasteleyn-Swendsen-Wang representation and monte carlo algorithm. *Phys. Rev. D* **38**(6) 2009–2012.
- [13] Fortuin, C. M. and Kasteleyn, P. W. (1972) On the random-cluster model: I. Introduction and relation to other models. *Physica* **57**(4) 536–564.
- [14] Galanis, A., Goldberg, L. A. and Stewart, J. (2021) Fast algorithms for general spin systems on bipartite expanders. *ACM Trans. Comput. Theory* **13**(4) 25:1–25:18.
- [15] Gupta, A., Harris, D. G., Lee, E. and Li, J. (2021) Optimal bounds for the k -cut problem. *ACM J. ACM (IACM)* **69**(1) 1–18.
- [16] Goldberg, L. A. and Jerrum, M. (2012) Approximating the partition function of the ferromagnetic Potts model. *J. ACM* **59**(5) 1–31.
- [17] Gruber, C. and Kunz, H. (1971) General properties of polymer systems. *Commun. Math. Phys.* **22**(2) 133–161.

- [18] Galanis, A., Štefankovič, D. and Vigoda, E. (2016) Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *Comb. Prob. Comput.* **25**(4) 500–559.
- [19] Galanis, A., Štefankovič, D., Vigoda, E. and Yang, L. (2016) Ferromagnetic Potts Model: Refined #BIS-hardness and related results. *SIAM J. Comput.* **45**(6) 2004–2065.
- [20] Häggström, O. (1996) The random-cluster model on a homogeneous tree. *Prob. Theory Relat. Fields* **104**(2) 231–253.
- [21] Helmuth, T., Jenssen, M. and Perkins, W. (2020) Finite-size scaling, phase coexistence, and algorithms for the random cluster model on random graphs.
- [22] Hoeffding, W. (1963) Probability inequalities for sums of bounded random variables. *J. Am. Stat. Assoc.* **58** 13–30.
- [23] Helmuth, T., Perkins, W. and Regts, G. (2019) Algorithmic Pirogov–Sinai theory. *Prob. Theory Relat. Fields*.
- [24] Jenssen, M., Keevash, P. and Perkins, W. (2020) Algorithms for #BIS-hard problems on expander graphs. *SIAM J. Comput.* **49**(4) 681–710.
- [25] Jenssen, M. and Perkins, W. (2020) Independent sets in the hypercube revisited. *J. Lond. Math. Soc.* **102**(2) 645–669.
- [26] Jenssen, M., Perkins, W. and Potukuchi, A. (2022) Approximately counting independent sets in bipartite graphs via graph containers. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings. Society for Industrial and Applied Mathematics, pp. 499–516.
- [27] Karger, D. R. (1993) Global min-cuts in RNC, and other ramifications of a simple min-cut algorithm. In *Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'93, USA, January 1993*. Society for Industrial and Applied Mathematics, pp. 21–30.
- [28] Knuth, D. E. (1998) *The Art of Computer Programming*. Addison-Wesley Longman Publishing Co., Inc, vol. 1–3.
- [29] Kotecký, R. and Preiss, D. (1986) Cluster expansion for abstract polymer models. *Commun. Math. Phys.* **103**(3) 491–498.
- [30] Kleitman, D. J. and Winston, K. J. (1980) The asymptotic number of lattices. In *Annals of Discrete Mathematics*, Vol. 6 of *Combinatorial Mathematics, Optimal Designs and Their Applications*, (J. Srivastava, ed.). Elsevier, pp. 243–249.
- [31] Kleitman, D. J. and Winston, K. J. (1982) On the number of graphs without 4-cycles. *Discrete Math.* **41**(2) 167–172.
- [32] Lubetzky, E. and Zhao, Y. (2015) On replica symmetry of large deviations in random graphs. *Random Struct. Algor.* **47**(1) 109–146.
- [33] Mantel, W. (1907) Problem 28. *Wiskundige Opgaven* **10** 60–61.
- [34] Potukuchi, A. and Yepremyan, L. (2021) Enumerating independent sets in Abelian Cayley graphs. arXiv, abs/2109.06152.
- [35] Sapozhenko, A. A. (1987) On the number of connected subsets with given cardinality of the boundary in bipartite graphs. *Metody Diskretnogo Analiza* **45** 42–70,96.
- [36] Sapozhenko, A. A. (2001) On the number of independent sets in extenders. *Diskretnaya Matematika* **13**(1) 56–62.
- [37] Sly, A. (2010) Computational transition at the uniqueness threshold. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, Las Vegas, NV, USA. IEEE, pp. 287–296.
- [38] Sokal, A. D. (2005) The multivariate Tutte polynomial (alias Potts model) for graphs and matroids. In *Surveys in Combinatorics 2005*, first edition (B. S. Webb, ed.). Cambridge University Press, pp. 173–226.
- [39] Sly, A. and Sun, N. (2014) Counting in two-spin models on d-regular graphs. *Ann. Prob.* **42**(6) 2383–2416.
- [40] Sah, A., Sawhney, M., Stoner, D. and Zhao, Y. (2020) A reverse Sidorenko inequality. *Invent. Math.* **221**(2) 665–711.