

Sampling in Uniqueness from the Potts and Random-Cluster Models on Random Regular Graphs

Antonio Blanca¹

School of Computer Science, Georgia Institute of Technology, Atlanta GA 30332, USA
ablanca@cc.gatech.edu

Andreas Galanis²

Department of Computer Science, University of Oxford, Parks Road, Oxford, OX1 3QD, UK
andreas.galanis@cs.ox.ac.uk

Leslie Ann Goldberg²

Department of Computer Science, University of Oxford, Parks Road, Oxford, OX1 3QD, UK
leslie.goldberg@cs.ox.ac.uk

Daniel Štefankovič³

Department of Computer Science, University of Rochester, Rochester, NY 14627, USA
stefanko@cs.rochester.edu

Eric Vigoda¹

School of Computer Science, Georgia Institute of Technology, Atlanta GA 30332, USA
vigoda@cc.gatech.edu

Kuan Yang²

Department of Computer Science, University of Oxford, Parks Road, Oxford, OX1 3QD, UK
kuan.yang@cs.ox.ac.uk

Abstract

We consider the problem of sampling from the Potts model on random regular graphs. It is conjectured that sampling is possible when the temperature of the model is in the so-called uniqueness regime of the regular tree, but positive algorithmic results have been for the most part elusive. In this paper, for all integers $q \geq 3$ and $\Delta \geq 3$, we develop algorithms that produce samples within error $o(1)$ from the q -state Potts model on random Δ -regular graphs, whenever the temperature is in uniqueness, for both the ferromagnetic and antiferromagnetic cases.

The algorithm for the antiferromagnetic Potts model is based on iteratively adding the edges of the graph and resampling a bichromatic class that contains the endpoints of the newly added edge. Key to the algorithm is how to perform the resampling step efficiently since bichromatic classes can potentially induce linear-sized components. To this end, we exploit the tree uniqueness to show that the average growth of bichromatic components is typically small, which allows us to use correlation decay algorithms for the resampling step. While the precise uniqueness threshold on the tree is not known for general values of q and Δ in the antiferromagnetic case, our algorithm works throughout uniqueness regardless of its value.

In the case of the ferromagnetic Potts model, we are able to simplify the algorithm significantly by utilising the random-cluster representation of the model. In particular, we demonstrate

¹ Research supported in part by NSF grants CCF-1617306 and CCF-1563838.

² The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) ERC grant agreement no. 334828. The paper reflects only the authors' views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.

³ Research supported in part by NSF grant CCF-1563757.



© Antonio Blanca, Andreas Galanis, Leslie Ann Goldberg, Daniel Štefankovič, Eric Vigoda, and Kuan Yang;

licensed under Creative Commons License CC-BY

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018).

Editors: Eric Blais, Klaus Jansen, José D. P. Rolim, and David Steurer; Article No. 33; pp. 33:1–33:15



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

that a percolation-type algorithm succeeds in sampling from the random-cluster model with parameters p, q on random Δ -regular graphs for all values of $q \geq 1$ and $p < p_c(q, \Delta)$, where $p_c(q, \Delta)$ corresponds to a uniqueness threshold for the model on the Δ -regular tree. When restricted to integer values of q , this yields a simplified algorithm for the ferromagnetic Potts model on random Δ -regular graphs.

2012 ACM Subject Classification Theory of computation \rightarrow Randomness, geometry and discrete structures, Theory of computation \rightarrow Design and analysis of algorithms

Keywords and phrases sampling, Potts model, random regular graphs, phase transitions

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2018.33

Related Version The full version is available at <https://arxiv.org/abs/1804.08111>. The theorem numbering here matches the full version.

1 Introduction

Random constraint satisfaction problems have been thoroughly studied in computer science in an effort to analyse the limits of satisfiability algorithms and understand the structure of hard instances. Analogously, understanding spin systems on random graphs [22, 23, 28, 3, 21, 7, 8] gives insights about the complexity of counting and the efficiency of approximate sampling algorithms. In this paper, we design approximate sampling algorithms for the Potts model on random regular graphs.

The Potts model is a fundamental spin system studied in statistical physics and computer science. The model has two parameters: an integer $q \geq 3$, which represents the number of states/colours of the model, and a real parameter $B > 0$, which corresponds to the so-called “temperature”. We denote the set of colours by $[q] := \{1, \dots, q\}$. For a graph $G = (V, E)$, configurations of the model are all possible assignments of colours to the vertices of the graph. Each assignment $\sigma : V \rightarrow [q]$ has a weight $w_G(\sigma)$ which is determined by the number $m(\sigma)$ of monochromatic edges under σ ; namely, $w_G(\sigma) = B^{m(\sigma)}$. The Gibbs distribution μ_G is defined on the space of all configurations σ by

$$\mu_G(\sigma) = B^{m(\sigma)} / Z_G, \text{ where } Z_G = \sum_{\sigma} B^{m(\sigma)}.$$

We also refer to μ_G as the Potts distribution; the quantity Z_G is known as the partition function. Well-known models closely related to the Potts model are the Ising and colourings models. The Ising model is the special case $q = 2$ of the Potts model, while the q -colourings model is the “zero-temperature” case $B = 0$ of the Potts model, where the distribution is supported on the set of proper q -colourings.

The behaviour of the Potts model has significant differences depending on whether B is less or larger than 1. When $B < 1$, configurations where most neighbouring vertices have different colours have large weight and the model is called *antiferromagnetic*; in contrast, when $B > 1$, configurations where most neighbouring vertices have the same colours have large weight and the model is called *ferromagnetic*. One difference between the two cases that will be relevant later is that the ferromagnetic Potts model admits a random-cluster representation – the details of this representation are given in Section 2.1.

Sampling from the Potts model is a problem that is frequently encountered in running simulations in statistical physics or inference tasks in computer science. To determine the efficiency and accuracy of sampling methods, it is relevant to consider the underlying phase

transitions, which signify abrupt changes in the properties of the Gibbs distribution when the underlying parameter changes. The so-called uniqueness phase transition captures the sensitivity of the state of a vertex to fixing far-away boundary conditions. As an example, in the case of the ferromagnetic Potts model on the Δ -regular tree, uniqueness holds when root-to-leaves correlations in the Potts distribution vanish as the height of the tree goes to infinity; it is known that this holds iff $B < B_c(q, \Delta)$, where $B_c(q, \Delta)$ is the “uniqueness threshold” (cf. (2) for its value). Connecting the uniqueness phase transition with the performance of algorithms is a difficult task that is largely under development. This connection is well-understood on the grid, where it is known that the mixing time of local Markov chains, such as the Glauber dynamics, switches from polynomial to exponential at the corresponding uniqueness threshold, see for example [18, 17, 27, 1, 19, 2].

For random Δ -regular graphs or, more generally, graphs with maximum degree Δ , the uniqueness threshold on the Δ -regular tree becomes relevant. For certain two-state models, such as the ferromagnetic Ising model and the hard-core model, it has been proved that Glauber dynamics mixes rapidly when the underlying parameter is in uniqueness and that the dynamics mixes slowly otherwise, see [22, 23, 8]. The same picture is conjectured to hold for the Potts model as well, but this remains open. More generally, there has been significant progress in understanding the complexity of sampling from the Gibbs distribution in two-state systems, but for multi-state systems progress has been slower, especially on the algorithmic side.

In this paper, for all integers $q \geq 3$ and $\Delta \geq 3$, we design approximate sampling algorithms for the q -state Potts model on random Δ -regular graphs (regular graphs with n vertices chosen uniformly at random), when the parameter B lies in the uniqueness regime of the regular tree, for both the ferromagnetic and antiferromagnetic cases. Our algorithms are not based on a Markov chain approach but proceed by iteratively adding the edges of the graph and performing a resampling step at each stage. As such, our algorithms can produce samples that are within error $1/n^\delta$ from the Potts distribution for some fixed constant $\delta > 0$ (which depends on B, q, Δ).

► **Remark.** There are certain “bad” Δ -regular graphs where the algorithms will fail to produce samples with the desired accuracy; saying that the algorithms work on random Δ -regular graphs means that the number of these “bad” graphs with n vertices is a vanishing fraction of all Δ -regular graphs with n vertices for large n . Moreover, we can recognise the “good” graphs (where our algorithms will successfully produce samples with the desired accuracy) in polynomial time.

Our approach is inspired by Efthymiou’s algorithm [6, 7] for sampling q -colourings on $G(n, d/n)$; the algorithm there also proceeds by iteratively adding the edges of the graph and exploits the uniqueness on the tree to show that the sampling error is small. However, for the antiferromagnetic Potts model, the resampling step turns out to be significantly more involved and we need substantial amount of work to ensure that it can be carried out efficiently, as we explain in detail in Section 3. Nevertheless, for the ferromagnetic case, we manage to give a far simpler algorithm by utilising the random-cluster representation of the model (see Section 2.1). In particular, we demonstrate that a percolation-type algorithm succeeds in sampling approximately from the random-cluster model with parameters p, q on random Δ -regular graphs for all values of $q \geq 1$ and $p < p_c(q, \Delta)$, where $p_c(q, \Delta)$ corresponds to a uniqueness threshold for the model on the Δ -regular tree. When restricted to integer values of q , this yields a simple algorithm for the ferromagnetic Potts model on random Δ -regular graphs.

To conclude this introductory section, we remark that, for many antiferromagnetic spin systems on random graphs, typical configurations in the Gibbs distribution display absence of long-range correlations even beyond the uniqueness threshold, up to the so-called reconstruction threshold [20, 12]. Note that uniqueness guarantees the absence of long-range correlations under a “worst-case” boundary, while non-reconstruction only asserts the absence of long-range correlations under “typical” boundaries; it is widely open whether this weaker notion is in fact sufficient for sampling on random graphs. On an analogous note, for the ferromagnetic Potts model on random regular graphs, the structure of typical configurations can be fairly well understood using probabilistic arguments for all temperatures (see, e.g., [5, 11]) and it would be very interesting to exploit this structure for the design of sampling algorithms beyond the uniqueness threshold.

2 Definitions and Main Results

To formally state our results, we first review in Section 2.1 the definition of the random-cluster model. In Section 2.2, we state results from the literature about uniqueness on the regular tree for the Potts and random-cluster models. Then, in Section 2.3, we state our algorithmic results for the ferromagnetic Potts and random-cluster models and, in Section 2.4, our result for the antiferromagnetic Potts model.

2.1 The random-cluster model

The random-cluster model has two parameters $p \in [0, 1]$ and $q > 0$; note that q in this case can take non-integer values. For a graph $G = (V, E)$, we denote the random-cluster distribution on G by φ_G ; this distribution is supported on the set of all edge subsets. In particular, for $S \subseteq E$, let $k(S)$ be the number of connected components in the graph $G' = (V, S)$ (isolated vertices do count). Then,

$$\varphi_G(S) = \frac{p^{|S|}(1-p)^{|E \setminus S|}q^{k(S)}}{Z_G^{\text{rc}}}, \text{ where } Z_G^{\text{rc}} = \sum_{S \subseteq E} p^{|S|}(1-p)^{|E \setminus S|}q^{k(S)}.$$

Following standard terminology, each edge in S will be called open, while each edge in $E \setminus S$ closed. For integer values of q , the random-cluster and ferromagnetic Potts models are connected as follows.

► **Lemma 1** (see, e.g., [13]). *Let $q \geq 2$ be an integer, $B > 1$, and $p = 1 - 1/B$. Then, the following hold for any graph $G = (V, E)$.*

- *Let $S \subseteq E$ be distributed according to the RC distribution φ_G with parameters p, q . Consider the configuration σ obtained from S by assigning each component in the graph (V, S) a random colour from $[q]$ independently. Then, σ is distributed according to the Potts distribution μ_G with parameter B .*
- *Conversely, suppose that $\sigma : V \rightarrow [q]$ is distributed according to the Potts distribution μ_G with parameter B . Consider $S \subseteq E$ obtained by adding to S each monochromatic edge under σ with probability p independently. Then, S is distributed according to the RC distribution φ_G with parameters p, q .*

2.2 Uniqueness for Potts and random-cluster models on the tree

In this section, we review uniqueness on the tree for the Potts and random-cluster models.

We start with the Potts model. For a configuration σ and a set U , we denote by σ_U the restriction of σ to the set U ; in the case of a single vertex u , we simply write σ_u to denote the colour of u . Denote by \mathbb{T}_Δ the infinite $(\Delta - 1)$ -ary tree with root vertex ρ and, for an integer $h \geq 0$, denote by T_h the subtree of \mathbb{T}_Δ induced by the vertices at distance $\leq h$ from ρ . Let L_h be the set of leaves of T_h .

► **Definition 2.** Let $B > 0$ and $q, \Delta \geq 3$ be integers. The q -state Potts model with parameter $B > 0$ has *uniqueness* on the infinite $(\Delta - 1)$ -ary tree if, for all colours $c \in [q]$, it holds that

$$\limsup_{h \rightarrow \infty} \max_{\tau: L_h \rightarrow [q]} \left| \mu_{T_h}(\sigma_\rho = c \mid \sigma_{L_h} = \tau) - \frac{1}{q} \right| = 0. \tag{1}$$

For the ferromagnetic q -state Potts model ($B > 1$), it is known that uniqueness holds on the $(\Delta - 1)$ -ary tree iff $B < B_c(q, \Delta)$, where

$$B_c(q, \Delta) = 1 + \inf_{y > 1} h(y), \text{ where } h(y) := \frac{(y - 1)(y^{\Delta-1} + q - 1)}{y^{\Delta-1} - y}. \tag{2}$$

For the antiferromagnetic Potts model ($B < 1$), the uniqueness threshold on the tree is not yet known in full generality. It is known that the model does not have uniqueness when $B < \frac{\Delta - q}{\Delta}$ [10] and therefore $B \geq (\Delta - q)/\Delta$ is a necessary condition for uniqueness to hold. It is also conjectured that this condition is sufficient but this has only been established for small values of q and Δ [9]. In the case $q = 3$, [9] also established the uniqueness threshold for all Δ : for $\Delta \geq 4$, uniqueness holds iff $B \in [(\Delta - 3)/\Delta, 1)$ and, for $\Delta = 3$ uniqueness holds iff $B \in (0, 1)$. For the q -colourings model ($B = 0$), Jonasson [16], building on work of Brightwell and Winkler [4], established that the model has uniqueness iff $q > \Delta$.

Uniqueness for the random-cluster model on the tree is less straightforward to define. Häggström [14] studied uniqueness of random-cluster measures on the infinite $(\Delta - 1)$ -ary tree where all infinite components are connected “at infinity” – we review his results in more detail in Section 7.1 of the full version. He showed that, for all $q \geq 1$, a sufficient condition for uniqueness is that $p < p_c(q, \Delta)$, where $p_c(q, \Delta)$ is given by⁴:

$$p_c(q, \Delta) = 1 - \frac{1}{1 + \inf_{y > 1} h(y)}, \text{ where } h(y) := \frac{(y - 1)(y^{\Delta-1} + q - 1)}{y^{\Delta-1} - y}. \tag{3}$$

Note that the critical values in (2) and (3) are connected for integer values of q via $p_c(q, \Delta) = 1 - \frac{1}{B_c(q, \Delta)}$. Häggström [14] also conjectured that uniqueness for the random-cluster model holds on \mathbb{T}_Δ when $p > \frac{q}{q + \Delta - 2}$ for all $q \geq 1$; this remains open but progress has been made in [15].

2.3 Sampling ferro Potts and random-cluster models on random regular graphs

We begin by stating our result for the random-cluster model on random regular graphs.

► **Theorem 6.** Let $\Delta \geq 3$, $q \geq 1$ and $p < p_c(q, \Delta)$. Then, there exists a constant $\delta > 0$ such that, for all sufficiently large n , the following holds with probability $1 - o(1)$ over the choice of a random Δ -regular graph $G = (V, E)$ with n vertices.

There is a polynomial-time algorithm which, on input the graph G , outputs a random set $S \subseteq E$ whose distribution ν_S is within total variation distance $O(1/n^\delta)$ from the RC distribution φ_G with parameters p, q , i.e., $\|\nu_S - \varphi_G\|_{TV} = O(1/n^\delta)$.

⁴ In [14], $p_c(q, \Delta)$ is defined in a different way, but the two definitions are equivalent for all $q \geq 1$.

For integer values of q , Theorem 6 combined with the translation between the random-cluster and Potts models (cf. Lemma 1) yields a sampling algorithm for the ferromagnetic q -state Potts model on random regular graphs. Since uniqueness for the ferromagnetic Potts model holds iff $B < B_c(q, \Delta)$ and $p_c(q, \Delta) = 1 - \frac{1}{B_c(q, \Delta)}$, we therefore have the following corollary of Theorem 6.

► **Corollary 7.** *Let $\Delta \geq 3$, $q \geq 3$ and $B > 1$ be in the uniqueness regime of the $(\Delta - 1)$ -ary tree. Then, there exists a constant $\delta > 0$ such that, for all sufficiently large n , the following holds with probability $1 - o(1)$ over the choice of a random Δ -regular graph $G = (V, E)$ with n vertices.*

There is a polynomial-time algorithm which, on input the graph G , outputs a random assignment $\sigma : V \rightarrow [q]$ whose distribution ν_σ is within total variation distance $O(1/n^\delta)$ from the Potts distribution μ_G with parameter B , i.e., $\|\nu_\sigma - \mu_G\|_{\text{TV}} = O(1/n^\delta)$.

2.4 Sampling antiferro Potts on random Δ -regular graphs

The algorithm of Corollary 7 for the ferromagnetic Potts model does not extend to the antiferromagnetic case since there is no analogous connection with the random-cluster model in this case. Nevertheless, we are able to design a sampling algorithm on random regular graphs when the parameter B is in uniqueness via a far more elaborate approach which consists of recolouring (large) bichromatic colour classes.

► **Theorem 8.** *Let $\Delta \geq 3$, $q \geq 3$ and $B \in (0, 1)$ be in the uniqueness regime of the $(\Delta - 1)$ -ary tree with $B \neq (\Delta - q)/\Delta$. Then, there exists a constant $\delta > 0$ such that, for all sufficiently large n , the following holds with probability $1 - o(1)$ over the choice of a random Δ -regular graph $G = (V, E)$ with n vertices.*

There is a polynomial-time algorithm which, on input the graph G , outputs a random assignment $\sigma : V \rightarrow [q]$ whose distribution ν_σ is within total variation distance $O(1/n^\delta)$ from the Potts distribution μ_G with parameter B , i.e., $\|\nu_\sigma - \mu_G\|_{\text{TV}} = O(1/n^\delta)$.

Note, the algorithm in the antiferromagnetic case works throughout uniqueness apart from the single point $(\Delta - q)/\Delta$, where uniqueness on the tree is expected to hold but the model is conjectured to be at criticality.

3 Proof Approach

In this section, we outline the main idea behind the algorithms of Theorems 6 and 8, and the key obstacles that we have to address. We focus on the antiferromagnetic Potts model where the details are much more complex and discuss how we get the simplification for the ferromagnetic case via the random-cluster model later.

► **Definition 9.** For an n -vertex graph $G = (V, E)$ with maximum degree Δ , a cycle is *short* if its length is at most $\frac{1}{5} \log_{\Delta-1} n$, and is *long* otherwise.

Let G be a uniformly random Δ -regular graph with n vertices. Following the approach of Efthymiou [7], our algorithm starts from the subgraph of G consisting of all short cycles, which we denote by G' . It is fairly standard to show that, with probability $1 - o(1)$ over the choice of G , the subgraph G' is a *disjoint* union of short cycles, see Lemma 12. It is therefore possible to sample a configuration σ' on G' which is distributed according to the Potts distribution $\mu_{G'}$ (exactly). This can be accomplished in several ways; in fact, since the cycles are disjoint and each cycle has logarithmic length, this initial sampling step can even be done via brute force in polynomial time (though it is not hard to come up with much faster algorithms).

After this initial preprocessing, the algorithm then proceeds by adding sequentially the edges that do not belong to short cycles. At each step, the current configuration is updated with the aim to preserve its distribution close to the Potts distribution of the new graph (with the edge that we just added). Key to this update procedure is a resampling step which is performed only when the endpoints of a newly added edge $\{u, v\}$ happen to have the same colours under the current configuration; intuitively, some action is required in this case because the weight of the current configuration reduces by a factor of $B < 1$ in the new graph (because of the added edge). The resampling step consists of recolouring a bichromatic class.

► **Definition 10.** Let $G = (V, E)$ be a graph and $\sigma : V \rightarrow [q]$ be a configuration. For colours $c_1, c_2 \in [q]$, let $\sigma^{-1}(c_1, c_2)$ be the set of vertices that have either colour c_1 or colour c_2 under σ . For distinct colours $c_1, c_2 \in [q]$, we say that $U = \sigma^{-1}(c_1, c_2)$ is the (c_1, c_2) -colour-class of σ and that U is a bichromatic class under σ . We refer to a connected component of $G[U]$ as a bichromatic component.

In the proper colourings case ($B = 0$), Efthymiou [7] demonstrated that the resampling step when adding an edge $e = \{u, v\}$ can be done by just flipping the colours of a bichromatic component chosen uniformly at random among those containing one of the vertices u and v (say u). The rough idea there is that, when the colourings model is in uniqueness, the bichromatic components on a random graph are typically small in size. At the same time, by the initial preprocessing step, the edge $e = \{u, v\}$ does not belong to a short cycle and therefore u and v are far away in the graph without e . Hence, u and v are unlikely to belong to the same bichromatic component and the flipping step will succeed in giving u and v different colours with good probability.

Unfortunately, this flipping method does not work for the antiferromagnetic Potts model. It turns out that when $q < \Delta$ and even when the Potts model is in uniqueness, bichromatic components can be large and therefore u and v may belong to the same bichromatic component. To make matters worse, these bichromatic components can be quite complicated (with many short/long cycles) and we need a more elaborate approach in our setting to succeed in giving u and v different colours without introducing significant bias to the sampler.

The key to overcoming these obstacles lies in the observation that the assignment of the two colours in a bichromatic component follows the Ising distribution, see Observation 22 for the precise formulation. Hence we can hope to use an approximate sampling algorithm for the Ising model in the resampling step. The natural implementation of this idea however fails: known algorithms for the antiferromagnetic Ising model, based on correlation decay, work as long as $B > (\Delta - 2)/\Delta$, where Δ is the maximum degree of the graph [25, 29]. In general, this inequality is not satisfied for us, i.e., there exist B in the uniqueness regime such that $B < \frac{\Delta-2}{\Delta}$.

Fortunately, we can employ fairly recent technology for two-state models [22, 26, 24] which demonstrates that the graph parameter that matters is not actually the maximum degree of the graph but rather the “average growth” of the graph. While we cannot apply any of the existing results in the literature directly, adapting these ideas to the antiferromagnetic Ising model is fairly straightforward, using results from Mossel and Sly [22]. The more difficult part in our setting is proving that the average growth of the bichromatic components that we consider for resampling is indeed small for “typical” configurations σ (note that in the worst case, the whole graph can be a bichromatic class which has large average growth for our purposes, so a probability estimate over σ is indeed due). Let us first formalise the notion of average growth that we use.

► **Definition 11.** Let M, b be positive constants and $G = (V, E)$ be a graph with n vertices. We say that G has average growth b up to depth $L = \lceil M \log n \rceil$ if for all vertices $v \in V$ the total number of paths with L vertices starting from v is less than b^L .

The notion of average growth is similar to the notion of connective constant for finite graphs used in [26, 24], the reason for the slightly different definition is that we will need an explicit handle on the constant M controlling the depth. Note that since we only consider paths with a fixed logarithmic length, this places a lower bound on the accuracy of the sampling algorithm. Nevertheless, by choosing the constant M sufficiently large, this will still be sufficient to make the error of our sampler polynomially small. In particular, as long as the inequality $b \frac{1-B}{1+B} < 1$ is satisfied, for all sufficiently large M , we obtain an approximate sampler for the antiferromagnetic Ising model with parameter B on graphs of average growth b up to depth $L = \lceil M \log n \rceil$, see Theorem 24 for details.

The next stage is to bound the average growth of bichromatic classes. Here, we utilise the tree uniqueness and the tree-like structure of random Δ -regular graphs (cf. Lemma 15) to provide an upper bound on the probability that a path is bichromatic. For paths of logarithmic length L , we show in Lemma 31 that this probability is bounded above by K^L , where K is roughly $(1+B)/(B+q-1)$. Since in a Δ -regular graph there are at most $\Delta(\Delta-1)^{L-2}$ paths with L vertices, we therefore obtain that, in most configurations σ , the average growth b of bichromatic components is bounded above by $(\Delta-1)K$. When B is in uniqueness, we have that $B > (\Delta-q)/\Delta$, and therefore the inequality $b \frac{1-B}{1+B} < 1$ that is required for the Ising sampler to work is satisfied (quite tightly in fact).

The final piece is to bound the error that is introduced by the resampling steps; note that, even if the resampling steps were perfect, some error is always introduced as a result of the placement of the new edge which reweights the probability that u and v have different colours. The idea now is to use the correlation decay properties on bichromatic components to show that, in the graph without the edge $\{u, v\}$, the correlation between the colours of u and v is relatively small since in that graph they are far apart (recall that $\{u, v\}$ does not belong to a short cycle in G). In Lemma 25, we show that the correlation between u and v can be upper bounded as a weighted sum over paths connecting u and v . Unfortunately, it is not possible to ensure that the correlation is small for each step separately since the natural union bound does not work. The right way to control the error of the sampler is to aggregate over all steps; we thus obtain a bound on the error via a weighted sum over logarithmically long cycles of the random graph G . Using a simple expectation argument of this latter quantity (and Markov's inequality), we obtain that the error will be small with probability $1 - o(1)$ over the choice of the graph G .

The algorithm that we described for the antiferromagnetic Potts model can actually be adapted to the ferromagnetic case as well. However, as mentioned earlier, we follow a different (and surprisingly simpler) route using the random-cluster representation of the model. At a very rough level, the reason behind the simplification is that the components in the random-cluster model provide a much better grip on capturing the properties of the Potts distribution than the bichromatic-component proxy we used earlier. Indeed, just as we described in the antiferromagnetic case, bichromatic components for the ferromagnetic Potts model can also be linear-sized. However, once we translate the Potts configuration to its random-cluster representation (cf. Lemma 1), the components in the latter are small in size (when the model is in the uniqueness region $p < p_c(q, \Delta)$) and therefore vertices that are far away do not belong to the same component. This allows us to perform the resampling step in the random-cluster model by a simple percolation procedure. The details can be found in Section 5.

4 Properties of random regular graphs

In this section, we state and prove structural properties of random Δ -regular graphs which ensure that our algorithms for the random-cluster and Potts models have the desired accuracy (cf. Remark 1). The proofs of these lemmas are fairly standard and are obtained by working in the well-known configuration model (see full version for details). The following lemma guarantees that short cycles are disjoint in a random Δ -regular graph.

► **Lemma 12.** *Let $\Delta \geq 3$ be an integer. Then, with probability $1 - o(1)$ over the choice of a uniformly random Δ -regular graph with n vertices, any two distinct cycles of length $\leq \frac{1}{5} \log_{\Delta-1} n$ are disjoint, i.e., they do not share any common vertices or edges.*

The next lemma guarantees that certain weighted sums over cycles are small; this bound will be used to show that the aggregate error of our samplers is small (cf. Section 3).

► **Lemma 13.** *Let $\Delta \geq 3$. Then, for any constant $W > \Delta - 1$ and any constant $\ell_0 > 0$, there exists a constant $\delta > 0$ such that the following holds with probability $1 - O(1/n^\delta)$ over the choice of $G \sim \mathcal{G}_{n,\Delta}$. Let C_ℓ denote the number of cycles of length ℓ . Then, $\sum_{\ell \geq \ell_0 \log n} \ell C_\ell / W^\ell \leq 1/(2n^\delta)$.*

Our next lemma captures the tree-like structure of random Δ -regular graphs that will be relevant for us. In particular, we give a description of the neighbourhood structure around a path. To do this accurately, we will need a few definitions. Let $G = (V, E)$ be a graph. For a vertex $v \in V$ and integer $h \geq 0$, we denote by $\Gamma_h(G, v)$ the set of vertices at distance $\leq h$ from v .

► **Definition 14.** Let G be a graph and P be a path in G with vertices u_1, \dots, u_ℓ . Let $G \setminus P$ be the graph obtained from G by removing the edges of the path P . Then, for an integer $h \geq 0$, the h -graph-neighbourhood of the path P is the subgraph of $G \setminus P$ induced by the vertex set $\bigcup_{i \in [\ell]} \Gamma_h(G \setminus P, u_i)$. A connected component of the h -graph-neighbourhood will be called *isolated* if it contains *exactly one* of the vertices u_1, \dots, u_ℓ .

► **Lemma 15.** *Let $\Delta \geq 3$. Then, for any constant integer $h \geq 0$ and any $\epsilon > 0$, there exists a constant $\ell_1 > 0$ such that the following holds for all sufficiently large n .*

With probability $1 - O(1/n^2)$ over the choice of $G \sim \mathcal{G}_{n,\Delta}$, every path P in G with ℓ vertices with $\ell_1 \leq \ell \leq n^{9/10}$ has an h -graph-neighbourhood with at least $(1 - \epsilon)\ell$ isolated tree components.

To conclude this section, we clarify a small point relevant to Remark 1. We will only utilise Lemma 15 for paths of logarithmic length (despite that the lemma is stated for convenience for much longer paths) and therefore the property can be checked in polynomial time. Similarly, the sum in Lemma 13 will only be considered for cycles of logarithmic length and hence the (restricted) inequality can also be checked in polynomial time.

5 Algorithm for the random-cluster model

Theorem 6 is obtained by analysing the following algorithm when the input is a random Δ -regular graph $G = (V, E)$. The detailed description of the algorithm is given in Figure 1 of the full version.

Theorem 21 of the full version details the performance of the algorithm when the input is a random Δ -regular graph; the statement is analogous to that of Theorem 6 so we omit it from this short version.

```

 $E' := \{e \in E \mid e \text{ belongs to a short cycle}\}$ 
if  $G' = (V, E')$  contains a component which is neither a cycle nor an isolated vertex then
  FAIL
else { Sample an RC configuration  $S' \subseteq E'$  on  $G'$  (according to  $\varphi_{G'}$ );
  Add to  $S'$  each edge in  $E \setminus E'$  independently with probability  $p/(q + (1 - p)q)$ ;
  Output the resulting set  $S \subseteq E$  }

```

To motivate the algorithm, we first display how to update an RC configuration when we add a single edge $\{u, v\}$. In fact, to control the effect of adding an edge, we need a bound on the probability of the event that there is an open path between u and v ; we denote this event by $u \leftrightarrow v$.

► **Lemma 17.** *Let $p \in (0, 1)$ and $q \geq 1$, and consider arbitrary $\epsilon \in (0, 1/q)$.*

Let $G = (V, E)$ be a graph and u, v be two vertices such that $\{u, v\} \notin E$ and $\varphi_G(u \leftrightarrow v) \leq \epsilon$. Consider the graph $G' = (V, E')$ obtained from G by adding the edge $\{u, v\}$. Sample a random subset of edges $Y \subseteq E'$ as follows: first, sample a subset of edges $X \subseteq E$ according to the RC measure φ_G and, then, set $Y = X \cup \{e\}$ with probability $p/(q + (1 - p)q)$, and $Y = X$ otherwise.

Then, the distribution of Y , denoted by ν_Y , is within total variation distance $2q\epsilon$ from the RC distribution $\varphi_{G'}$ on G' with parameters p, q , i.e., $\|\nu_Y - \varphi_{G'}\|_{\text{TV}} \leq 2q\epsilon$.

To give the intuition behind the sampling procedure of Lemma 17, suppose that u and v were in distinct components of G . Then $\varphi_G(u \leftrightarrow v) = 0$ and the sampling procedure in Lemma 17 would actually produce an exact sample from $\varphi_{G'}$. In the setting of Lemma 17, we do not have this ideal situation but we have a close analogue, i.e., the probability of the event that u and v belong to the same component in a random-cluster configuration is less than ϵ .

To utilise Lemma 17, we need to upper bound the probability that two vertices belong to the same component in a RC configuration. In turn, it suffices to bound the probability that there is an open path between the vertices. To this end, we utilise the uniqueness and the tree-like structure around paths (cf. Definition 14) to show the following.

► **Lemma 18.** *Let $\Delta \geq 3$ be an integer, $q \geq 1$ and $p < p_c(q, \Delta)$. There exist constants $K < 1/(\Delta - 1)$ and $\epsilon > 0$ such that the following holds for all sufficiently large integers ℓ and h . Let G be a Δ -regular graph and P be a path with ℓ vertices whose h -graph-neighbourhood contains $(1 - \epsilon)\ell$ isolated tree components. Let φ_G be the RC distribution on G with parameters p, q . Then, $\varphi_G(\text{path } P \text{ is open}) \leq K^\ell$.*

The proof of Lemma 18 is the most intricate part of our analysis for the random-cluster model algorithm. The difficulty is that the random-cluster model cannot be treated as a spin system due to the dependence of the model on the number of components. The rough intuition behind the proof is that, in the graph without the edges of the path P , the tree-like structure combined with uniqueness on the tree imply that (most) vertices in P belong to distinct components. As a result, when we add the edges of the path P , the trees decorating the path P do not have a significant effect on the number of components that vertices in P belong to and can be treated essentially as “not being there” in the analysis. Of course, the actual argument needs to account for all these effects quite carefully, see also Section 7.2 of the full version.

Using monotonicity properties of the random-cluster model (see Lemma 19 in the full version), we can extend the bound in Lemma 18 to arbitrary subgraphs of a target graph G . In particular, suppose that G, P are as in Lemma 18 and that G' is a subgraph of G which contains the path P . Then it also holds that $\varphi_{G'}(P \text{ is open}) \leq K^\ell$. Using this and Lemma 13, we obtain the following.

► **Lemma 20.** *Let $\Delta \geq 3$ be an integer, $q \geq 1$ and $p < p_c(q, \Delta)$. Then, there exists a constant $\delta > 0$ such that, for all sufficiently large n , the following holds with probability $1 - o(1)$ over the choice of a uniformly random Δ -regular graph $G = (V, E)$ with n vertices.*

Let e_1, \dots, e_t be the edges of G that do not belong to short cycles. For $j \in [t]$, let $e_j = \{u_j, v_j\}$ and G_j be the subgraph $G \setminus \{e_1, \dots, e_j\}$. Then, it holds that $\sum_{j=1}^t \varphi_{G_j}(u_j \leftrightarrow v_j) \leq 1/n^\delta$.

Combining Lemmas 17 and 20, the proof of Theorem 6 is concluded in Section 5.3 of the full version.

6 Algorithm for the antiferromagnetic Potts model

In this section, following the outline of Section 3, we give some more precise technical details for the algorithm in the case of the antiferromagnetic Potts model (see also the relevant Section 6 of the full version).

6.1 Connection between Potts on bichromatic classes and the Ising model.

Our sampling algorithm for the Potts model uses an approximate sampler for the Ising model as a subroutine. Recall that the Ising model is the special case $q = 2$ of the Potts model; to distinguish between the models, we will use π_G to denote the Ising distribution on G with parameter B . Sometimes we will need to replace the binary set of states $\{1, 2\}$ in the Ising model by other binary sets to facilitate the arguments; we use $\pi_G^{c_1, c_2}$ to denote the Ising distribution with binary set of states $\{c_1, c_2\}$ (we will have that $c_1, c_2 \in [q]$). For vertices u, v , we also write $\pi_{G, u, v}^{c_1, c_2}$ to denote the Ising distribution on G conditioned on u taking the state c_1 and v the state c_2 . The following observation details the connection between the distribution of the Potts model on bichromatic classes and the Ising model.

► **Observation 22.** *Let $q \geq 3$ and $B > 0$. Let $G = (V, E)$ be a graph, U be a subset of V and c, c' be distinct colours in $[q]$. Then, for any configuration $\eta : U \rightarrow \{c, c'\}$, it holds that $\mu_G(\sigma_U = \eta \mid \sigma^{-1}(c, c') = U) = \pi_{G[U]}^{c, c'}(\eta)$, i.e., conditioned on U being the (c, c') -colour-class in the Potts distribution μ_G , the marginal distribution on U is the Ising distribution $\pi_{G[U]}$ (with set of states $\{c, c'\}$).*

Adapting results of [22], we show the following algorithm for the Ising model in Section 9 of the full version.

► **Theorem 24.** *Let $B \in (0, 1)$ and $b > 0$ be constants such that $b \frac{1-B}{1+B} < 1$, and let $\Delta \geq 3$ be an integer. Then, there exists $M_0 > 0$ such that the following holds for all $M > M_0$.*

There is a polynomial-time algorithm that, on input an n -vertex graph G with maximum degree at most Δ and average growth b up to depth $L = \lceil M \log n \rceil$, outputs a configuration $\tau : V \rightarrow \{1, 2\}$ whose distribution ν_τ is within total variation distance $1/n^{10}$ from the Ising distribution on G with parameter B , i.e., $\|\nu_\tau - \pi_G\|_{\text{TV}} \leq 1/n^{10}$.

Moreover, the algorithm, when given as additional input two vertices u and v in G , outputs a configuration $\tau : V \rightarrow \{1, 2\}$ such that $\tau_u = 1$ and $\tau_v = 2$, and whose distribution ν_τ satisfies $\|\nu_\tau - \pi_{G, u, v}(\cdot)\|_{\text{TV}} \leq 1/n^{10}$, where $\pi_{G, u, v}$ is the Ising distribution on G conditioned on u having state 1 and v having state 2.

In addition, we will use the following spatial mixing result to analyse the accuracy of our algorithm for the antiferromagnetic Potts model. The proof is given in Section 9.3 of the full version.

► **Lemma 25.** *Let $B \in (0, 1)$ and $b > 0$ be constants such that $b \frac{1-B}{1+B} < 1$. Then, there exists $M'_0 > 0$ such that the following holds for all $M > M'_0$. Let G be an n -vertex graph with average growth b up to depth $L = \lceil M \log n \rceil$, and let u, v be distinct vertices in G . Then*

$$\left| \pi_G(\sigma_u = 1 \mid \sigma_v = 1) - \pi_G(\sigma_u = 1 \mid \sigma_v = 2) \right| \leq \frac{1}{n^{10}} + \sum_{\ell=1}^L P_\ell(G, u, v) \left(\frac{1-B}{1+B} \right)^\ell$$

where $P_\ell(G, u, v)$ is the number of paths with ℓ vertices in G that connect u and v .

6.2 Average growth of bichromatic components in the Potts distribution

To utilise the algorithm of Theorem 24 for our Potts sampler, we will need to bound the average growth of bichromatic classes in a typical Potts configuration on a random regular graph. Our key lemma to achieve this will bound the probability that a path is bichromatic⁵ when the parameter B is in uniqueness, provided that the local neighbourhood around the path (in the sense of Definition 14) has a tree-like structure. The following lemma quantifies this probability bound and is proved in Section 8 of the full version. The proof technique resembles that of Lemma 18: uniqueness on the tree guarantees that, on the graph $G \setminus P$, the colour of most vertices on the path P is roughly uniformly distributed on $[q]$ and therefore, when we attach the edges of the path P back, the hanging trees can essentially be ignored and we can just focus on the edges of the path.

► **Lemma 31.** *Let $\Delta, q \geq 3$ be integers, and $B \in (0, 1)$ be in the uniqueness regime of the $(\Delta - 1)$ -ary tree with $B \neq (\Delta - q)/\Delta$. Then, for any $\epsilon' > 0$, there exists a positive constant $K < \frac{1+B}{B+q-1} + \epsilon'$ and $\epsilon > 0$ such that the following holds for all sufficiently large integers ℓ and h .*

Let G be a graph of maximum degree Δ and P be a path with ℓ vertices whose h -graph-neighbourhood contains $(1 - \epsilon)\ell$ isolated tree components. Let μ_G be the Potts measure on G with parameter B . Then, $\mu_G(\text{path } P \text{ is bichromatic}) \leq K^\ell$.

Recall that for a random Δ -regular graph G , paths do have the tree-like structure of Lemma 31 (cf. Lemma 15), and hence we can aggregate over all paths emanating from an arbitrary vertex (roughly $(\Delta - 1)^\ell$ of them) and get a bound of roughly $(\Delta - 1)K < \frac{(\Delta-1)(1+B)}{B+q-1}$ for the average growth of bichromatic components in a typical configuration σ . This will allow us to use the RESAMPLE subroutine from Section 6.3.

6.3 The resampling subroutine

In Section 3, we discussed that key to the algorithm for the antiferromagnetic case is the resampling subroutine which is invoked when we add a new edge. Here, we give more details about this subroutine and highlight the crucial steps in its analysis.

Let us first abstract somewhat the setting. In particular, let $G = (V, E)$ be a graph and suppose that u, v are vertices in G such that $\{u, v\} \notin E$. Consider the graph $G' = (V, E')$ obtained from G by adding the edge $\{u, v\}$. Given σ distributed according to μ_G , our goal is to produce σ' distributed according to $\mu_{G'}$, perhaps with some small error.

⁵ Let $G = (V, E)$ be a graph and $\sigma : V \rightarrow [q]$. We call a path P bichromatic under σ if there exist colours $c_1, c_2 \in [q]$ such that every vertex u of P satisfies $\sigma_u \in \{c_1, c_2\}$.

We consider the following procedure. If $\sigma_u \neq \sigma_v$, we set $\sigma' = \sigma$. Otherwise, we flip a coin with heads probability $\frac{qB}{B+q-1}$. If the coin comes up heads, we set $\sigma' = \sigma$. Otherwise, suppose that $\sigma_u = \sigma_v = c$ for some colour $c \in [q]$. Then, we choose uniformly at random a bichromatic class U containing u and v , i.e., we pick uniformly at random a colour $c' \in [q] \setminus c$ and set $U = \sigma^{-1}(c, c')$. Then, we resample the colours on U using a sample from the distribution $\pi_{G[U],u,v}^{c,c'}$ (recall that the latter distribution is the Ising distribution with states $\{c, c'\}$ on the subgraph $G[U]$, conditioned on u taking the colour c and v taking the colour c').

Our key lemma shows that the procedure produces σ' whose distribution $\nu_{\sigma'}$ is close to $\mu_{G'}$, provided that the “average correlation” between u and v in a random bichromatic class containing them is small. To make this precise, for a set $U \subseteq V$ such that $u, v \in U$, let

$$\text{Corr}_G(U, u, v) = \left| \frac{\pi_{G[U]}(\eta_u = 1, \eta_v = 1)}{\pi_{G[U]}(\eta_u = 1, \eta_v = 2)} - 1 \right|.$$

i.e., $\text{Corr}_G(U, u, v)$ measures the correlation between u and v in the Ising distribution with parameter B on the subgraph $G[U]$. For a configuration σ , we let $U_\sigma \subseteq V$ be a bichromatic class under σ which contains u and v , chosen uniformly at random among the set of all such classes if there is more than one. Then, the term “average correlation” refers to $\mathbf{E}[\text{Corr}_G(U_\sigma, u, v)]$ where the expectation is over the choice of $\sigma \sim \mu_G$ and the choice of the bichromatic class U_σ containing u and v under σ . If this quantity is small, we show in Lemma 29 of the full version that the distance between the distributions $\nu_{\sigma'}$ and $\mu_{G'}$ is also small.

At this stage, there are two points we need to address. First, the procedure assumed that sampling from the distribution $\pi_{G[U],u,v}^{c,c'}$ can be done efficiently which is not always the case. The way to rectify efficiency is the following: once we are given σ , we first check whether all bichromatic classes in σ have average growth less than $b = (\Delta - 1)K$ up to depth $L = \lceil M \log n \rceil$, where K is the constant in Lemma 31 and M is a sufficiently large constant. Using Lemma 31, we can show that this will be the case with very high probability and therefore we can use the algorithm of Theorem 24 to sample from the distribution $\pi_{G[U],u,v}^{c,c'}$; since the sampler is approximate, this introduces a small sampling error ($\leq 1/n^{10}$), which can be safely ignored. (In the unlikely event where σ contains a bichromatic class with average growth larger than b , we can just set σ' to be a configuration chosen uniformly at random). A detailed description of the algorithm can be found in Figures 3 and 4 of the full version.

The second point we need to address is how to bound the complicated-looking quantity $\mathbf{E}[\text{Corr}_G(U_\sigma, u, v)]$ which is relevant for the accuracy of the algorithm. To do this, we first show that, for any U containing u and v , the quantity $\text{Corr}_G(U, u, v)$ is within a constant factor from the quantity

$$\widehat{\text{Corr}}_G(U, u, v) := \left| \pi_{G[U]}(\eta_u = 1 \mid \eta_v = 1) - \pi_{G[U]}(\eta_u = 1 \mid \eta_v = 2) \right|.$$

We can bound the latter quantity by Lemma 25, using the number of paths connecting u and v in $G[U]$. In particular, for $\sigma \sim \mu_G$, we can upper bound the (random variable) $\widehat{\text{Corr}}_G(U_\sigma, u, v)$ by a sum, ranging over paths P of logarithmic length, of indicator functions that P is bichromatic. Then, Lemma 31 allows us to upper bound the expectation of this sum. We thus obtain an upper bound on $\mathbf{E}[\text{Corr}_G(U_\sigma, u, v)]$ by a weighted sum over the paths P connecting u and v in the graph G .

The proof of Theorem 32 in the full version gives the formal details behind these arguments.

References

- 1 K. S. Alexander. Mixing properties and exponential decay for lattice systems in finite volumes. *Ann. Probab.*, 32(1A):441–487, 2004.
- 2 V. Beffara and H. Duminil-Copin. The self-dual point of the two-dimensional random-cluster model is critical for $q \geq 1$. *Probability Theory and Related Fields*, 153(3):511–542, 2012.
- 3 M. Bordewich, C. Greenhill, and V. Patel. Mixing of the Glauber dynamics for the ferromagnetic Potts model. *Random Structures & Algorithms*, 48(1):21–52, 2016.
- 4 G. R. Brightwell and P. Winkler. Random colorings of a Cayley tree. *Contemporary combinatorics*, 10:247–276, 2002.
- 5 A. Dembo, A. Montanari, A. Sly, and N. Sun. The replica symmetric solution for Potts models on d -regular graphs. *Communications in Mathematical Physics*, 327(2):551–575, 2014.
- 6 C. Efthymiou. A simple algorithm for random colouring $G(n, d/n)$ using $(2 + \epsilon)d$ colours. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '12*, pages 272–280, 2012.
- 7 C. Efthymiou. A simple algorithm for sampling colorings of $G(n, d/n)$ up to the Gibbs uniqueness threshold. *SIAM Journal on Computing*, 45(6):2087–2116, 2016.
- 8 C. Efthymiou, T. P. Hayes, D. Štefankovič, and E. Vigoda. Sampling random colorings of sparse random graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '18*, pages 1759–1771, 2018.
- 9 A. Galanis, L. A. Goldberg, and K. Yang. Uniqueness of the 3-state antiferromagnetic Potts model on the tree. *arXiv/1804.03514*, 2018.
- 10 A. Galanis, D. Štefankovič, and E. Vigoda. Inapproximability for antiferromagnetic spin systems in the tree nonuniqueness region. *J. ACM*, 62(6):50:1–50:60, 2015.
- 11 A. Galanis, D. Štefankovič, E. Vigoda, and L. Yang. Ferromagnetic Potts model: Refined #BIS-hardness and related results. *SIAM Journal on Computing*, 45(6):2004–2065, 2016.
- 12 A. Gerschenfeld and A. Montanari. Reconstruction for models on random graphs. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science, FOCS '07*, pages 194–204, 2007.
- 13 G. Grimmett. *The Random-Cluster Model*. Springer, 2006.
- 14 O. Häggström. The random-cluster model on a homogeneous tree. *Probability Theory and Related Fields*, 104(2):231–253, 1996.
- 15 J. Jonasson. The random cluster model on a general graph and a phase transition characterization of nonamenability. *Stochastic Processes and their Applications*, 79(2):335–354, 1999.
- 16 J. Jonasson. Uniqueness of uniform random colorings of regular trees. *Statistics & Probability Letters*, 57(3):243–248, 2002.
- 17 F. Martinelli and E. Olivieri. Approach to equilibrium of Glauber dynamics in the one phase region. I. The attractive case. *Comm. Math. Phys.*, 161(3):447–486, 1994.
- 18 F. Martinelli and E. Olivieri. Approach to equilibrium of Glauber dynamics in the one phase region. II. The general case. *Communications in Mathematical Physics*, 161(3):487–514, 1994.
- 19 F. Martinelli, E. Olivieri, and R. H. Schonmann. For 2-D lattice spin systems weak mixing implies strong mixing. *Communications in Mathematical Physics*, 165(1):33–47, 1994.
- 20 M. Mézard and A. Montanari. *Information, Physics, and Computation*. Oxford University Press, 2009.
- 21 E. Mossel and A. Sly. Rapid mixing of Gibbs sampling on graphs that are sparse on average. *Random Structures & Algorithms*, 35(2):250–270, 2009.

- 22 E. Mossel and A. Sly. Exact thresholds for Ising–Gibbs samplers on general graphs. *Ann. Probab.*, 41(1):294–328, 2013.
- 23 E. Mossel, D. Weitz, and N. Wormald. On the hardness of sampling independent sets beyond the tree threshold. *Probability Theory and Related Fields*, 143(3):401–439, 2009.
- 24 A. Sinclair, P. Srivastava, D. Štefankovič, and Y. Yin. Spatial mixing and the connective constant: optimal bounds. *Probability Theory and Related Fields*, 168(1):153–197, 2017.
- 25 A. Sinclair, P. Srivastava, and M. Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. *Journal of Statistical Physics*, 155(4):666–686, 2014.
- 26 A. Sinclair, P. Srivastava, and Y. Yin. Spatial mixing and approximation algorithms for graphs with bounded connective constant. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013*, pages 300–309, 2013.
- 27 L. E. Thomas. Bound on the mass gap for finite volume stochastic Ising models at low temperature. *Communications in Mathematical Physics*, 126(1):1–11, 1989.
- 28 Y. Yin and C. Zhang. Sampling in Potts model on sparse random graphs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2016*, pages 47:1–47:22, 2016.
- 29 J. Zhang, H. Liang, and F. Bai. Approximating partition functions of the two-state spin system. *Information Processing Letters*, 111(14):702–710, 2011.