Local Convergence of Random Graph Colorings

Amin Coja-Oghlan[*1], Charilaos Efthymiou[†2], and Nor Jaafari[1]

1   Mathematics Institute, Goethe University of Frankfurt
    10 Robert Mayer St, Frankfurt 60325, Germany
    {acoghlan,jaafari}@math.uni-frankfurt.de
2   College of Computing, Georgia Institute of Technology
    266 Ferst Drive, Atlanta, 30332, USA
    efthymiou@gmail.com

──── **Abstract** ────────────────────────────────

Let $\mathbf{G} = \mathbf{G}(n, m)$ be a random graph whose average degree $d = 2m/n$ is below the k-colorability threshold. If we sample a k-coloring σ of $\mathbf{G}$ uniformly at random, what can we say about the correlations between the colors assigned to vertices that are far apart? According to a prediction from statistical physics, for average degrees below the so-called *condensation threshold* $d_{k,\mathrm{cond}}$, the colors assigned to far away vertices are asymptotically independent [Krzakala et al: PNAS 2007]. We prove this conjecture for k exceeding a certain constant k_0. More generally, we determine the joint distribution of the k-colorings that σ induces locally on the bounded-depth neighborhoods of a fixed number of vertices.
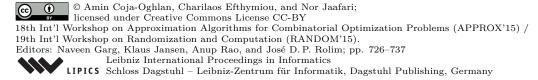
1   Introduction

We let $\mathbf{G} = \mathbf{G}(n, m)$ denote the random graph on the vertex set $[n] = \{1, \ldots, n\}$ with precisely m edges. Unless specified otherwise, we assume that $m = m(n) = \lceil dn/2 \rceil$ for a fixed number $d > 0$.

Going back to the seminal paper of Erdős and Rényi [18], the problem of coloring the random graph $\mathbf{G}$ remains one of the longest-standing challenges in probabilistic combinatorics. Over the past half-century, efforts have been devoted to problems ranging from determining the likely value of the chromatic number [4, 9, 26, 29] and its concentration [5, 27, 36] to algorithmic ones such as constructing or sampling colorings of the random graph [3, 12, 14, 16, 23, 24, 38].

The last few years have witnessed substantial progress w.r.t. estimating $\chi(\mathbf{G})$ accurately. Achlioptas and Friedgut [2] proved that for any number $k \geq 3$ of colors there exists a *sharp threshold sequence* $d_{k-\mathrm{col}} = d_{k-\mathrm{col}}(n)$ such that for any fixed $\varepsilon > 0$, $\mathbf{G}$ is k-colorable w.h.p. if $2m/n < d_{k-\mathrm{col}}(n) - \varepsilon$, while $\chi(\mathbf{G}) > k$ w.h.p. if $2m/n > d_{k-\mathrm{col}}(n) + \varepsilon$. Furthermore, starting

─────────────────────

from [4] there is a sequence of results which improve on the bounds for $d_{k-\text{col}}(n)$. The best current bounds read

$$(2k-1)\ln k - 2\ln 2 + \delta_k \leq \liminf_{n\to\infty} d_{k-\text{col}}(n) \leq \limsup_{n\to\infty} d_{k-\text{col}}(n) \leq (2k-1)\ln k - 1 + \delta_k, \quad (1)$$

where $\lim_{k\to\infty} \delta_k = 0$, see [10, 11]. In addition, the understanding of the geometry of the set of k-colorings has advanced significantly [1, 32], too.

Much of the recent work has been inspired by ideas from statistical physics. Indeed, based on a systematic but non-rigorous approach, the "cavity method" [30, 31], physicists have derived "predictions" on a wide range of problems in combinatorics. More specifically, from the statistical mechanics viewpoint, problems such as random graph coloring are examples of "disordered systems", and the principal interest is in the associated *Gibbs distributions*. For instance, if $k \geq 3$ is an integer and $G = (V, E)$ is a k-colorable graph, this is a probability distribution on $[k]^V$, namely the uniform distribution $\mu_{G,k}$ on the set $\mathcal{S}_k(G)$ of all k-colorings of G. The fact that graph coloring is computationally hard suggests that $\mu_{G,k}$ is a quite a complex object to study, and so it is. Not only does the Gibbs distribution reflect structural properties of the problem such as the geometry of the space of colorings. But it is also expected that its *spatial mixing properties*, i.e. the nature of correlations under the Gibbs distribution, have a substantial impact on the performance of algorithms both for finding and sampling (counting) colorings [13, 16, 28, 38]. Moreover, the existence of long range correlation phenomena under the Gibbs distribution have been related to the hardness of certain computational problems, e.g. see [37, 19, 20].

The usual setting where we study spatial mixing under the Gibbs distribution is as follows: Considering a small region of vertices $\Lambda \subset V(G)$, we analyze how the coloring of Λ is correlated with that of vertices at some distance ω around Λ, as ω increases. It turns out that spatial mixing is the most fundamental variant in understanding the behavior of Gibbs distribution.

As far as the case where $\mathbf{G} = G(n, m)$ is regarded, it seems that the local treelike structure of the graph plays a prominent role in the study of spatial mixing. It is well known that the neighborhood structure within a fixed radius around some vertex $v \in \mathbf{G}$ is asymptotically distributed as a Galton-Watson tree with offspring distribution $Po(d)$, where $d = 2m/n$. This observation motivates the question of whether it is possible to study the spatial mixing properties of the Gibbs distribution on $\mathbf{G}$ by means of the Gibbs distribution on Galton-Watson trees with $Po(d)$ offspring. For certain questions like the so-called Gibbs uniqueness (e.g. see [21]) the reduction from $\mathbf{G}$ to a Galton-Watson tree is trivially true. However, for different kinds of spatial mixing this reduction turns out to be far from trivial, e.g. the reconstructibility we consider here.

In this work we study the above question in its full generality. We investigate the asymptotic relation between the Gibbs distribution over the k-colorings of a $Po(d)$ Galton-Watson tree and the marginal of the Gibbs distribution over the k-colorings of $\mathbf{G}$ on a fixed size neighborhood around some vertex v. We show that these two distributions converge to each other in a very specific way, provided that the average degree d is smaller than the so-called *condensation threshold*

$$d_{k,\text{cond}} = \inf\left\{ d > 0 : \limsup_{n\to\infty} \mathbb{E}[Z_k(G(n,m))^{1/n}] < k(1-1/k)^{d/2} \right\}.$$

The precise value of $d_{k,\text{cond}}$ has recently been determined rigorously (for k exceeding a certain constant k_0) [7]. An asymptotic expansion of that precise formula yields the explicit

expression

$$d_{k,\mathrm{cond}} = (2k - 1)\ln k - 2\ln 2 + \delta_k \quad \text{with } \delta_k \to 0,$$

matching the best current lower bound on $d_{k-\mathrm{col}}$ (cf. (1)). Based on non-rigorous considerations from statistical physics, we do not expect that the aforementioned convergence is true for $d > d_{k,\mathrm{cond}}$ (see further discussion in Section 3).

As a matter of fact we show that the above convergence does hold even if we consider a fixed number of neighborhoods jointly. That is, for some fixed r, we consider r many vertices in $\mathbf{G}$ and a small radius neighborhood around each vertex. We show that the joint distribution of the k-colorings of these neighborhoods under $\mu_{\mathbf{G},k}$ converges to the product of r Gibbs distributions over the k-colourings of a $Po(d)$ Galton-Watson tree.

All the above imply that the reduction from $\mathbf{G}$ to the Galton-Watson tree is indeed correct for studying any generic spatial mixing condition as long as $d < d_{k,\mathrm{cond}}$. A direct corollary of our results is to verify the existence of a certain kind of decorrelation between the color assignments of distant vertices, that was predicted by physicists using the cavity method, see [25]. Furthermore, there is an impressing consequence that comes from the observation that $d_{k,\mathrm{cond}}$ is asymptotically equal to the k-colorability threshold for $\mathbf{G}$. Our result shows that the long-range effects that drive up the chromatic number of $\mathbf{G}$ are not noticeable at a "local scale" (see further discussion in Section 3).

The challenging task in our analysis is that even though we speak about the local behavior of the Gibbs distribution of $\mathbf{G}$, we still need to argue about its global properties. To this end we make a novel use of the "planted trick", introduced in [1]. In particular, we generalize the planted model by introducing what we call the "planted replica model". The planting, now, involves two independent k-colorings rather than just one. Our formalization of the problem as well as the use of planting are quite generic. For this reason, we expect that our approach can be extended to other models on random (hyper)graphs such as the hard-core model, hypergraph two coloring, random k-SAT, etc.

2 Results

For a k-colorable graph $G = (V, E)$, a vertex v and an integer $\omega \geq 0$ we let $\partial^\omega(G, v)$ signify the depth-ω neighborhood of v, i.e., the graph obtained from G by deleting all vertices at a distance greater than ω from v. Additionally, for a set $U \subset V$ we let $\mu_{k,G|U}$ denote the projection of $\mu_{k,G}$ onto $[k]^U$, i.e.,

$$\mu_{k,G|U}(\sigma_0) = \mu_{k,G}\left(\left\{\sigma \in [k]^V : \forall u \in U : \sigma(u) = \sigma_0(u)\right\}\right) \qquad (\sigma_0 \in [k]^U).$$

If H is a subgraph of G we briefly write $\mu_{k,G|H} = \mu_{k,G|V(H)}$. Additionally, given a graph G, we let $\mathbf{v}_1, \mathbf{v}_2, \ldots$ denote vertices of G that are chosen uniformly and independently at random. Finally, let $\|\cdot\|_{\mathrm{TV}}$ be the total variation norm. The main result of this paper is

▶ **Theorem 1.** *There is a constant $k_0 > 0$ such that for any $k \geq k_0$, $d < d_{k,\mathrm{cond}}$, $l \geq 1$, $\omega \geq 0$ we have*

$$\lim_{n \to \infty} \mathbb{E} \left\| \mu_{k,\mathbf{G}|\partial^\omega(\mathbf{G},\mathbf{v}_1) \cup \cdots \cup \partial^\omega(\mathbf{G},\mathbf{v}_l)} - \bigotimes_{i=1}^{l} \mu_{k,\partial^\omega(\mathbf{G},\mathbf{v}_i)} \right\|_{\mathrm{TV}} = 0. \tag{2}$$

In words, suppose that $k \geq k_0$ is not too small, let $d < d_{k,\mathrm{cond}}$ and fix integers l, ω. Choose a random graph $\mathbf{G}$ and pick l vertices $\mathbf{v}_1, \ldots, \mathbf{v}_l$ uniformly and independently at random. Standard properties of the random graph ensure that w.h.p. their depth-ω neighborhoods

$\partial^\omega(\mathbf{G}, \mathbf{v}_1), \ldots, \partial^\omega(\mathbf{G}, \mathbf{v}_l)$ are pairwise disjoint and acyclic. Hence, each $\partial^\omega(\mathbf{G}, \mathbf{v}_i)$ is a tree w.h.p. However, w.h.p. in $\mathbf{G}$ there are paths of length $\Omega(\log n)$ joining (most of) the vertices at distance precisely ω from the randomly choosen "roots" $\mathbf{v}_i$. Now, (2) states that w.h.p. the total variation distance of the following two distributions tends to 0 as $n \to \infty$. Under the first distribution, choose a k-coloring σ of the *entire* random graph $\mathbf{G}$ uniformly at random and consider its projection onto the forest $\partial^\omega(\mathbf{G}, \mathbf{v}_1) \cup \cdots \cup \partial^\omega(\mathbf{G}, \mathbf{v}_l)$. In particular, σ has to respect the constraints imposed by the "long paths" that join the different components $\partial^\omega(\mathbf{G}, \mathbf{v}_i)$. Under the second distribution, ignore these long-range effects and obtain a k-coloring of $\partial^\omega(\mathbf{G}, \mathbf{v}_1) \cup \cdots \cup \partial^\omega(\mathbf{G}, \mathbf{v}_l)$ simply by picking a k-coloring of each tree $\partial^\omega(\mathbf{G}, \mathbf{v}_i)$ independently and uniformly at random (a task that can be performed efficiently by dynamic programming).

Setting $\omega = 0$ in Theorem 1 yields the following statement, which is of interest in its own right.

▶ **Corollary 2.** *There is a number $k_0 > 0$ such that for all $k \geq k_0$, $d < d_{k,\mathrm{cond}}$ and any integer $l > 0$ we have*

$$\lim_{n \to \infty} \mathbb{E} \left\| \mu_{k,\mathbf{G}|\{\mathbf{v}_1, \ldots, \mathbf{v}_l\}} - \bigotimes_{i=1}^{l} \mu_{k,\mathbf{G}|\{\mathbf{v}_i\}} \right\|_{\mathrm{TV}} = 0. \tag{3}$$

By the symmetry of the colors, for each vertex v the "marginal distribution" $\mu_{k,\mathbf{G}|\{v\}}$ is just the uniform distribution on $[k]$ for every vertex v. Hence, Corollary 2 states that for $d < d_{k,\mathrm{cond}}$ w.h.p. in the random graph $\mathbf{G}$ for randomly chosen vertices $\mathbf{v}_1, \ldots, \mathbf{v}_l$ the following is true: if we choose a k-coloring σ of $\mathbf{G}$ at random, then $(\sigma(\mathbf{v}_1), \ldots, \sigma(\mathbf{v}_l)) \in [k]^l$ is asymptotically uniformly distributed. Prior results of Montanari and Gershenfeld [22] and of Montanari, Restrepo and Tetali [33] imply that (3) holds for $d < 2(k-1)\ln(k-1)$, about an additive $\ln k$ below $d_{k,\mathrm{cond}}$.

Another interesting special case occurs if we set $l = 1$ in Theorem 1. In this case we obtain a result about the well-known *reconstruction problem*. Suppose we draw a k-coloring σ of $\mathbf{G}$ at random and we reveal the assignment of a vertex v. The reconstruction problem amounts to studying how the information about the assignment at v biases the distribution of assignments of other vertices in $\mathbf{G}$, i.e. point to set correlation. It is straightforward that the assignments at the neighbors of v are correlated with that of v since they must be distinct. More generally, the reconstruction problem considers the bias on the assignments of vertices on a *fixed* "radius" ω from v. If this correlation persists as $\omega \to \infty$ we say that *reconstruction is possible* on $\mathbf{G}$. Otherwise we say that *reconstruction is impossible*. A similar notion can be defined on a random Galton-Watson tree $\mathbf{T}(k, d)$ with offspring distribution $\mathrm{Po}(d)$. That is, the reconstruction problem considers how the color assignment at the root biases the distribution of the assignments of the vertices at level ω in a random k-colouring of $\mathbf{T}(k, d)$. [1]

▶ **Corollary 3.** *There is a number $k_0 > 0$ such that for all $k \geq k_0$ and $d < d_{k,\mathrm{cond}}$ the following is true.*

Reconstruction is possible on $\mathbf{G}$ $\Leftrightarrow$ tree reconstruction is possible on $T(k, d)$. $\qquad$ (4)

The point of Corollary 3 is that it reduces the reconstruction problem on a combinatorially extremely intricate object, namely the random graph $\mathbf{G}$, to the same problem on a much

---

[1] Formal definitions can be found in [22].

simpler structure, namely the Galton-Watson tree. That said, the reconstruction problem on $\mathbf{T}(d)$ is far from trivial. The best current bounds show that there exists a sequence $(\delta_k)_k \to 0$ such that non-reconstruction holds in $\mathbf{T}(d)$ if $d < (1 - \delta_k)k \ln k$, while reconstruction occurs if $d > (1 + \delta_k)k \ln k$ [15].

Montanari, Restrepo and Tetali [33] proved (4) for $d < 2(k-1)\ln(k-1)$, about an additive $\ln k$ below $d_{k,\mathrm{cond}}$. This gap could alternatively be plugged by invoking recent results on the geometry of the set of k-colorings [6, 10, 32]. However, Corollary 3 is an immediate consequence of Theorem 1.

3   Discussion and related work

Erdős' observation [17] that the random graph $\mathbf{G}$ provides an example of a graph with high girth and high chromatic number is, quite literally, a textbook application of the "probabilistic method". One possible proof is to combine (1) with the well-known observation that for any $\ell \geq 3$ the probability that $\mathbf{G}$ fails to contain a cycle of length at most ℓ remains bounded away from 0 as $n \to \infty$. The difficulty of coming up with a deterministic construction of such a graph highlights the extent to which the phenomenon baffles intuition [34].

Theorem 1 goes one step further. It shows that the long-range effects that drive up the chromatic number of $\mathbf{G}$ are "locally elusive". Indeed, suppose we project a random k-coloring of $\mathbf{G}$ to the depth-ω neighborhood of a bounded number of vertices. Then Theorem 1 shows that asymptotically no traces of the long-range effects that drive up the chromatic number remain as the induced coloring is distributed as though no such effects were present. In particular, w.h.p. *any* k-coloring of the neighborhood of a given vertex v extends in asymptotically the same number of ways to the entire graph. Needless to say, we are unaware of any deterministic construction that exhibits this property.

It is instructive to discuss the relationship between Theorem 1 and the geometry of the set $\mathcal{S}_k(\mathbf{G})$ of k-colorings. For $(1 + \eta_k)k \ln k < d < d_{k,\mathrm{cond}}$, where $\lim_{k\to\infty} \eta_k = 0$, w.h.p. the set $\mathcal{S}_k(\mathbf{G})$ shatters into an exponential number of disjoint "clusters" $\mathcal{C}_1, \ldots, \mathcal{C}_N$, each containing merely an exponentially small fraction of the set $\mathcal{S}_k(\mathbf{G})$ [1]. Additionally, w.h.p. a randomly chosen coloring σ of $\mathbf{G}$ belongs to a cluster $\mathcal{C}_i$ that is "frozen" [1, 32]. That is, there is a number $\Omega(n)$ of vertices v that receive the same color under *all* the k-colorings in $\mathcal{C}_i$. Conversely, the cluster $\mathcal{C}_i$ can be characterised (almost) completely by a map $\tau : V \to [k] \cup \{*\}$ that assigns all frozen vertices their color and that sets all other variables to the "joker color" $*$ [10, 32]. In effect, the *internal* structure of a typical cluster is subject to strong long-range correlations. For instance, if we consider the projection of a random k-coloring τ_i chosen from the cluster $\mathcal{C}_i$ to a bounded number $\mathbf{v}_1, \ldots, \mathbf{v}_l$ of randomly chosen vertices, then the distribution of the resulting color vector $(\tau_i(\mathbf{v}_1), \ldots, \tau_i(\mathbf{v}_l))$ would be far from uniform. What Theorem 2 and Corollary 2 show is that these long range correlations "cancel out perfectly" due to the large overall number of clusters.

No such cancellation is expected to occur for d beyond $d_{k,\mathrm{cond}}$ but below the k-colorability threshold [25]. In fact, non-rigorous but sophisticated arguments from physics predict that in this regime, the so-called *condensed phase*, the set $\mathcal{S}_k(\mathbf{G})$ is dominated by a *bounded* number of frozen clusters. Hence, the joint distribution $(\sigma(\mathbf{v}_1), \ldots, \sigma(\mathbf{v}_l))$ will be a mixture of a bounded number of "frozen" distributions $(\tau_i(\mathbf{v}_1), \ldots, \tau_i(\mathbf{v}_l))$. Consequently, we expect Theorem 1 to be best-possible in terms of the range of d.

Theorem 1 deals with the absence of "long-range correlations" in the random graph coloring problem. It makes sense to raise similar questions in a wide variety of other problems as well. In fact, the method that we develop in this paper is rather generic, and we expect

that it will generalise to various other problems; the following section will provide a detailed discussion. Immediate examples that spring to mind include random hypergraph coloring or models from physics such as the Potts model. Another potential candidate may be the hardcore model on the random regular graph, which was studied by Bhatnagar, Sly and Tetali via a rather different approach [8], and its generalisations [35].

4 Outline

Throughout this section we assume that $k \geq k_0$ and that $d < d_{k,\text{cond}}$.

4.1 Spatial mixing via replicas

The proof of Theorem 1 consists of three components. The first part is to reduce the proof of Theorem 1 to studying the distribution called *random replica model*. This is a distribution over triples $(\mathbf{G}, \sigma_1, \sigma_2)$ such that $\mathbf{G}$ is a k-colorable graph and $\sigma_1, \sigma_2 \in \mathcal{S}_k(\mathbf{G})$. In particular, it is induced by the following experiment.

RR1 Choose a random graph $\mathbf{G} = \mathbf{G}(n, m)$ subject to the condition that $\mathbf{G}$ is k-colorable.

RR2 Sample two k-colorings σ_1, σ_2 of $\mathbf{G}$ uniformly and independently.

In an analogous manner we define the distribution $\mathbf{T}^{\otimes}(k, d)$ as follows: The distribution is over triples $(\mathbf{T}, \tau_1, \tau_2)$, where $\mathbf{T}$ is an instance of the Galton-Watson tree with offspring distribution $\text{Po}(d)$ and τ_1, τ_2 are two independent random k-colorings of $\mathbf{T}$.

So as to proceed with our arguments we need a bit of terminology. A *rooted graph* is a graph $G = (V, E)$ with vertex set $V \subset \mathbb{R}$ together with a distinguished vertex $v_0 \in V$, the *root*. Further, a *dicolored rooted graph* is a k-colorable rooted graph (G, v_0) together with two k-colorings $\sigma_1, \sigma_2 \in \mathcal{S}_k(G)$. An *isomorphism* between two dicolored rooted graphs $(G, v_0, \sigma_1, \sigma_2)$, $(G', v_0', \sigma_1', \sigma_2')$ is a graph isomorphism $\varphi : G \to G'$ such that $\varphi(v_0) = v_0'$, $\sigma_1 = \sigma_1' \circ \varphi$, $\sigma_2 = \sigma_2' \circ \varphi$ and $\varphi(v) < \varphi(w)$ iff $v < w$ for all vertices v, w of G. In words, φ preserves the roots, both colorings and the order of the vertices (which are real numbers). For $\omega \geq 0$ we denote by $\partial^{\omega}(G, v, \sigma_1, \sigma_2)$ the rooted dicolored graph obtained from $(G, v, \sigma_1, \sigma_2)$ by deleting all vertices whose distance from v exceeds ω.

Given Γ a rooted dicolored graph and $\omega \geq 0$, for some graph $G = (V, E)$, $\sigma_1, \sigma_2 \in \mathcal{S}_k(G)$, we let

$$Q_{\Gamma, \omega}(G) = \frac{1}{|V|} \sum_{v \in V} \mathbf{1}\left\{ \partial^{\omega}(G, v, \sigma_1, \sigma_2) \cong \partial^{\omega}\Gamma \right\}.$$

That is, $Q_{\Gamma, \omega}(G)$ is the fraction of vertices of G whose depth-ω neighbourhood, dicolored as in (σ_1, σ_2), are isomorphic to $\partial^{\omega}\Gamma$. The following proposition reduces the proof of Theorem 1 to studying the quantity $Q_{\Gamma, \omega}$ under the random replica model. In particular, it holds

▶ **Proposition 4.** *Assume that for any dicolored rooted tree θ and any integer $\omega \geq 0$, we have*

$$Q_{\theta, \omega}(\mathbf{G}, \sigma_1, \sigma_2) \xrightarrow{\mathbb{P}} \mathbb{P}\left[\partial^{\omega}\mathbf{T}^{\otimes}(d, k) \cong \partial^{\omega}\theta \right]. \tag{5}$$

Then (2) holds for any $\omega \geq 0$, $l \geq 0$.

The proof of Proposition 4 is based on an averaging argument that generalises a very elegant argument from [22]. More details can be found in Section 5.

4.2 Planting replicas

The above implies that we need to study the random variable $Q_{\theta,\omega}(\mathbf{G}, \sigma_1, \sigma_2)$ with $(\mathbf{G}, \sigma_1, \sigma_2)$ chosen from the random replica model. It turns out that studying the random replica model is a formidable task. E.g. for half the range of d we consider there is not even a practical way of finding a k-coloring not to mention generating one at random (cf. the discussion in [1]). For this reason we study the random replica model by means of the *planted replica model*. The planted replica model, is a probability distribution that is easy both to implement and to analyse.

For two maps $\sigma_1, \sigma_2 : [n] \to [k]$ let

$$\mathcal{F}(\sigma_1, \sigma_2) = \sum_{i=1}^{k} \left[\binom{|\sigma_1^{-1}(i)|}{2} + \binom{|\sigma_2^{-1}(i)|}{2} \right] - \sum_{i,j=1}^{k} \binom{|\sigma_1^{-1}(i) \cap \sigma_2^{-1}(j)|}{2},$$

i.e. $\mathcal{F}(\sigma_1, \sigma_2)$ is the number of edges of the complete graph on $[n]$ that are monochromatic in at least one of σ_1 or σ_2.

The planted replica model is induced by the following experiment.

PR1 Sample two maps $\hat{\sigma}_1, \hat{\sigma}_2 : [n] \to [k]$ independently and uniformly at random subject to the condition that $\mathcal{F}(\hat{\sigma}_1, \hat{\sigma}_2) \le \binom{n}{2} - m$.

PR2 Choose a graph $\hat{\mathbf{G}}$ on $[n]$ with precisely m edges uniformly at random, subject to the condition that both $\hat{\sigma}_1, \hat{\sigma}_2$ are proper k-colorings of $\hat{\mathbf{G}}$.

The above experiment is easy to get a handle on. Indeed, for large enough n the conditioning in **PR1** is essentially void as $\mathbb{E}[\mathcal{F}(\hat{\sigma}_1, \hat{\sigma}_2)] \sim (1 - 1/k)^2 n$ and $m = O(n)$. In addition, **PR2** amounts to simply choosing a random set of m edges out of the $\binom{n}{2} - \mathcal{F}(\hat{\sigma}_1, \hat{\sigma}_2)$ "allowed" edges of the complete graph.

The random replica model and the planted replica model are a priori two different distributions. However, we show that they are closely related in the following sense.

▶ **Proposition 5.** *For any sequence $(\mathcal{A}_n)_n$ of events we have*

$$\lim_{n \to \infty} \mathbb{P}\left[(\hat{\mathbf{G}}, \hat{\sigma}_1, \hat{\sigma}_2) \in \mathcal{A}_n \right] = 0 \quad \Rightarrow \quad \lim_{n \to \infty} \mathbb{P}\left[(\mathbf{G}, \sigma_1, \sigma_2) \in \mathcal{A}_n \right] = 0.$$

In probability jargon, Proposition 5 states that the random replica model is contiguous with respect to the planted replica model. That is, triples (G, σ_1, σ_2) that are typical w.r.t. the first distribution are typical w.r.t. the second distribution, too. The above proposition generalises a result from [6]. For more details about Proposition 5, see Section 6.

4.3 A coupling argument

The following proposition summarises the third and last ingredient to the proof of Theorem 1.

▶ **Proposition 6.** *Let θ be a dicolored rooted tree θ and let $\omega \ge 0$. Then $Q_{\theta,\omega}(\hat{\mathbf{G}}, \hat{\sigma}_1, \hat{\sigma}_2)$ converges in probability to $\mathbb{P}[\partial^\omega \mathbf{T}^{\otimes}(k, d) \cong \partial^\omega \theta]$ as $n \to \infty$.*

The proof of Proposition 6 is based on a coupling argument that illustrates how convenient it is to work with the planted replica model. Namely, to prove Proposition 6 it is merely necessary to couple the distribution of the breadth first search tree from a random vertex v in $(\hat{\mathbf{G}}, \hat{\sigma}_1, \hat{\sigma}_2)$ up to depth ω with the truncated random tree $\partial^\omega \mathbf{T}^{\otimes}(k, d)$. The coupling is rather immediate from the definitions of these two distributions.

To complete the proof of Theorem 1, we combine Propositions 5 and 6 to conclude that for any $\theta, \omega,$. $Q_{\theta,\omega}(\mathbf{G}, \sigma_1, \sigma_2)$ converges in probability to $\mathbb{P}[\partial^\omega \mathbf{T}^{\otimes}(d, k) \cong \partial^\omega \theta]$. Hence, the assertion follows from Proposition 4.

5 Proof of Proposition 4

Let G be a graph and $Z_k(G) = |\mathcal{S}_k(G)|$. For some $X : \mathcal{S}_k(G) \to \mathbb{R}$, we write

$$\langle X(\sigma) \rangle_G = \frac{1}{Z_k(G)} \sum_{\sigma \in \mathcal{S}_k(G)} X(\sigma).$$

That is, $\langle X(\sigma) \rangle_G$ denotes the expectation of X over the choice of random colorings of G.

Let some integers $l, \omega \geq 0$, let $\theta_1, \ldots, \theta_l$ be rooted trees and let $\tau_1 \in \mathcal{S}_k(\theta_1), \ldots, \tau_l \in \mathcal{S}_k(\theta_l)$. For a graph G, let $\mathcal{U} = \mathcal{U}(G)$ denote the number of vertex sequences $v_1, \ldots, v_l$ such that $\partial^\omega(G, v_i) \cong \theta_i$ for each $i \in [l]$. Let $Y(G)$ denote the set of vertex sequences $v_1, \ldots, v_l$ that in addition to $\partial^\omega(G, v_i) \cong \theta_i$ for each $i \in [l]$, also satisfy

$$\left| \left\langle \prod_{i=1}^{l} \mathbf{1}\{\partial^\omega(G, v_i, \sigma) \cong (\theta_i, \tau_i)\} \right\rangle_G - \prod_{i=1}^{l} Z_k(\theta_i)^{-1} \right| > \delta,$$

for fixed $\delta > 0$. Conditional on the convergence condition in (5) we show that

$$n^{-l}|Y(\mathbf{G})| \xrightarrow{\mathbb{P}} 0. \tag{6}$$

Then, the proposition follows by using (6) and noting that the random graph converges locally to the Galton-Watson tree with offspring distribution $\mathrm{Po}(d)$, i.e. w.h.p. we have that

$$n^{-l}|\mathcal{U}| = o(1) + \prod_{i=1}^{l} \mathbb{P}[\mathbf{T}(d) \cong \theta_i].$$

For (6), we extend an argument from [22, Proposition 3.2]. Given a sequence $\varepsilon = \varepsilon(n)$, we let $\mathcal{X}_{\theta_1, \ldots, \theta_l}(G, [l], \omega)$ be the set of all vertex sequences $u_1, \ldots, u_l$ such that $\partial^\omega(G, u_i) \cong \theta_i$ while

$$\left| \left\langle \prod_{i \in [l]} \left(\mathbf{1}\{\partial^\omega(G, u_i, \sigma) \cong (\theta_i, \tau_i)\} - \frac{1}{Z_k(\theta_i)} \right) \right\rangle_G \right| > \varepsilon.$$

Conditional on the convergence assumption in (5), we show that there is a sequence $\varepsilon = \varepsilon(n) = o(1)$ such that w.h.p. (over the graph instances) it holds that $|\mathcal{X}_{\theta_1, \ldots, \theta_l}(\mathbf{G}, [l], \omega)| \leq \varepsilon n^l$. In particular, let $z_i = Z_k(\theta_i)$ and $t_i(v, \sigma) = \mathbf{1}\{\partial^\omega(\mathbf{G}, v, \sigma) \cong (\theta_i, \tau_i)\}$. Moreover, set

$$Q_i(v) = \mathbf{1}\{\partial^\omega(\mathbf{G}, v) \cong \theta_i\} \cdot \left\langle (t_i(v, \sigma_1) - z_i^{-1})(t_i(v, \sigma_2) - z_i^{-1}) \right\rangle_\mathbf{G}, \quad Q_i = \frac{1}{n} \sum_{v \in [n]} Q_i(v).$$

The convergence assumption in (5) implies that there exists $\varepsilon = \varepsilon(n) = o(1)$ such that $\sum_{i \in [l]} Q_i \leq \varepsilon^3$. Then fixing an arbitrary $i_0 \in [l]$ we get that

$$\frac{\varepsilon^2}{n^l} |\mathcal{X}_{\theta_1, \ldots, \theta_l}(\mathbf{G}, [l], \omega)| \leq \frac{1}{n^l} \sum_{u_1, \ldots, u_l \in [n]} \left\langle \prod_{i \in [l]} (t_i(u_i, \sigma) - z_i^{-1}) \right\rangle_\mathbf{G}^2 \prod_{i=1}^{l} \mathbf{1}\{\partial^\omega(\mathbf{G}, u_i) \cong \theta_i\}$$

$$\leq \frac{1}{n^l} \sum_{u_1, \ldots, u_l \in [n]} \left\langle (t_{i_0}(u_{i_0}, \sigma_1) - z_{i_0}^{-1})(t_{i_0}(u_{i_0}, \sigma_2) - z_{i_0}^{-1}) \right\rangle_\mathbf{G} \prod_{i=1}^{l} \mathbf{1}\{\partial^\omega(\mathbf{G}, u_i) \cong \theta_i\}$$

[as σ_1, σ_2 are independent]

$$\leq \frac{1}{n^l} \sum_{u_1, \ldots, u_l \in [n]} Q_{i_0}(u_{i_0}) = Q_{i_0} \leq \varepsilon^3, \tag{7}$$

which implies that w.h.p. $|\mathcal{X}_{\theta_1,\ldots,\theta_l}(\mathbf{G},[l],\omega)| \leq \varepsilon n^l$. Now, we consider the sequences of l vertices which do not belong to $\mathcal{X}_{\theta_1,\ldots,\theta_l}(\mathbf{G},[l],\omega)$, i.e. the majority of the l-tuples. We show that all of them are, somehow, well behaved.

▶ **Claim 7.** *Let $\mathcal{E}_{\theta_1,\ldots,\theta_l}$ be the set of all l-tuples $(v_1,\ldots,v_l)$ of distinct vertices such that $\partial^\omega(\mathbf{G}, v_i) \cong \theta_i$ for all $i \in [l]$. Under the assumption in (5) the following is true: There is a number $C > 0$ such that for all $(v_1,\ldots,v_l) \in \mathcal{E}_{\theta_1,\ldots,\theta_l} \setminus \mathcal{X}_{\theta_1,\ldots,\theta_l}(\mathbf{G},[l],\omega)$*

$$\left| \left\langle \prod_{i \in [l]} \mathbf{1}\left\{ \partial^\omega(\mathbf{G}, v_i, \sigma) \cong (\theta_i, \tau_i) \right\} \right\rangle_{\mathbf{G}} - \prod_{i \in [l]} z_i^{-1} \right| \leq C\varepsilon^{1/2}.$$

Then we get (6) from (7) and Claim 7. The proposition follows.

6 Proof of Proposition 5

Before proving the proposition, we consider the following: Given two maps $\sigma_1, \sigma_2 : [n] \to [k]$, we let the *overlap* matrix $\rho(\sigma_1, \sigma_2)$ be a $k \times k$ matrix such that $\rho_{ij}(\sigma_1, \sigma_2) = \frac{1}{n}\left|\sigma_1^{-1}(i) \cap \sigma_2^{-1}(j)\right|$.

▶ **Claim 8.** *Let $\bar{\rho}$ be the uniform distribution on $[k]^2$. Then, there is $k_0 > 0$ such that for all $k \geq k_0$ and all $d < d_{k,\text{cond}}$, it holds that $\mathbb{E}[\langle\|\rho(\sigma_1, \sigma_2) - \bar{\rho}\|_F\rangle_{\mathbf{G}}] = o(1)$.*

In words, the above claim asserts that for typical instances of $\mathbf{G}$ the expectation over the choice of the random graph $\mathbf{G}$ (the outer $\mathbb{E}$) of the average ℓ_2-distance of the overlap of two randomly chosen k-colorings of $\mathbf{G}$ from $\bar{\rho}$ goes to 0 as $n \to \infty$. The $d < 2(k-1)\ln(k-1)$ case of Claim 8 was previously proved in [33] by way of the second moment analysis from [4]. As it turns out, the regime $2(k-1)\ln(k-1) < d < d_{k,\text{cond}}$ requires a somewhat more sophisticated argument.

In addition to Claim 8, we need the following concentration result from [6].

▶ **Theorem 9** ([6]). *There is $k_0 > 0$ such that for all $k \geq k_0$ and all $d < d_{k,\text{cond}}$ we have*

$$\lim_{\omega \to \infty} \lim_{n \to \infty} \mathbb{P}\left[|\ln Z_k(\mathbf{G}) - \ln \mathbb{E}[Z_k(\mathbf{G})]| \leq \omega\right] = 1.$$

Proof. The proof of the proposition is by contradiction. Assume that $(\mathcal{A}'_n)_{n \geq 1}$ is a sequence of events such that for some fixed number $\varepsilon > 0$ we have

$$\lim_{n \to \infty} \mathbb{P}\left[(\hat{\mathbf{G}}, \hat{\sigma}_1, \hat{\sigma}_2) \in \mathcal{A}'_n\right] = 0 \quad \text{while} \quad \limsup_{n \to \infty} \mathbb{P}\left[(\mathbf{G}, \sigma_1, \sigma_2) \in \mathcal{A}'_n\right] > 2\varepsilon. \tag{8}$$

Setting $\omega(n) = \ln\ln\left(1/\mathbb{P}\left[(\hat{\mathbf{G}}, \hat{\sigma}_1, \hat{\sigma}_2) \in \mathcal{A}'_n\right]\right)$, we let $\mathcal{B}_n$ be the set of all pairs (σ_1, σ_2) of maps $[n] \to [k]$ such that $\|\rho(\sigma_1, \sigma_2) - \bar{\rho}\|_2 \leq \sqrt{\omega/n}$. Also we let

$$\mathcal{A}_n = \left\{(G, \sigma_1, \sigma_2) \in \mathcal{A}'_n : (\sigma_1, \sigma_2) \in \mathcal{B}_n\right\}.$$

We observe that

$$\omega(n) = \ln\ln\left(1/\mathbb{P}\left[(\hat{\mathbf{G}}, \hat{\sigma}_1, \hat{\sigma}_2) \in \mathcal{A}_n\right]\right) \to \infty. \tag{9}$$

Then Claim 8 and (8) imply that

$$\lim_{n \to \infty} \mathbb{P}\left[(\hat{\mathbf{G}}, \hat{\sigma}_1, \hat{\sigma}_2) \in \mathcal{A}_n\right] = 0 \quad \text{while} \quad \limsup_{n \to \infty} \mathbb{P}\left[(\mathbf{G}, \sigma_1, \sigma_2) \in \mathcal{A}_n\right] > \varepsilon.$$

The assumption that $\lim_{n\to\infty} \mathbb{P}\left[(\hat{\mathbf{G}}, \hat{\sigma}_1, \hat{\sigma}_2) \in \mathcal{A}_n\right] = 0$, implies that

$$
\begin{aligned}
\mathbb{E}[Z_k(\mathbf{G})^2 \mathbf{1}\{\mathcal{A}_n\}] &= \sum_{(\sigma_1,\sigma_2)\in\mathcal{B}_n} \mathbb{P}\left[\sigma_1, \sigma_2 \in \mathcal{S}_k(\mathbf{G}), (\mathbf{G}, \sigma_1, \sigma_2) \in \mathcal{A}_n\right] \\
&= \sum_{(\sigma_1,\sigma_2)\in\mathcal{B}_n} \mathbb{P}\left[(\mathbf{G}, \sigma_1, \sigma_2) \in \mathcal{A}_n | \sigma_1, \sigma_2 \in \mathcal{S}_k(\mathbf{G})\right] \mathbb{P}\left[\sigma_1, \sigma_2 \in \mathcal{S}_k(\mathbf{G})\right] \\
&\leq q_n \sum_{(\sigma_1,\sigma_2)\in\mathcal{B}_n} \mathbb{P}\left[(\mathbf{G}, \sigma_1, \sigma_2) \in \mathcal{A}_n | \sigma_1, \sigma_2 \in \mathcal{S}_k(\mathbf{G})\right], \quad (10)
\end{aligned}
$$

where $q_n = \max\{\mathbb{P}\left[\sigma_1, \sigma_2 \in \mathcal{S}_k(\mathbf{G})\right] : (\sigma_1, \sigma_2) \in \mathcal{B}_n\}$. Using the definition of the planted replica model, (10) writes as follows:

$$
\mathbb{E}[Z_k(\mathbf{G})^2 \mathbf{1}\{\mathcal{A}_n\}] \leq k^{2n} q_n \mathbb{P}\left[(\hat{\mathbf{G}}, \hat{\sigma}_1, \hat{\sigma}_2) \in \mathcal{A}_n\right]. \quad (11)
$$

Furthermore, since $\sum_{j=1}^{k} \rho_{ij}(\sigma_1, \sigma_2)^2, \sum_{i=1}^{k} \rho_{ij}(\sigma_1, \sigma_2)^2 \geq 1/k$ for all $i, j \in [k]$, inclusion/exclusion principle implies that

$$
\begin{aligned}
\frac{1}{n} \ln \mathbb{P}\left[\sigma_1, \sigma_2 \in \mathcal{S}_k(\mathbf{G})\right] &\leq \frac{d}{2} \ln\left(1 - \frac{2}{k} + \|\rho(\sigma_1, \sigma_2)\|_2^2\right) + O(1/n) \\
&= d \ln(1 - 1/k) + O(\omega/n) \qquad\qquad \text{for all } (\sigma_1, \sigma_2) \in \mathcal{B}_n.
\end{aligned}
$$

Hence, $q_n \leq (1 - 1/k)^{2m} \exp(O(\omega))$. Plugging this bound into (11) and setting $\bar{z} = \mathbb{E}[Z_k(\mathbf{G})]$, we get that

$$
\begin{aligned}
\mathbb{E}[Z_k(\mathbf{G})^2 \mathbf{1}\{\mathcal{A}_n\}] &\leq k^{2n}(1 - 1/k)^{2m} \exp(O(\omega)) \mathbb{P}\left[(\hat{\mathbf{G}}, \hat{\sigma}_1, \hat{\sigma}_2) \in \mathcal{A}_n\right] \\
&= \bar{z}^2 \exp(O(\omega)) \mathbb{P}\left[(\hat{\mathbf{G}}, \hat{\sigma}_1, \hat{\sigma}_2) \in \mathcal{A}_n\right]. \quad (12)
\end{aligned}
$$

On the other hand, if $\mathbb{P}\left[(\mathbf{G}, \sigma_1, \sigma_2) \in \mathcal{A}_n'\right] > \varepsilon$, then Theorem 9 implies that

$$
\mathbb{P}\left[(\mathbf{G}, \sigma_1, \sigma_2) \in \mathcal{A}_n' \cap \{Z_k(\mathbf{G}) \geq \bar{z}/\omega\}\right] > \varepsilon/2.
$$

Hence, the distribution of the random replica model yields

$$
\mathbb{E}[Z_k(\mathbf{G})^2 \mathbf{1}\{\mathcal{A}_n\}] \geq \frac{\varepsilon}{2}\left(\frac{\bar{z}}{\omega}\right)^2. \quad (13)
$$

But due to (9), (13) contradicts (12). The proposition follows. ◄

—— **References** ——

1   Dimitris Achlioptas and Amin Coja-Oghlan. Algorithmic barriers from phase transitions. In *Foundations of Computer Science, 2008. FOCS'08. IEEE 49th Annual IEEE Symposium on*, pages 793–802. IEEE, 2008.
2   Dimitris Achlioptas and Ehud Friedgut. A sharp threshold for *k*-colorability. *Random Structures Algorithms*, 14(1):63–70, 1999.
3   Dimitris Achlioptas and Michael Molloy. The analysis of a list-coloring algorithm on a random graph. In *Foundations of Computer Science, 1997. Proceedings., 38th Annual Symposium on*, pages 204–212. IEEE, 1997.

4 Dimitris Achlioptas and Assaf Naor. The two possible values of the chromatic number of a random graph. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 587–593. ACM, 2004.

5 Noga Alon and Michael Krivelevich. The concentration of the chromatic number of random graphs. *Combinatorica*, 17(3):303–313, 1997.

6 Victor Bapst, Amin Coja-Oghlan, and Charilaos Efthymiou. Planting colourings silently. *arXiv preprint arXiv:1411.0610*, 2014.

7 Victor Bapst, Amin Coja-Oghlan, Samuel Hetterich, Felicia Raßmann, and Dan Vilenchik. The condensation phase transition in random graph coloring. *arXiv preprint arXiv:1404.5513*, 2014.

8 Nayantara Bhatnagar, Allan Sly, and Prasad Tetali. Decay of correlations for the hardcore model on the *d*-regular random graph. *arXiv preprint arXiv:1405.6160*, 2014.

9 Béla Bollobás. The chromatic number of random graphs. *Combinatorica*, 8(1):49–55, 1988.

10 Amin Coja-Oghlan. Upper-bounding the k-colorability threshold by counting covers. *arXiv preprint arXiv:1305.0177*, 2013.

11 Amin Coja-Oghlan and Dan Vilenchik. Chasing the k-colorability threshold. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 380–389. IEEE, 2013.

12 Martin Dyer and Alan Frieze. Randomly coloring random graphs. *Random Structures & Algorithms*, 36(3):251–272, 2010.

13 Martin Dyer, Alistair Sinclair, Eric Vigoda, and Dror Weitz. Mixing in time and space for lattice spin systems: A combinatorial view. *Random Structures & Algorithms*, 24(4):461–479, 2004.

14 Charilaos Efthymiou. MCMC sampling colourings and independent sets of $G(n, d/n)$ near uniqueness threshold. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 305–316. SIAM, 2014.

15 Charilaos Efthymiou. Reconstruction/non-reconstruction thresholds for colourings of general Galton-watson trees. *arXiv preprint arXiv:1406.3617*, 2014.

16 Charilaos Efthymiou. Switching colouring of $G(n, d/n)$ for sampling up to Gibbs uniqueness threshold. In *Algorithms-ESA 2014*, pages 371–381. Springer, 2014.

17 Paul Erdős. Graph theory and probability. *canad. J. Math*, 11:34G38, 1959.

18 Paul Erdős and A Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci*, 5:17–61, 1960.

19 Andreas Galanis, Qi Ge, Daniel Štefankovič, Eric Vigoda, and Linji Yang. Improved inapproximability results for counting independent sets in the hard-core model. *Random Structures & Algorithms*, 45(1):78–110, 2014.

20 Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability for antiferromagnetic spin systems in the tree non-uniqueness region. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 823–831. ACM, 2014.

21 Hans-Otto Georgii. *Gibbs measures and phase transitions*, volume 9. Walter de Gruyter, 2011.

22 Antoine Gerschenfeld and Andrea Montanari. Reconstruction for models on random graphs. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 194–204. IEEE, 2007.

23 Geoffrey R Grimmett and Colin JH McDiarmid. On colouring random graphs. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 77, pages 313–324. Cambridge Univ Press, 1975.

24 Michael Krivelevich and Benny Sudakov. Coloring random graphs. *Information Processing Letters*, 67(2):71–74, 1998.

25   Florent Krzakała, Andrea Montanari, Federico Ricci-Tersenghi, Guilhem Semerjian, and Lenka Zdeborová. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proceedings of the National Academy of Sciences*, 104(25):10318–10323, 2007.

26   Tomasz Łuczak. The chromatic number of random graphs. *Combinatorica*, 11(1):45–54, 1991.

27   Tomasz Łuczak. A note on the sharp concentration of the chromatic number of random graphs. *Combinatorica*, 11(3):295–297, 1991.

28   Fabio Martinelli, Alistair Sinclair, and Dror Weitz. Fast mixing for independent sets, colorings, and other models on trees. *Random Structures & Algorithms*, 31(2):134–172, 2007.

29   David W Matula. Expose-and-merge exploration and the chromatic number of a random graph. *Combinatorica*, 7(3):275–284, 1987.

30   Marc Mézard and Andrea Montanari. *Information, physics, and computation.* Oxford University Press, 2009.

31   Marc Mézard, Giorgio Parisi, and Riccardo Zecchina. Analytic and algorithmic solution of random satisfiability problems. *Science*, 297(5582):812–815, 2002.

32   Michael Molloy. The freezing threshold for k-colourings of a random graph. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 921–930. ACM, 2012.

33   Andrea Montanari, Ricardo Restrepo, and Prasad Tetali. Reconstruction and clustering in random constraint satisfaction problems. *SIAM Journal on Discrete Mathematics*, 25(2):771–808, 2011.

34   Jaroslav Nešetřil. A combinatorial classic—sparse graphs with high chromatic number. In *Erdős Centennial*, pages 383–407. Springer, 2013.

35   Olivier Rivoire, Giulio Biroli, Olivier C Martin, and Marc Mézard. Glass models on bethe lattices. *The European Physical Journal B-Condensed Matter and Complex Systems*, 37(1):55–78, 2004.

36   Eli Shamir and Joel Spencer. Sharp concentration of the chromatic number on random graphs $G(n, p)$. *Combinatorica*, 7(1):121–129, 1987.

37   Allan Sly. Computational transition at the uniqueness threshold. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 287–296. IEEE, 2010.

38   Yitong Yin and Chihao Zhang. Sampling colorings almost uniformly in sparse random graphs. *arXiv preprint arXiv:1503.03351*, 2015.