

SPATIAL MIXING AND DETERMINISTIC APPROXIMATE COUNTING OF MULTI-SPIN SYSTEMS BEYOND BOUNDED DEGREE GRAPHS

ZHIDAN LI, KUAN YANG

ABSTRACT. We develop a framework for deterministic approximate counting of multi-spin systems beyond bounded-degree graphs. The algorithm recursively constructs rational polytopes containing the true marginal vectors and uses linear-fractional programming to obtain certified bounds on marginal ratios. For positive interactions on graphs of polynomial connective constant D , we establish strong spatial mixing and a fully polynomial-time approximation scheme (**FPTAS**) whenever $Dc < 1$, where c bounds the Birkhoff contraction coefficients of the interactions.

We further extend the framework to proper colorings of sparse Erdős-Rényi random graphs using recursion on permissive blocks. For every fixed $\eta \in (0, 1)$, sufficiently large fixed d , and fixed integer $q \geq (2 + \eta)d$, we obtain an **FPTAS** for counting proper q -colorings of $G \sim \mathcal{G}(n, d/n)$ with high probability over G . This improves the leading constant 3 in the earlier counting guarantee of Yin and Zhang (APPROX/RANDOM, 2016) to 2, and asymptotically matches the spatial mixing regime established by Yin (ICALP, 2014).

1. INTRODUCTION

The multi-spin system is a statistical physics model reflecting the global behavior of local interactions. Their partition functions encode many counting problems in combinatorics, as well as key thermodynamic quantities in statistical mechanics. Exact evaluation of partition functions is $\#\mathbf{P}$ -hard for broad classes of interactions [CC19, CCL13], making efficient approximation algorithms a central object of study in theoretical computer science.

An instance of q -spin systems is a tuple $\Phi = (G = (V, E), [q], \mathbf{A} = (A_e)_{e \in E}, \boldsymbol{\lambda} = (\lambda_v)_{v \in V})$ where $G = (V, E)$ is the underlying graph, the edge interactions $A_e \in \mathbb{R}_{\geq 0}^{q \times q}$ is a non-zero symmetric matrix for every $e \in E$, and $\lambda_v \in \mathbb{R}_{> 0}^q$ is the external field associated with a vertex $v \in V$. The partition function of Φ is defined as

$$\mathcal{Z}(\Phi) = \mathcal{Z}_G(\boldsymbol{\lambda}, \mathbf{A}) := \sum_{\sigma: V \rightarrow [q]} w(\sigma), \quad \text{where } w(\sigma) = \prod_{v \in V} \lambda_v(\sigma(v)) \prod_{e=(u,v) \in E} A_e(\sigma(u), \sigma(v))$$

gives the weight of configuration σ . The corresponding Gibbs distribution μ_Φ is then defined as

$$(1.1) \quad \forall \sigma \in [q]^V, \quad \mu_\Phi(\sigma) = \frac{w(\sigma)}{\mathcal{Z}_G(\boldsymbol{\lambda}, \mathbf{A})}.$$

Correlation decay provides a powerful route to deterministically approximating partition functions. For a self-reducible problem, the partition function can be recovered from a sequence of marginal probabilities. A correlation-decay algorithm estimates these marginals by truncating a suitable recursion and controlling the effect of the resulting boundary error. This approach has led to fully polynomial-time deterministic approximation schemes (**FPTASes**) for the hard-core model [Wei06], matchings [BGK⁺07], colorings and multi-spin systems [GK12, LY13], and counting problems in the Holant framework [YZ13, LWZ14]. An underlying probabilistic property is *strong*

(Zhidan Li) SCHOOL OF COMPUTER SCIENCE, SHANGHAI JIAO TONG UNIVERSITY, SHANGHAI, CHINA. Email: yueEnTeRnAL@sjtu.edu.cn.

(Kuan Yang) JOHN HOPCROFT CENTER FOR COMPUTER SCIENCE, SHANGHAI JIAO TONG UNIVERSITY, SHANGHAI, CHINA. Email: kuan.yang@sjtu.edu.cn.

spatial mixing (SSM), which indicates that the influence of a boundary condition decays with the distance to the vertices where that condition changes.

Unfortunately, when the number of spins q is greater than 2, turning spatial mixing into an efficient algorithm often presents a challenge. To obtain an **FPTAS**, a recursive marginal estimator must control errors for the approximate messages that it actually produces. However, prior analyses over independent coordinate ranges may lose constraints shared by the coordinates of a probability vector, including normalization and model-specific marginal bounds. The resulting contraction condition can be stronger than the condition needed to compare true marginals. In addition, if the underlying graphs are not bounded-degree, vertices of large degree can cause both local error amplification and a large number of recursive calls. The connective-constant approach controls the growth of self-avoiding walks and has allowed correlation-decay methods to handle certain models beyond bounded-degree graphs [SSY13, SSŠY17]. Extending this approach to multi-spin systems requires simultaneous control of the admissible messages and the cost of the recursion. For example, counting proper colorings of sparse random graphs illustrates both difficulties. For $G \sim \mathcal{G}(n, d/n)$, Yin [Yin14] established, with high probability, strong spatial mixing at an arbitrarily fixed vertex when $q \geq \alpha d + \beta$ for every fixed $\alpha > 2$ and a sufficiently large constant β . The deterministic counting method of Yin and Zhang [YZ16] applies in a regime with leading constant 3. Closing this gap requires a recursion that preserves more of the structure of marginal distributions while accommodating vertices of unbounded degree.

In this work, we study strong spatial mixing and deterministic approximate counting for multi-spin systems on graphs with potentially unbounded degree, but with bounded *connective constant*. The connective constant is a quantity associated with the number of self-avoiding walks. Bounded connective constant typically allows for locally dense parts in the underlying graphs, while from the global perspective, such graphs remain similar to those with bounded degree.

We first generalize the result in [SSY13, SSŠY17] to establish strong spatial mixing of multi-spin systems with positive interactions on graphs with bounded connective constant. Then our main result is developing a new framework to design the **FPTAS** based on correlation decay by recursively constructing *rational polytopes containing the true marginal vector*. The polytopes retain constraints among marginal coordinates, and linear-fractional optimization produces certified bounds for the factors in the marginal-ratio recursion. In addition, we apply our algorithm to approximate the number of proper q -colorings on sparse random graphs and improves the best known algorithmic result.

1.1. Main results. Our first result treats multi-spin systems with positive interactions on graphs of bounded connective constant. Writing D for a bound on the growth of self-avoiding walks and c for a uniform bound on the Birkhoff contraction coefficients of the interactions, we establish SSM whenever $Dc < 1$. We then give a fully polynomial-time approximation scheme (**FPTAS**) for the partition function under the same condition. The external fields may vary arbitrarily across vertices, and the graph need not have bounded maximum degree.

Our second result concerns proper colorings of sparse Erdős–Rényi random graphs. For every fixed $\eta \in (0, 1)$, sufficiently large fixed d , and fixed integer $q \geq (2 + \eta)d$, we give a deterministic FPTAS that succeeds with high probability over $G \sim \mathcal{G}(n, d/n)$. This improves the leading constant in the counting guarantee of [YZ16] and asymptotically matches the parameter range of the fixed-vertex spatial-mixing result of [Yin14]¹. Here the high-probability guarantee is over the input graph; the counting algorithm itself is deterministic.

Computation model. Before formally stating our main result, we clarify the computation model here. In this work, we employ the classical *deterministic Turing machine*, which can be described

¹We mention here that spatial mixing in [Yin14] is slightly different: it considers strong spatial mixing only on a *fixed* vertex.

as follows. Rational numbers are encoded by binary bits, and we allow all common bit operations, *e.g.*, arithmetic operations and shifting operations. For an input Φ , the length of Φ is the total number of bits to encode Φ in the model. We use $|\Phi|$ to denote its length. For a single element x in the model, we use $\text{bit}(x)$ to denote the number of bits to encode x . The time complexity of an algorithm is the bit complexity of it, *i.e.*, the number of bit operations used by it.

Now we formally introduce our results. Our first result is on the *spatial mixing* of multi-spin system with all positive interactions. For two probability distributions p, p' on Ω , the total variation distance between p and p' is defined as

$$\|p - p'\|_{\text{TV}} = \frac{1}{2} \sum_{c \in \Omega} |p(c) - p'(c)|.$$

Definition 1.1 (Strong spatial mixing). Fix a graph $G = (V, E)$ and a Gibbs distribution on G . We say μ exhibits *strong spatial mixing* of rate $\rho < 1$ if for two different partial assignments σ, τ on a vertex subset $\Lambda \subseteq V$ and $v \in V \setminus \Lambda$,

$$\|\mu(v \leftarrow \cdot \mid \Lambda \leftarrow \sigma) - \mu(v \leftarrow \cdot \mid \Lambda \leftarrow \tau)\|_{\text{TV}} \leq \text{poly}(|V|) \rho^{\text{dist}_G(v, \Lambda(\sigma, \tau))}$$

where $\Delta(\sigma, \tau) := \{v \in \Lambda \mid \sigma(v) \neq \tau(v)\}$ and $\text{dist}_G(\cdot, \cdot)$ is the distance metric on G .

For a symmetric matrix A with positive entries, define its *Birkhoff projective diameter* as

$$(1.2) \quad \Delta_B(A) := \ln \max_{i, j, k, \ell \in [q]} \frac{A_{ik} A_{j\ell}}{A_{i\ell} A_{jk}}$$

and the *Birkhoff contraction coefficient* as

$$(1.3) \quad c_B(A) := \tanh\left(\frac{\Delta_B(A)}{4}\right).$$

The coefficient $c_B(A)$ measures the contraction of positive vectors under multiplication by A in the Hilbert metric.

We will show that the Gibbs distribution of a multi-spin system exhibits the following spatial mixing property related to the *connective constant*². For convenience, we use the following characterization of connective constants. A self-avoiding walk of $G = (V, E)$ of length ℓ is a simple path $v_0, \dots, v_\ell$ in G . Denote by $\text{SAW}_G(v, \ell)$ the set of self-avoiding walks of length ℓ starting from v , and define

$$S_\ell(G) := \max_{v \in V} |\text{SAW}_G(v, \ell)|$$

with convention $S_0(G) = 1$. We say $\mathcal{G} = (\mathcal{G}_n)_{n \geq 1}$ where $\mathcal{G}_n$ is a family of n -vertex graphs has a *polynomial connective constant* $D \geq 1$ if there are constants $K \geq 1$ and $\alpha \geq 0$ such that

$$(1.4) \quad \forall n \geq 1, \ell > 0, G \in \mathcal{G}_n, \quad S_\ell(G) \leq K n^\alpha D^\ell.$$

Theorem 1.2 (Spatial mixing). Fix an integer $q \geq 3$, and constants $K \geq 1, \alpha > 0, c > 0$ and $D \geq 1$ such that $Dc < 1$. Let $\mathcal{G} = (\mathcal{G}_n)_{n \geq 1}$ satisfy (1.4). For every instance $\Phi = (G, [q], \mathbf{A}, \boldsymbol{\lambda})$ with positive interactions, $G \in \mathcal{G}_n$ and $c_B(A_e) \leq c$ for every $e \in E$, the Gibbs distribution μ_Φ exhibits strong spatial mixing of rate Dc .

The same sufficient condition yields a deterministic approximation scheme to the partition function $Z_G(\boldsymbol{\lambda}, \mathbf{A})$.

²For simplicity to design the algorithm, we use an alternative definition of connective constant in this work. The difference between it and the standard one does no harm on Theorem 1.3.

Theorem 1.3 (Multi-spin systems). *Fix an integer $q \geq 3$, and constants $K \geq 1, \alpha > 0, c > 0$ and $D \geq 1$ such that $Dc < 1$. Let $\mathcal{G} = (\mathcal{G}_n)_{n \geq 1}$ satisfy (1.4). There exists a deterministic algorithm that, given an instance $\Phi = (G, [q], \mathbf{A}, \boldsymbol{\lambda})$ with positive interactions, $G \in \mathcal{G}_n$ and $c_B(A_e) \leq c$ for every $e \in E$, and a binary rational $\varepsilon \in (0, 1)$, it outputs an ε -approximation to $Z_G(\boldsymbol{\lambda}, \mathbf{A})$ in*

$$\left(\frac{n + |\Phi| + \text{bit}(\varepsilon)}{\varepsilon} \right)^C$$

bit operations, where C depends only on q, K, α, D and c .

For the q -state Potts models at edge interaction $\beta > 0$ with any external field $\boldsymbol{\lambda}$, all edges share the matrix A with diagonal entries β and off-diagonal entries 1. Then we have $\Delta_B(A) = 2|\ln \beta|$ and $c_B(A) = \frac{|\beta-1|}{\beta+1}$, which gives the following direct corollary.

Corollary 1.4 (Potts model). *Fix an integer $q \geq 3$, and constants $K \geq 1, \alpha > 0, D \geq 1$ and a rational number β such that $D \frac{|\beta-1|}{\beta+1} < 1$. For every graph family $\mathcal{G} = (\mathcal{G}_n)_{n \geq 1}$ satisfying (1.4), there is a deterministic algorithm that, given a Potts instance Φ on $G \in \mathcal{G}_n$ with positive rational external fields and a binary rational $\varepsilon \in (0, 1)$, outputs an ε -approximation to its partition function in*

$$\left(\frac{n + |\Phi| + \text{bit}(\varepsilon)}{\varepsilon} \right)^C$$

bit operations, where C depends only on q, K, α, D and β .

Remark 1.5. Theorems 1.2 and 1.3 need the bounded connective constant of underlying graphs. We list some amendable graph families here:

- (1) all n -vertex graphs of maximum degree d have a polynomial connective constant $d - 1$;
- (2) for every fixed $d > 1$ and $\delta \in (0, 1)$, with high probability over G drawn from the Erdős-Rényi random graph $\mathcal{G}(n, d/n)$, it has a polynomial connective constant $d(1 + \delta)$;
- (3) some specific lattices own known polynomial connective constants; for example, finite subgraphs of the honeycomb lattice satisfy (1.4) for every fixed $D > \sqrt{2} + \sqrt{2}$, using the exact standard connective constant established in [DCS12].

The second specific spin system is proper colorings on sparse Erdős-Rényi random graphs. A proper q -coloring of $G = (V, E)$ is an assignment $\sigma : V \rightarrow [q]$ satisfying $\sigma(u) \neq \sigma(v)$ on every edge $(u, v) \in E$. Denote by $\Omega(G, q)$ the collection of all proper q -colorings of G . This is the spin system with $\lambda_v(c) = 1$ and $A_e(i, j) = \mathbb{1}[i \neq j]$. The zero entries introduce feasibility constraints, so this application requires an extension of the positive-interaction construction.

Theorem 1.6 (Colorings on random graphs). *For every $\eta \in (0, 1)$, there exists $d_0 = d_0(\eta)$ such that the following holds. For any fixed real $d \geq d_0$ and fixed integer $q \geq (2 + \eta)d$, there are a deterministic algorithm $\mathbf{A} = \mathbf{A}_{\eta, d, q}$ and a positive integer $n_0 = n_0(\eta, d, q)$ with the property: for every $n \geq n_0$, with probability $1 - o_n(1)$ over $G \sim \mathcal{G}(n, d/n)$, the graph is q -colorable and, for any binary rational $\varepsilon \in (0, 1)$, on input G and ε , the algorithm $\mathbf{A}$ outputs an ε -approximation to $|\Omega(G, q)|$ in*

$$\left(\frac{n + \text{bit}(\varepsilon)}{\varepsilon} \right)^C$$

bit operations, where C depends only on η, d and q .

1.2. Technical overview.

Spatial mixing along self-avoiding walks. The marginal-ratio recursion expresses the influence on a vertex through marginals in smaller instances. For positive interactions, the Birkhoff contraction bounds the contribution of each child by its Hilbert distance multiplied by the coefficient of the corresponding edge. Expanding these bounds produces a sum over self-avoiding walks. The coefficients contribute at most c^ℓ along a walk of length ℓ , while (1.4) bounds the number of walks by $Kn^\alpha D^\ell$. Their combined contribution therefore decays at rate $Dc < 1$, without requiring contraction of the sum of influences at every vertex.

Deterministic approximate counting. We design **FPTASes** via estimating marginal probabilities. At a query (Φ, v, t) , the algorithm constructs a rational polytope $\mathcal{P}_t(\Phi, v)$ containing the true marginal vector at v . The polytope enforces normalization and bounds on every pairwise marginal ratio. For a fixed pair of spins a, b , the exact recursion has the form

$$\frac{\mu_\Phi(v \leftarrow a)}{\mu_\Phi(v \leftarrow b)} = \frac{\lambda_v(a)}{\lambda_v(b)} \prod_i \frac{A_i(a, \cdot) p_i}{A_i(b, \cdot) p_i},$$

where each child marginal p_i depends on the pair a, b . Each factor is a linear-fractional function of a single child marginal. Minimizing and maximizing it over the child polytope gives certified bounds, whose products yield linear constraints on the parent marginal vector. This optimization uses all constraints in the child polytope simultaneously.

The error measure is the Hilbert diameter of the feasible polytope. Since the true marginal belongs to the polytope, a small diameter guarantees a multiplicative approximation from any feasible point. For positive interactions, Birkhoff contraction and the self-avoiding-walk bound control the propagation of these diameters. We round every computed lower bound downward and every upper bound upward. This preserves containment and keeps the coefficient encoding lengths polynomial. A logarithmic truncation depth and logarithmic rounding precision suffice for the desired accuracy; a separate bound on the full computation tree establishes polynomial running time.

Permissive blocks of random graph colorings. For proper colorings, a locally proper partial assignment need not extend to a coloring of the graph. We use the permissive blocks of [Yin14, YZ16] to obtain a recursion on blocks instead of vertices, for which local feasibility suffices. The polytope now describes the joint distribution of a block. A lower bound on each feasible block probability ensures positivity, while upper bounds on the single-vertex marginals sharpen the contraction of the ratio factors. Specifically, if the relevant marginal coordinates are at most $1/s$, the Hilbert contraction coefficient is at most $1/(s-1)$.

The random graph analysis has two roles. Weighted self-avoiding-walk estimates control the accumulated error, and bounds on exceptional components and weighted block walks control the total computation cost. The latter is necessary because polynomial work at each block does not alone bound the number of recursive calls. Together these estimates yield Theorem 1.6 at $q \geq (2 + \eta)d$.

1.3. Discussion. We put here some discussions and future directions.

Strong spatial mixing and approximate counting. As mentioned before, strong spatial mixing is closely related to approximate counting while there usually remains a gap between them. When the underlying graph is bounded degree, by [CFG⁺25], it seems that there exists a path from strong spatial mixing to approximate counting via *coupling independence* and marginal lower bounds, which however fails on unbounded-degree graphs. On the other hand, our work provides **FPTASes** for multi-spin systems with positive interactions on graphs with bounded connective constant and proper colorings on sparse Erdős-Rényi random graphs asymptotically matching the current regime of strong spatial mixing on both models respectively. The problem whether strong spatial mixing

implies efficient approximation schemes still remains in-depth discussion. We state the following conjecture on this topic.

Conjecture 1.7. *For multi-spin systems on a family of graphs satisfying some regularities, when the Gibbs distribution exhibits exponentially strong spatial mixing, there is a deterministic approximation schemes to the partition function with running time at most (quasi-)polynomial in the size of input.*

Recovering earlier correlation-decay guarantees. Our algorithm provides a route to retaining the guarantees of earlier recursive correlation-decay algorithms. We expect the framework to recover, in principle, all previous results based on recursive correlation decay. A polynomial-time recovery also requires that the recursion, any change of message coordinates, the enclosure representation, and outward rounding preserve the earlier accuracy and complexity guarantees. The Birkhoff contraction in Hilbert metric used in our positive-interaction theorem is one analysis of the framework; recovering a different result may require its original potential function or other model-specific bounds. The full coverage statement is a methodological expectation, rather than a universal simulation theorem established here. Our coloring result demonstrates a concrete gain beyond the parameter range of the earlier recursive counting method of [YZ16].

Picture of approximate counting. Besides the correlation-decay method, there are different ways to derive **FPTAS** for multi-spin systems: polynomial interpolations from zero-freeness [LSS19, LSS25, BBR26], marginal estimators by derandomizing Markov chain Monte Carlo methods [FGW⁺25], Moitra’s linear-programming methods on coupling independence [CFG⁺25] and algorithmic methods from statistical physics [HPR20, JKP20]. It remains an open problem whether the algorithm proposed in this paper can cover even surpass the results by other techniques.

Deterministic and randomized algorithms. Randomized algorithm, such as Markov chain Monte Carlo, has been widely studied in approximate algorithms. It has been commonly believed that deterministic algorithms could have same capacities as randomized algorithms. A recent work [CFG⁺25] on Moitra’s linear-programming methods on coupling independence provides several deterministic counting results matching sampling results for spin systems on graphs with bounded degree. However, when the bounded-degree assumption is missing, there still remains a gap to reach current results of randomized algorithms, for example, [EHŠV18] shows the rapid mixing of Glauber dynamics of random graph colorings when $q > \alpha d$ where $\alpha \approx 1.7632$ is the root of $\alpha = e^{1/\alpha}$ on $(1, \infty)$ but we are only able to achieve $q \geq (2 + o(1))d$. In spite of the gap, we hope that our algorithm can give some inspiration on this topic.

Turing machine vs. arithmetic models. The computation model here is the deterministic Turing machine. The choice of this model relies on the running time to solve a linear-fractional programming, which is in **P** under Turing machine but remains unknown under the arithmetic models.

1.4. Organization. Section 2 introduces the notation, Hilbert metric, random graphs, and some necessary tools. Spatial mixing Theorem 1.2 for positive interactions is given in Section 3, and our **FPTAS** and its approximation guarantee are provided in Section 4. Section 5 extends our algorithm to proper colorings of sparse random graphs. The proofs of the Hilbert contraction inequalities appear in Appendix A.

1.5. AI disclosure. The main idea of the algorithm framework was provided by authors. The proofs in this paper were generated after supplying ChatGPT Pro 5.6 Sol Ultra with an earlier research draft containing a preliminary algorithmic idea. The agents suggested using the polynomial bit complexity of linear-fractional programming and provided parameter choices in Section 5.3 and some calculations. The authors have verified all proofs and take full responsibility for the paper.

2. PRELIMINARIES

This section fixes basic notations and introduces necessary toolkits and preliminaries in the work.

2.1. Mathematical notations. Let $\mathbb{R}$ be all real numbers, $\mathbb{R}_{>0}$ be all positive real numbers and $\mathbb{N}$ be all natural numbers. We use e to denote the natural logarithm base and $\ln$ to denote the natural logarithm function. For a natural number $n \in \mathbb{N}$, we use $[n]$ to denote the set $\{1, \dots, n\}$.

For a graph G , $V(G)$ denotes its vertex set and $E(G)$ denotes its edge set. When the context is clear, we simply use V and E . We give an order to V and use the alphabetical order to compare paths. The notation $\deg_G(v)$ is the degree of v in G , *i.e.*, the number of neighbors of v in G . For a vertex subset $B \subseteq V$, let $\partial_E B$ be the edge-boundary of B in G . We always list $\partial_E B = \{e_1 = (u_1, v_1), \dots, e_m = (u_m, v_m)\}$ in a fixed order on E that $u_i \in B$ and $v_i \notin B$ for $i = 1, \dots, m$.

Without any additional statement, we define a vector $u \in \mathbb{R}^d$ as a column vector in $\mathbb{R}^{d \times 1}$. For a matrix $A \in \mathbb{R}^{d \times d}$, we use $A(i, \cdot) \in \mathbb{R}^{1 \times d}$ to denote the *row* vector of A at row i . To take the value of an entry in a vector u or a matrix A , the notations $u(i)$ and $A(i, j)$ are employed together with the standard notation u_i .

For an assignment (or coloring) $\sigma : U \rightarrow [q]$ and a subset $S \subseteq U$, we use $\sigma(S)$ to denote the partial assignment of σ on S , *i.e.*, $\sigma(S) = \sigma|_S$ and $\sigma(v)$ to denote $\sigma(\{v\})$ when $S = \{v\}$. For a single element $v \in U$, we also use $\sigma(v)$ to denote the assigned value on v . We use $v \leftarrow c$ to describe the term that v is assigned the value c and use $B \leftarrow \sigma$ to describe the term that the subset B is assigned the partial assignment σ .

2.2. Hilbert distance. We introduce the following metric used by [GK12, GKM15, Yin14, YZ16] to describe the distance of two positive vectors. For two positive vectors p, p' supported on a common state space Ω , define the Hilbert distance between p, p' as

$$(2.1) \quad d_H(p, p') := \max_{a \in \Omega} \ln \frac{p(a)}{p'(a)} - \min_{b \in \Omega} \ln \frac{p(b)}{p'(b)} = \max_{a, b \in \Omega} \ln \frac{p(a)p'(b)}{p'(a)p(b)}.$$

It is trivial to see that any normalization of vectors does not change the Hilbert distance. The following inequality uses the Hilbert metric to bound the total variation distance.

Lemma 2.1. *For two positive probability vectors p, p' supported on a common state space Ω , it holds that*

$$\|p - p'\|_{\text{TV}} \leq \tanh\left(\frac{d_H(p, p')}{4}\right) \leq \frac{d_H(p, p')}{4}.$$

Proof. When $p = p'$, the lemma is trivial. Otherwise, define the ratio vector of p and p' as $R(c) = p(c)/p'(c)$, and set $m = \min_{c \in \Omega} R(c)$ and $M = \max_{c \in \Omega} R(c)$. It is clear that $m \leq 1 \leq M$. We compute the total variation distance as

$$\|p - p'\|_{\text{TV}} = \sum_{c \in \Omega} \max\{p(c) - p'(c), 0\} = \sum_{c \in \Omega} p'(c) \max\{R(c) - 1, 0\}.$$

By the secant-line bound

$$\max\{0, x - 1\} \leq \frac{M - x}{M - m} \max\{0, m - 1\} + \frac{x - m}{M - m} \max\{0, M - 1\} = \frac{(M - 1)(x - m)}{M - m},$$

it holds that

$$\|p - p'\|_{\text{TV}} \leq \sum_{c \in \Omega} p'(c) \frac{(M - 1)(R(c) - m)}{M - m} = \frac{M - 1}{M - m} (1 - m).$$

Observe that $M/m = e^t$ where $t = d_H(p, p')$. By the AM-GM inequality $e^t m + \frac{1}{m} \geq 2e^{t/2}$, we have

$$\|p - p'\|_{\text{TV}} \leq \frac{e^t + 1 - (e^t m + 1/m)}{e^t - 1} \leq \frac{e^t - 2e^{t/2} + 1}{e^t - 1} = \frac{e^{t/2} - 1}{e^{t/2} + 1} = \tanh(t/4).$$

This proves the first inequality. For the second inequality, just observe that $\tanh(x) \leq x$ on $[0, \infty)$ and we conclude it. $\square$

The following Birkhoff contraction for the Hilbert distance plays a crucial role. It is actually a standard finite-dimensional argument in [Bir57] and for completeness, we include its proof in Appendix A.

Lemma 2.2. *For every positive symmetric matrix A and positive vectors x, y , it holds that*

$$d_H(Ax, Ay) \leq c_B(A) d_H(x, y).$$

2.3. Self-reducibility. In the design of approximate algorithms, self-reducibility is an important property. [JVV86] shows that for a family of self-reducible instances, it suffices to estimate the marginal probability to approximate partition functions. For simplicity, we list some self-reducible instances here instead of introducing the precise concept.

- Multi-spin systems on graphs with bounded degree or with bounded connective constant are self-reducible.
- Although the proper q -coloring model is not self-reducible, the proper *list-coloring* model is self-reducible.

2.4. Rounding length and rounding oracle. To maintain the bit complexity, we introduce the rounding length and rounding oracles. Let $P \geq 3$ be a picked rounding length and set $\xi_P = 2^{2-P}$ be the rounding accuracy. We define two oracles $\text{round}_P^-(x)$ and $\text{round}_P^+(x)$ on $(0, \infty)$ as following. Let s be the unique integer satisfying $2^s \leq x < 2^{s+1}$ and set $h_P(x) = 2^{s-P+1}$. We define $\text{round}_P^-(x) = h_P(x) \lfloor \frac{x}{h_P(x)} \rfloor$ and $\text{round}_P^+(x) = h_P(x) \lceil \frac{x}{h_P(x)} \rceil$. It holds that the bit complexity of computing $\text{round}_P^+(x)$ and $\text{round}_P^-(x)$ is at most $O((\text{bit}(x) + P)^2)$. Furthermore, assume that $x \in [\ell, u]$ where ℓ and u can be encoded by at most L bits. Then it takes at most $O(L + P)$ bits to encode the rounded rational numbers $\text{round}_P^-(x)$ and $\text{round}_P^+(x)$. Meanwhile $0 < \text{round}_P^-(x) \leq x \leq \text{round}_P^+(x)$ and

$$(2.2) \quad \max \left\{ \ln \frac{x}{\text{round}_P^-(x)}, \ln \frac{\text{round}_P^+(x)}{x} \right\} \leq \xi_P.$$

2.5. Linear-fractional programming. Now we state some standard results for linear-fractional programmings. A linear-fractional programming is of the form

$$(2.3) \quad \begin{aligned} & \max/\min \frac{\mathbf{c}^\top \mathbf{x} + \alpha}{\mathbf{d}^\top \mathbf{x} + \beta} \\ & \text{subject to } A\mathbf{x} \leq \mathbf{b}. \end{aligned}$$

where $\mathbf{x}$ represents the variables of the system, $\mathbf{c}, \mathbf{d} \in \mathbb{R}^n$ and $\alpha, \beta \in \mathbb{R}$ are given parameters about the optimization, and $A \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$ are given parameters to describe the constraints.

The Charnes-Cooper transformation [CC62] shows that a linear-fractional programming can be converted to a linear programming when the denominator $\mathbf{d}^\top \mathbf{x} + \beta$ is positive uniformly over the feasible region.

Theorem 2.3 (Charnes-Cooper transformation [CC62]). *For an instance Φ of form (2.3) representing a bounded rational polytope, if $\mathbf{d}^\top \mathbf{x} + \beta > 0$ for every feasible $\mathbf{x}$, then it can be transformed to a linear programming Φ' of encoding length $O(|\Phi|)$. As a result, a rational optimum can be solved in $O(n^3 m^3 |\Phi|^3)$ bit operations.*

2.6. Erdős-Rényi random graphs. The Erdős-Rényi graph $G = (V, E) \sim \mathcal{G}(n, p)$ on n vertices of rate $p \in [0, 1]$ is a random graph generated as:

- $V = [n]$; and
- for every $(i, j) \in \binom{[n]}{2}$, E includes it with probability p independently.

The following lemma gives an algorithm to find a proper q -coloring on random graphs.

Lemma 2.4. *For every $d \geq 1$ and $q > \max d, 3$, with probability $1 - o_n(1)$ over $G \sim \mathcal{G}(n, d/n)$, G is q -colorable and a proper q -coloring on G can be output in $\text{poly}(n, q)$ bit operations.*

Proof. By [PSW96], with probability $1 - o_n(1)$, G has no q -core, i.e., an induced subgraph of minimal degree at least q . Consider the following algorithm:

- (1) initialize $H = G$, $S = \emptyset$ and $Q = \{v \in V(H) : \deg_H(v) \leq q - 1\}$;
- (2) repeat the following process until $Q = \emptyset$:
 - (a) pick a smallest vertex $v \in Q$;
 - (b) append v to the tail of S and remove v from Q ;
 - (c) remove v from H and delete all incident edges to v ;
 - (d) $Q \leftarrow Q \cup \{u \in V(H) : \deg_H(u) \leq q - 1\}$;
- (3) from the tail of S to the head, color the vertex with the smallest color which is not used by assigned neighbors.

When G has no q -core, it holds that when $Q = \emptyset$, H is empty and thus S contains all vertices in $V(G)$. The feasibility of the coloring comes directly from the order to construct S and the total bit operations is at most $O(|V(G)| + |E(G)| + q(|V(G)| + |E(G)|) \ln q) = \text{poly}(n, q)$. $\square$

3. SPATIAL MIXING OF MULTI-SPIN SYSTEMS

In this section, we show the strong spatial mixing of multi-spin systems Theorem 1.2. The key ingredient is the *ratio recursion* of spin systems, which has been widely used in the analysis of proper coloring models [GK12, GKM15].

To describe the recursion, we first introduce the concept of pinnings.

Definition 3.1. Given an instance $\Phi = (G = (V, E), [q], \mathbf{A} = (A_e)_{e \in E}, \boldsymbol{\lambda} = (\lambda_v)_{v \in V})$ and a vertex $v \in V$, assume that its neighbors are $v_1, \dots, v_d$ where $d = \deg_G(v)$. For two distinct colors $a, b \in [q]$ and $i \in [d]$, define the child instance $\Phi_{i,a,b}^v = (G', [q], \mathbf{A}', \boldsymbol{\lambda}')$ as:

- the underlying graph $G' = G[V \setminus \{v\}]$ is the induced subgraph of G by removing v ;
- $\mathbf{A}' = (A_e)_{e \in E \setminus \{(v, v_1), \dots, (v, v_d)\}}$;
- $\boldsymbol{\lambda}' = (\lambda'_u)_{u \in V \setminus \{v\}}$ is defined as: for $j = 1, \dots, i - 1$, $\lambda'_{v_j}(c) = A_{(v, v_j)}(a, c) \lambda_{v_j}(c)$; for $j = i + 1, \dots, d$, $\lambda'_{v_j}(c) = A_{(v, v_j)}(b, c) \lambda_{v_j}(c)$; and $\lambda'_u = \lambda_u$ for remaining vertices.

Let $p_{i,a,b}^v$ be the marginal vector of v_i in $\Phi_{i,a,b}^v$, i.e., $p_{i,a,b}^v(c) = \mu_{\Phi_{i,a,b}^v}(v_i \leftarrow c)$ for every $c \in [q]$. When the context is clear, we simply use Φ_i to denote $\Phi_{i,a,b}^v$ and p_i to denote $p_{i,a,b}^v$.

Based on Definition 3.1, we have the following recursion for the marginal ratio, which plays a central role in the proof of Theorems 1.2 and 1.3.

Lemma 3.2. *Given an instance $\Phi = (G = (V, E), [q], \mathbf{A}, \boldsymbol{\lambda})$, for every vertex $v \in V$ and $a, b \in [q]$, it holds that*

$$\frac{\mu_\Phi(v \leftarrow a)}{\mu_\Phi(v \leftarrow b)} = \frac{\lambda_v(a)}{\lambda_v(b)} \prod_{i=1}^d \frac{\sum_{c \in [q]} A_{(v, v_i)}(a, c) \mu_{\Phi_i}(v_i \leftarrow c)}{\sum_{c \in [q]} A_{(v, v_i)}(b, c) \mu_{\Phi_i}(v_i \leftarrow c)}.$$

Proof. The proof is similar to [LY13, Proposition 1]. For simplicity, we use A_i to denote $A_{(v,v_i)}$. For a partial assignment π on $V \setminus \{v\}$, define

$$W(\pi) := \prod_{u \in V \setminus \{v\}} \lambda_u(\pi(u)) \prod_{(u,w) \in E \setminus \{(v,v_1), \dots, (v,v_d)\}} A_{(u,w)}(\pi(u), \pi(w)).$$

For every $c \in [q]$, let

$$F(c) := \sum_{\pi: V \setminus \{v\} \rightarrow [q]} W(\pi) \prod_{i=1}^d A_{(v,v_i)}(c, \pi(v_i)).$$

By (1.1),

$$\frac{\mu_\Phi(v \leftarrow a)}{\mu_\Phi(v \leftarrow b)} = \frac{\lambda_v(a)F(a)}{\lambda_v(b)F(b)}.$$

To telescope $F(a)/F(b)$, for $i = 0, \dots, d$, define

$$(3.1) \quad F_i := \sum_{\pi: V \setminus \{v\} \rightarrow [q]} W(\pi) \prod_{j=1}^i A_{(v,v_j)}(a, \pi(v_j)) \prod_{j=i+1}^d A_{(v,v_j)}(b, \pi(v_j)).$$

Then it holds that $F(a)/F(b) = \prod_{i=1}^d F_i/F_{i-1}$ since $F(a) = F_d$ and $F(b) = F_0$. For a fixed $i = 1, \dots, d$, write

$$W_i(\pi) = W(\pi) \prod_{j=1}^{i-1} A_{(v,v_j)}(a, \pi(v_j)) \prod_{j=i+1}^d A_{(v,v_j)}(b, \pi(v_j)).$$

Observe that the partition function of Φ_i is equal to

$$\mathcal{Z}(\Phi_i) = \sum_{\pi: V \setminus \{v\} \rightarrow [q]} W_i(\pi)$$

and for the marginal vector p_i on v_i ,

$$\forall c \in [q], \quad p_i(c) = \frac{1}{\mathcal{Z}(\Phi_i)} \sum_{\pi: V \setminus \{v\} \rightarrow [q], \pi(v_i)=c} W_i(\pi).$$

Then we have

$$\begin{aligned} F_i &= \sum_{\pi: V \setminus \{v\} \rightarrow [q]} A_{(v,v_i)}(a, \pi(v_i)) W_i(\pi) \\ &= \mathcal{Z}(\Phi_i) \sum_{c \in [q]} A_{(v,v_i)}(a, c) p_i(c) \end{aligned}$$

and similarly $F_{i-1} = \mathcal{Z}(\Phi_i) \sum_{c \in [q]} A_{(v,v_i)}(b, c) p_i(c)$. Therefore

$$\frac{F_i}{F_{i-1}} = \frac{\sum_{c \in [q]} A_{(v,v_i)}(a, c) p_i(c)}{\sum_{c \in [q]} A_{(v,v_i)}(b, c) p_i(c)}.$$

Plugging it into (3.1), we conclude the lemma. $\square$

Proof of Theorem 1.2. Fix the boundary set Λ and two different partial assignments σ, τ on Λ . Let $\Gamma = \Delta(\sigma, \tau)$, i.e., the discrepancy of σ and τ . Set $\overline{\Delta} = 4 \operatorname{arctanh} c$ and $\rho = Dc$. For convenience, denote by $H_\Phi(v, \Lambda, \sigma, \tau)$ the Hilbert distance $d_H(\mu_\Phi(v \leftarrow \cdot \mid \Lambda \leftarrow \sigma), \mu_\Phi(v \leftarrow \cdot \mid \Lambda \leftarrow \tau))$.

We will show that

$$(3.2) \quad H_\Phi(v, \Lambda, \sigma, \tau) \leq \frac{Kn^\alpha D \bar{\Delta}}{1 - \rho} \rho^{\text{dist}_G(v, \Gamma) - 1}$$

and therefore by Lemma 2.1, we conclude the theorem.

We first introduce *local activities* on $V \setminus \Lambda$ after a pinning π on Λ . Let

$$\forall x \in V \setminus \Lambda, c \in [q], \quad \lambda_x^\pi(c) = \lambda_x(c) \prod_{y \in \Lambda \cap N_G(x)} A_{(x,y)}(c, \pi(y)).$$

The local Hilbert discrepancy is induced by local activities conditional on different pinnings:

$$H^{\text{local}}(x) := d_H(\lambda_x^\sigma, \lambda_x^\tau).$$

By direct calculation on $d_H(\lambda_x^\sigma, \lambda_x^\tau)$, it holds that

$$H^{\text{local}}(x) \leq \sum_{y \in \Gamma \cap N_G(x)} \Delta_B(A_{(x,y)}).$$

We remark here $H^{\text{local}}(x) = 0$ unless x is adjacent to Γ .

To show (3.2), we specify a sub-instance $\mathcal{I}$ by a vertex subset $S = S_{\mathcal{I}}$ and an external field $\mathbf{h} = (h_x : [q] \rightarrow \mathbb{R}_{>0})_{x \in S}$ with the original edge interactions restricted on $G[S]$. We say a pair of instances $(\mathcal{I} = (S, \mathbf{h}), \mathcal{I}' = (S, \mathbf{h}'))$ are *admissible* if $S \subseteq V \setminus \Lambda$ and for every $x \in S$,

$$\forall c \in [q], \quad \frac{h_x(c)}{h'_x(c)} = \frac{\lambda_x^\sigma(c)}{\lambda_x^\tau(c)}.$$

Consider the Hilbert distance $d_H(\mu_{\mathcal{I}}(x \leftarrow \cdot), \mu_{\mathcal{I}'}(x \leftarrow \cdot))$. Assume that $x_1, \dots, x_d$ be the neighbors of x in $G[S]$. For two positive probability p, p' and $a, b \in [q]$, by Lemma 2.2,

$$\begin{aligned} \left| \ln \frac{\sum_{c \in [q]} A_{(x, x_i)}(a, c) p(c)}{\sum_{c \in [q]} A_{(x, x_i)}(b, c) p(c)} - \ln \frac{\sum_{c \in [q]} A_{(x, x_i)}(a, c) p'(c)}{\sum_{c \in [q]} A_{(x, x_i)}(b, c) p'(c)} \right| &\leq d_H(A_{(x, x_i)} p, A_{(x, x_i)} p') \\ &\leq c_B(A_{(x, x_i)}) d_H(p, p'). \end{aligned}$$

Apply Lemma 3.2, and we obtain that

$$d_H(\mu_{\mathcal{I}}(x \leftarrow \cdot), \mu_{\mathcal{I}'}(x \leftarrow \cdot)) \leq H^{\text{local}}(x) + \max_{a, b \in [q]} \sum_{i=1}^d c_B(A_{(x, x_i)}) d_H(\mu_{\mathcal{I}_i}(x_i \leftarrow \cdot), \mu_{\mathcal{I}'_i}(x_i \leftarrow \cdot))$$

where $\mathcal{I}_i, \mathcal{I}'_i$ are generated by Definition 3.1 from $\mathcal{I}, \mathcal{I}'$ respectively.

For a vertex subset S and a vertex $x \in S$, introduce the function $\text{INF}(S, x)$ in the following recursive form:

$$(3.3) \quad \text{INF}(S, x) = H^{\text{local}}(x) + \sum_{i=1}^d c_B(A_{(x, x_i)}) \text{INF}(S \setminus \{x\}, x_i)$$

with boundary condition $\text{INF}(S, x) = H^{\text{local}}(x)$ for $S = \{x\}$ or x is isolated in the graph $G[S]$. By induction, $d_H(\mu_{\mathcal{I}}(x \leftarrow \cdot), \mu_{\mathcal{I}'}(x \leftarrow \cdot))$ is upper bounded by $\text{INF}(S, x)$ for every pair of admissible $(\mathcal{I}, \mathcal{I}')$. In addition, note that the instance Φ with pinning π on Λ is specified by $V \setminus \Lambda$ and $\lambda^\pi = (\lambda_x^\pi)_{x \in V \setminus \Lambda}$. Therefore, it holds that

$$H_\Phi(v, \Lambda, \sigma, \tau) \leq \text{INF}(V \setminus \Lambda, v).$$

Unfold (3.3) and substitute the bound of $H^{\text{local}}(x)$. Every non-zero term is encoded by a self-avoiding walk $P = (v = x_0, \dots, x_{\ell-1})$ in $G[V \setminus \Lambda]$ and a vertex $y \in \Gamma$ incident to $x_{\ell-1}$. Append y

to the tail of P . Therefore P is a self-avoiding walk from v to Γ of length ℓ with contribution at most $c^{\ell-1} \cdot \bar{\Delta}$. Hence,

$$\begin{aligned} \text{INF}(V \setminus \Lambda, v) &\leq \sum_{\ell=\text{dist}_G(v, \Gamma)}^{n-1} c^{\ell-1} \cdot \bar{\Delta} \cdot |\text{SAW}_G(v, \ell)| \\ &\leq Kn^\alpha D \cdot \bar{\Delta} \sum_{\ell \geq \text{dist}_G(v, \Gamma)-1} \rho^\ell \\ &= \frac{Kn^\alpha D \bar{\Delta}}{1 - \rho} \rho^{\text{dist}_G(v, \Gamma)-1}. \end{aligned}$$

Combining all things together, we conclude the theorem. $\square$

4. APPROXIMATE COUNTING MULTI-SPIN SYSTEMS WITH POSITIVE INTERACTIONS

In this section, we design the algorithm in Theorem 1.3. Based on the recursion in Section 3, we show how to recursively construct rational polytopes which approximately estimate the true marginal vector and analyze them when the instance is defined on a graph of bounded connective constant in Section 4.1 and we derive the approximation scheme to the partition function by polytopes in Section 4.2..

4.1. Rational polytopes induced by the recursion. Recall Lemma 3.2. Consider a query $(\Phi = (G = (V, E), [q], \mathbf{A}, \mathbf{\lambda}), v, t)$ for an instance Φ , a vertex $v \in V$ and an integer $t \geq 0$ representing the recursive depth. We recursively construct a rational polytope $\mathcal{P} = \mathcal{P}_t(\Phi, v)$ to estimate the marginal vector $\mu_\Phi(v \leftarrow \cdot)$ as Definition 4.1.

Definition 4.1 (Rational polytope for spin systems). The variables of $\mathcal{P}_t(\Phi, v)$ is $(z_c)_{c \in [q]}$. When v is an isolated vertex in G , $\mathcal{P}_t(\Phi, v)$ is the exact polytope

$$\forall c \in [q], \quad z_c = \frac{\lambda_v(c)}{\sum_{k \in [q]} \lambda_v(k)}.$$

When v is not isolated, assume that its neighbors in G are $v_1, \dots, v_d$ where $d = \deg_G(v)$. The polytope $\mathcal{P}_t(\Phi, v)$ at first contains:

- **Non-negativity.**

$$(4.1) \quad \forall c \in [q], \quad z_c \geq 0.$$

- **Normalization.**

$$(4.2) \quad \sum_{c \in [q]} z_c = 1.$$

We separate cases based on the value of the depth t .

- $t = 0$. For every distinct spins $a, b \in [q]$, the polytope contains

$$(4.3) \quad L_{a,b}^0 \cdot z_b \leq z_a \leq U_{a,b}^0 \cdot z_b$$

where

$$L_{a,b}^0 := \frac{\lambda_v(a)}{\lambda_v(b)} \prod_{i=1}^d \min_{c \in [q]} \frac{A_{(v, v_i)}(a, c)}{A_{(v, v_i)}(b, c)}, \quad U_{a,b}^0 := \frac{\lambda_v(a)}{\lambda_v(b)} \prod_{i=1}^d \max_{c \in [q]} \frac{A_{(v, v_i)}(a, c)}{A_{(v, v_i)}(b, c)}.$$

- $t > 0$. For every distinct spins $a, b \in [q]$ and $i = 1, \dots, d$, construct an instance $\Phi_i = \Phi_{i,a,b}^v$. Recursively obtain the polytope $\mathcal{P}_i = \mathcal{P}_{t-1}(\Phi_i, v_i)$. Solve the following two LFPs:

$$\ell_{i,a,b}^* := \min_{z \in \mathcal{P}_i} \frac{\sum_{c \in [q]} A_{(v,v_i)}(a, c) z_c}{\sum_{c \in [q]} A_{(v,v_i)}(b, c) z_c}, \quad u_{i,a,b}^* := \max_{z \in \mathcal{P}_i} \frac{\sum_{c \in [q]} A_{(v,v_i)}(a, c) z_c}{\sum_{c \in [q]} A_{(v,v_i)}(b, c) z_c}.$$

and set $\ell_{i,a,b} = \text{round}_P^-(\ell_{i,a,b}^*)$, $\bar{u}_{i,a,b} = \text{round}_P^+(u_{i,a,b}^*)$. Setting $L_{a,b}^t = \frac{\lambda_v(a)}{\lambda_v(b)} \prod_{i=1}^d \ell_{i,a,b}$ and $U_{a,b}^t = \frac{\lambda_v(a)}{\lambda_v(b)} \prod_{i=1}^d \bar{u}_{i,a,b}$, the polytope then contains

$$(4.4) \quad L_{a,b}^t \cdot z_b \leq z_a \leq U_{a,b}^t \cdot z_b.$$

Fact 4.2 (Well-definedness and soundness of polytopes). *At every depth the recursive construction is well-defined, the true marginal vector at v belongs to $\mathcal{P}_t(\Phi, v)$ and the bit complexity to encode $\mathcal{P}_t(\Phi, v)$ is at most $O(q^3(|\Phi| + nP))$.*

Proof. We prove the fact by induction. When the polytope contains only an exact point (the case v is isolated), the fact is trivial. Otherwise, at depth 0, it is direct to see the true marginal vector p_v at v satisfies constraints (4.1) and (4.2). By Lemma 3.2, it is easy to verify p_v satisfies (4.3). The bit complexity to encode $\mathcal{P}_0(\Phi, v)$ is at most $O(q^3|\Phi|)$ since we have q variables and $O(q^2)$ constraints with each constraint encoded with at most $O(|\Phi|)$ bits.

Suppose that the assertion holds through depth $t - 1$. It is trivial to see the true marginal vector p_v satisfies (4.1) and (4.2). Since the true marginal vector p_i belongs to $\mathcal{P}_i$ by induction, by Lemma 3.2, p_v also satisfies the constraint (4.4) even after applying rounding oracles. Therefore we conclude that $p_v \in \mathcal{P}_t(\Phi, v)$. To compute the number of bits to encode $\mathcal{P}_t(\Phi, v)$, note that we have q variables and $O(q^2)$ linear constraints. For a constraint (4.4), note that for each $i = 1, \dots, d$, it holds that

$$\min_{c \in [q]} \frac{A_{(v,v_i)}(a, c)}{A_{(v,v_i)}(b, c)} \leq \ell_{i,a,b}^* \leq u_{i,a,b}^* \leq \max_{c \in [q]} \frac{A_{(v,v_i)}(a, c)}{A_{(v,v_i)}(b, c)}.$$

Hence, it takes $O(\text{bit}(A_{(v,v_i)}) + P)$ bits to encode the rounded rational number $\ell_{i,a,b}$, $\bar{u}_{i,a,b}$ and thus $O(\text{bit}(\lambda_v) + dP + \sum_{i=1}^d \text{bit}(A_{(v,v_i)}))$ bits to encode $L_{a,b}^t$ and $U_{a,b}^t$. Therefore, the bit complexity to encode $\mathcal{P}_t(\Phi, v)$ is at most $O(q^3(|\Phi| + nP))$. $\square$

For an admissible polytope $\mathcal{P}$, we use $\text{diam}_H(\mathcal{P})$ to denote the Hilbert diameter of $\mathcal{P}$, i.e.,

$$\text{diam}_H(\mathcal{P}) := \max_{x, y \in \mathcal{P}} d_H(x, y).$$

Based on the Birkhoff contraction Lemma 2.2, the following recursion is on the control of the accuracy in $\mathcal{P}_t(\Phi, v)$ by its child polytopes.

Proposition 4.3. *Fix an admissible feasible query (Φ, v, t) . When v is isolated, $\text{diam}_H(\mathcal{P}_t(\Phi, v)) = 0$. Otherwise, assume that its neighbors are $v_1, \dots, v_d$. It holds that*

$$\text{diam}_H(\mathcal{P}_t(\Phi, v)) \leq \begin{cases} \sum_{i=1}^d \Delta_B(A_{(v,v_i)}), & t = 0 \\ \max_{a, b \in [q]} \sum_{i=1}^d \left(c_B(A_{(v,v_i)}) \text{diam}_H(\mathcal{P}_{t-1}(\Phi_{i,a,b}^v, v_i)) + 2\xi_P \right), & t > 0 \end{cases}.$$

Proof. When v is isolated, $\mathcal{P}_t(\Phi, v)$ contains only a single point and thus $\text{diam}_H(\mathcal{P}_t(\Phi, v)) = 0$. Suppose that v is not isolated and its neighbors are $v_1, \dots, v_d$. For $t = 0$, it holds that

$$\ln \frac{U_{a,b}^0}{L_{a,b}^0} = \sum_{i=1}^d \left(\ln \max_{c, d \in [q]} \frac{A_{(v,v_i)}(a, c) A_{(v,v_i)}(b, d)}{A_{(v,v_i)}(a, d) A_{(v,v_i)}(b, c)} \right) \leq \sum_{i=1}^d \Delta_B(A_{(v,v_i)}).$$

Then for every $x, y \in \mathcal{P}_0(\Phi, v)$, it holds that

$$d_H(x, y) = \max_{a, b \in [q]} \ln \frac{x(a)y(b)}{x(b)y(a)} \leq \max_{a, b \in [q]} \ln \frac{U_{a,b}^0}{L_{a,b}^0} \leq \sum_{i=1}^d \Delta_B(A_{(v, v_i)}).$$

For $t > 0$, for every $i = 1, \dots, d$, by Lemma 2.2,

$$\begin{aligned} \ln \frac{u_{i,a,b}^*}{\ell_{i,a,b}^*} &\leq \max_{x, y \in \mathcal{P}_{t-1}(\Phi_i, v_i)} d_H(A_{(v, v_i)}x, A_{(v, v_i)}y) \\ &\leq c_B(A_{(v, v_i)}) \text{diam}_H(\mathcal{P}_{t-1}(\Phi_i, v_i)). \end{aligned}$$

Together with (2.2), it holds that

$$\ln \frac{U_{a,b}^t}{L_{a,b}^t} \leq \sum_{i=1}^d (c_B(A_{(v, v_i)}) \text{diam}_H(\mathcal{P}_{t-1}(\Phi_{i,a,b}^v, v_i)) + 2\xi_P).$$

A similar argument to the case $t = 0$ shows

$$\text{diam}_H(\mathcal{P}_t(\Phi, v)) \leq \max_{a, b \in [q]} \ln \frac{U_{a,b}^t}{L_{a,b}^t}$$

and thus we conclude the full proposition. $\square$

The following bound of accuracy is a corollary of Proposition 4.3.

Corollary 4.4. *Given an instance $\Phi = (G = (V, E), [q], \mathbf{A}, \boldsymbol{\lambda})$ such that G satisfies (1.4) and $D \cdot c_B(A_e) < 1$, for every integer $R \geq 0$ and $v \in V$, it holds that*

$$(4.5) \quad \text{diam}_H(\mathcal{P}_R(\Phi, v)) \leq Kn^\alpha D \left(\overline{\Delta} \rho^R + \frac{2\xi_P}{1-\rho} \right)$$

where $c = \max_{e \in E} c_B(A_e)$, $\overline{\Delta} = 4 \arctanh c$ and $\rho = D \cdot c$. Moreover, set

$$(4.6) \quad R = R_\varepsilon = \left\lceil \frac{\ln(4Kn^\alpha D \max\{1, \overline{\Delta}\} / \varepsilon)}{\ln(1/\rho)} \right\rceil, \quad P = P_\varepsilon = \max \left\{ 3, \left\lceil \ln_2 \frac{32Kn^\alpha D}{\varepsilon(1-\rho)} \right\rceil \right\}.$$

For every feasible point $z \in \mathcal{P}_R(\Phi, v)$, it holds that for every $c \in [q]$,

$$e^{-\varepsilon} \mu_\Phi(v \leftarrow c) \leq z_c \leq e^\varepsilon \mu_\Phi(v \leftarrow c).$$

Proof. The proof is similar to Theorem 1.2. Fix an instance $\Phi = (G = (V, E), [q], \mathbf{A}, \boldsymbol{\lambda})$, a vertex $v \in V$ and a depth $R \geq 0$. First we have

$$\Delta_B(A_e) = 4 \arctanh c_B(A_e) \leq 4 \arctanh c = \overline{\Delta}.$$

Consider the recursion in Proposition 4.3. We unfold the recursion into a rooted tree $\mathcal{T}$ with root ϱ . A node ω in $\mathcal{T}$ at level $0 \leq r \leq R$ is labelled by a derived instance Φ_ω on the underlying graph G_ω , a vertex v_ω and a remaining depth $R - r$. At every node ω at level $r < R$, we choose any pair of (a_ω, b_ω) to attain the maximum in the recursion. Then for each vertex u incident to v in the instance Φ_ω , there is a child node θ of ω in $\mathcal{T}$ of level $r + 1$ with label corresponding to the recursive query. Assign to the edge (ω, θ) the value

$$\kappa_{\omega, \theta} = c_B(A_{(v, u)}) \leq c.$$

Now we define the weight function F on $\mathcal{T}$ as $F(\varrho) = 1$ and

$$\forall \omega \neq \varrho, \quad F(\omega) = \prod_e \kappa_e$$

where the product takes over all edges in the unique path from ϱ to ω . Using the above bound for κ_e , we have $F(\omega) \leq c^r$ for every node ω at level r . By Proposition 4.3, it holds that

$$\begin{aligned} \text{diam}_H(\mathcal{P}_R(\Phi, v)) &\leq \sum_{\omega \in \mathcal{T} \text{ at level } R} F(\omega) \sum_{u \text{ is a neighbor of } v_\omega \text{ in } \Phi_\omega} \Delta_B(A_{(v_\omega, u)}) \\ &\quad + 2\xi_P \sum_{r=0}^{R-1} \sum_{\omega \in \mathcal{T} \text{ at level } r} F(\omega) \deg_{G_\omega}(v_\omega). \end{aligned}$$

On the other hand, observe that there is an one-to-one correspondence between nodes in $\mathcal{T}$ to all self-avoiding walks from v . That is to say, a node $\omega \in \mathcal{T}$ at level r corresponds to a unique self-avoiding walk from v of length r . Therefore,

$$\begin{aligned} \sum_{\omega \in \mathcal{T} \text{ at level } R} F(\omega) \sum_{u \text{ is a neighbor of } v_\omega \text{ in } \Phi_\omega} \Delta_B(A_{(v_\omega, u)}) &\leq \bar{\Delta} \cdot c^R |\text{SAW}_G(v, R+1)| \\ &\leq \bar{\Delta} \cdot c^R S_{R+1}(G) \\ &\leq Kn^\alpha D \bar{\Delta} (D \cdot c)^R \\ &= Kn^\alpha D \bar{\Delta} \rho^R \end{aligned}$$

and

$$\begin{aligned} 2\xi_P \sum_{r=0}^{R-1} \sum_{\omega \in \mathcal{T} \text{ at level } r} F(\omega) \deg_{G_\omega}(v_\omega) &\leq 2\xi_P \sum_{r=0}^{R-1} c^r |\text{SAW}_G(v, r+1)| \\ &\leq 2\xi_P \sum_{r=0}^{R-1} c^r \cdot Kn^\alpha D^{r+1} \\ &\leq 2Kn^\alpha D \xi_P \sum_{r=0}^{\infty} (D \cdot c)^r \\ &= Kn^\alpha D \cdot \frac{2\xi_P}{1-\rho}. \end{aligned}$$

Combining all things together, we conclude (4.5).

For the second argument, substitute the choice of R and P in (4.5) and we obtain it. $\square$

4.2. Approximation scheme to the partition function. Now we are ready to show Theorem 1.3.

Proof of Theorem 1.3. For a fixed assignment $\sigma : V \rightarrow [q]$, pin vertices in a fixed order $v_1, \dots, v_n$. Let p_j be the true probability of $\sigma(v_j)$ and $\Phi^{(j)}$ be the j -th derived instance. Then

$$\mathcal{Z}_G(\boldsymbol{\lambda}, \mathbf{A}) = w(\sigma) \prod_{j=1}^n p_j^{-1}.$$

Set $\varepsilon' = \varepsilon/(32n)$ and choose $R = R_{\varepsilon'}, P = P_{\varepsilon'}$ as (4.6). For every $j = 1, \dots, n$, construct the polytope $\mathcal{P}_R(\Phi^{(j)}, v_j)$. Pick $\hat{p}_j$ to minimize the value of entry $\sigma(v_j)$ in $\mathcal{P}_R(\Phi^{(j)}, v_j)$. By Fact 4.2 and Corollary 4.4, it holds that

$$\forall j = 1, \dots, n, \quad e^{-\varepsilon/(32n)} \leq \hat{p}_j/p_j \leq 1.$$

Write $\hat{Z} = w(\sigma) \prod_{j=1}^n \hat{p}_j^{-1}$. We know that $1 \leq \hat{Z}/\mathcal{Z}_G(\boldsymbol{\lambda}, \mathbf{A}) \leq e^{\varepsilon/32} \leq 1 + \varepsilon$.

Now we consider the bit complexity. When $R \geq n$, it must hold that $\varepsilon = e^{-\Omega(n)}$. Then we exactly compute $\mathcal{Z}_G(\boldsymbol{\lambda}, \mathbf{A})$ by brute force enumerating all assignments and the bit complexity is bounded

by desired number. Otherwise, note that $R = O_{K,\alpha,D,c}(\ln(n/\varepsilon))$ and $P = O_{K,\alpha,D,c}(\ln(n/\varepsilon))$. The number of polytopes is at most

$$O\left(n \sum_{r=0}^R q^{2r} S_r(G)\right) = O(Kn^{\alpha+1}(R+1)(q^2D)^R) \leq (n/\varepsilon)^{C_1(q,K,\alpha,D,c)}.$$

By Theorem 2.3 and Fact 4.2, for every polytope, it can be encoded by $O(q^3(|\Phi| + n \ln(n/\varepsilon)))$ bits and takes at most $O(q^{20}(|\Phi| + n \ln(n/\varepsilon))^3)$ bit operations to solve induced LFPs. Therefore, the total bit complexity is at most

$$O\left((n/\varepsilon)^{C_1(q,K,\alpha,D,c)} \cdot q^{20}(|\Phi| + n \ln(n/\varepsilon))^3\right) \leq \left(\frac{n + |\Phi| + \text{bit}(\varepsilon)}{\varepsilon}\right)^C$$

for some constant $C > 0$ depending only on q, K, α, D and c . $\square$

Proof of Corollary 1.4. It is a direct result of Theorem 1.3 when $A_e = A$ for every $e \in E$ where $A(i, i) = \beta$ and $A(i, j) = 1$ for $i \neq j$. $\square$

5. PROPER COLORINGS ON SPARSE RANDOM GRAPHS

In this section, we prove Theorem 1.6. Fix $\eta \in (0, 1)$. Throughout this section, $d \geq d_0(\eta)$ and the integer $q \geq (2 + \eta)d$ are fixed, where $d_0(\eta)$ is supplied by Lemma B.1. All asymptotic statements are for $n \rightarrow \infty$ with η, d, q fixed.

We first construct the marginal recursion on permissive blocks and bound the Hilbert diameter of the resulting polytopes. We then establish random graph estimates that control both the accumulated error and the total cost of the recursion. The degree thresholds and auxiliary parameters are introduced where they enter the analysis.

5.1. List-coloring instances and permissive blocks. As mentioned in Section 2.3, proper q -coloring model is not self-reducible. Instead, we consider the list-coloring instances as [Yin14, YZ16].

A list-coloring instance Φ consists of an underlying graph G_Φ and color lists L_Φ such that $L_\Phi(v) \subseteq [q]$ collects feasible colors on v . An assignment $\sigma : V(G_\Phi) \rightarrow [q]$ is a proper list-coloring in Φ if and only if it is a proper q -coloring and for every vertex $v \in V(G_\Phi)$, $\sigma(v) \in L_\Phi(v)$. Let $\Omega(\Phi)$ be the collection of all proper list-colorings in Φ and set $\mathcal{Z}(\Phi) = |\Omega(\Phi)|$. With abuse of notation, for a vertex subset $S \subseteq G_\Phi$, we use Ω_B to denote all proper partial list-coloring on B ³.

Throughout the construction, $G = (V, E)$ denotes the original graph. Define

$$(5.1) \quad X := \{v \in V \mid \deg_G(v) \geq q - 1\}.$$

This set is used to construct permissive blocks. It remains fixed when vertices are removed or colors are deleted from the lists.

Then we introduce the following marginal upper bound for $v \notin X$ which has been established in [Yin14].

Lemma 5.1 ([Yin14]). *For a feasible list-coloring instance Φ and every v such that $|L_\Phi(v)| > \deg_{G_\Phi}(v)$, it holds that*

$$\forall c \in L_\Phi(v), \quad \mu_\Phi(v \leftarrow c) \leq \frac{1}{|L_\Phi(v)| - \deg_{G_\Phi}(v)}.$$

For brevity, set

$$\forall v \notin X, \quad s_v := q - \deg_G(v), \quad \kappa_v := \frac{1}{q - \deg_G(v) - 1}$$

where G is the original graph.

³A proper partial list-coloring on B here is not necessarily able to extend a proper list-coloring in Φ .

Now we introduce the concept of *permissive blocks* in [Yin14, YZ16]. Given a list-coloring instance Φ , we say a connected vertex subset $B \subseteq V(G_\Phi)$ is a *permissive block* for every $(u_i, v_i) \in \partial_E B$, $|L_\Phi(v_i)| > \deg_{G_\Phi}(v_i) + 1$. It is direct to see that a permissive block B preserves permissiveness under pinning feasible partial colorings.

We state here the following property of permissive blocks.

Proposition 5.2. *For a feasible list-coloring instance Φ and a permissive block $B \subseteq V(G_\Phi)$, every partial proper list-coloring on B extends to a proper list-coloring in Φ . Furthermore, for every proper list-coloring σ on B , it holds that*

$$\mu_\Phi(B \leftarrow \sigma) \geq b_B := 2^{-|\partial_E B|} q^{-|B|}.$$

Proof. The feasibility of a feasible partial coloring on B directly comes from [Yin14, Lemma 3.4]. For the marginal lower bound, fix a feasible partial coloring σ on B . Construct intermediate instances $\Psi_0, \dots, \Psi_{|\partial_E B|}$ as: Ψ_0 is the instance derived by removing B from G_Φ , $L_\Phi(v)$ from color lists for every $v \in B$ and keeping all remaining color lists unchanged; for every $j = 1, \dots, |\partial_E B|$, generate Ψ_j from Ψ_{j-1} by removing $\sigma(u_j)$ from the color list on v_j . For any intermediate instance Ψ , suppose that we need to delete color c from the list of v to obtain the instance Ψ' . Since B is a permissive block, by Lemma 5.1,

$$\mu_\Psi(v \leftarrow c) \leq \frac{1}{2}.$$

Therefore,

$$\begin{aligned} \mathcal{Z}(\Psi') &= \mathcal{Z}(\Psi) - |\{\sigma \text{ is a proper list-coloring in } \Psi : \sigma(v) = c\}| \\ &= \mathcal{Z}(\Psi) - \mu_\Psi(v \leftarrow c) \cdot \mathcal{Z}(\Psi) \\ &\geq \mathcal{Z}(\Psi)/2. \end{aligned}$$

Observe that $\mathcal{Z}(\Phi) \leq |\Omega_B| \cdot \mathcal{Z}(\Psi_0)$ and $\mathcal{Z}(\Psi_{\partial_E B}) = \mathcal{Z}(\Phi \mid B \leftarrow \sigma)$. Therefore,

$$\mathcal{Z}(\Phi \mid B \leftarrow \sigma) \geq 2^{-|\partial_E B|} \cdot \mathcal{Z}(\Psi_0) \geq 2^{-|\partial_E B|} q^{-|B|} \mathcal{Z}(\Phi).$$

Then we conclude the marginal lower bound. $\square$

For every $\sigma, \tau \in \Omega_B$ and $i = 1, \dots, m := |\partial_E B|$, define the intermediate instance $\Phi_{i,\sigma,\tau}$ be the instance derived from Φ by:

- the underlying graph is the induced subgraph of G_Φ by $V(G_\Phi) \setminus B$;
- the color lists are obtained by deleting $\sigma(u_j)$ from $L_\Phi(v_j)$ for $j = 1, \dots, i-1$, deleting $\tau(u_j)$ from $L_\Phi(v_j)$ for $j = i+1, \dots, m$ and keeping unchanged for other lists.

Then we have the following recursion of probability ratios.

Lemma 5.3 ([Yin14]). *For every $\sigma, \tau \in \Omega_B$, it holds that*

$$(5.2) \quad \frac{\mu_\Phi(B \leftarrow \sigma)}{\mu_\Phi(B \leftarrow \tau)} = \prod_{i=1}^m \frac{1 - \mu_{\Phi_{i,\sigma,\tau}}(v_i \leftarrow \sigma(u_i))}{1 - \mu_{\Phi_{i,\sigma,\tau}}(v_i \leftarrow \tau(u_i))}.$$

We mention here that by Lemma 5.1, it is clear that every denominator in the right side of (5.2) is at least $1/2$. Also, by [Yin14], for every intermediate list-coloring instance Φ generated from pinning a feasible partial coloring on a vertex subset in G , it holds that Φ is feasible and

$$(5.3) \quad |L_\Phi(v)| \geq q - (\deg_G(v) - \deg_{G_\Phi}(v)).$$

Thus by Lemma 5.12, for $v \in V(G_\Phi) \setminus X$, $|L_\Phi(v)| - \deg_{G_\Phi}(v) \geq s_v$ and the marginal upper bound is at most $1/s_v$. Also it is trivial to see that Proposition 5.2 and Lemma 5.3 hold for Φ .

5.2. Rational polytope for list-coloring instances. Denote by $\mathcal{C}(H, v)$ the maximal connected component containing v in H . For every $v \in V(G_\Phi)$, we define the permissive block of v in Φ as

$$(5.4) \quad B_\Phi(v) = \begin{cases} \mathcal{C}(G[X], v) \cap V(G_\Phi), & v \in X \\ \{v\} \cup (\bigcup_{C \sim v} (C \cap V(G_\Phi))), & v \notin X \end{cases}$$

where $C \sim v$ enumerates all connected components in $G[X]$ containing a neighbor of v in G_Φ . It is clear that $B_\Phi(v)$ is connected in G_Φ . Given a q -coloring instance Φ_0 on a graph G , we call a list-coloring instance Φ *accessible* if it can be obtained by repeatedly applying the recursion in Lemma 5.3, starting from an instance obtained by pinning a feasible partial coloring of Φ_0 . Now we show how to construct the rational polytope $\mathcal{P}_t(\Phi, B)$ with an accessible list-coloring instance $\Phi = (G, L_\Phi)$, a permissive block B of Φ and a depth t .

Definition 5.4 (Rational polytope for list colorings). The variables of $\mathcal{P}_t(\Phi, B)$ is $(z_\sigma)_{\sigma \in \Omega_B}$. When $\partial_E B = \emptyset$, $\mathcal{P}_t(\Phi, B)$ is the exact uniform polytope

$$\forall \sigma \in \Omega_B, \quad z_\sigma = \frac{1}{|\Omega_B|}.$$

When $\partial_E B \neq \emptyset$, list $\partial_E B$ in order $(u_1, v_1), \dots, (u_m, v_m)$. The polytope $\mathcal{P}_t(\Phi, B)$ at first contains:

- **Marginal lower bound.**

$$(5.5) \quad \forall \sigma \in \Omega_B, \quad z_\sigma \geq b_B = 2^{-m} q^{-|B|}.$$

- **Normalization.**

$$(5.6) \quad \sum_{\sigma \in \Omega_B} z_\sigma = 1.$$

- **Marginal upper bound.**

$$(5.7) \quad \forall v \in B \setminus X, c \in L_\Phi(v), \quad \sum_{\sigma \in \Omega_B: \sigma(v)=c} z_\sigma \leq \frac{1}{s_v}.$$

We separate cases based on the value of the depth t .

- $t = 0$. For every distinct partial colorings $\sigma, \tau \in \Omega_B$, the polytope contains

$$(5.8) \quad 2^{-m} z_\tau \leq z_\sigma \leq 2^m z_\tau.$$

- $t > 0$. For every distinct partial colorings $\sigma, \tau \in \Omega_B$ and $i = 1, \dots, m$, construct an instance $\Phi_i = \Phi_{i, \sigma, \tau}$. Recursively obtain the polytope $\mathcal{P}_i = \mathcal{P}_{t-1}(\Phi_i, B_{\Phi_i}(v_i))$. For a probability vector $y \in \mathcal{P}_i$, we define

$$p_c(y) := \sum_{\pi: \pi(v_i)=c} y_\pi$$

with $p_c(y) = 0$ if c is absent. Solve the following two LFPs:

$$\ell_{i, \sigma, \tau}^* := \min_{y \in \mathcal{P}_i} \frac{1 - p_{\sigma(u_i)}(y)}{1 - p_{\tau(u_i)}(y)}, \quad u_{i, \sigma, \tau}^* := \max_{y \in \mathcal{P}_i} \frac{1 - p_{\sigma(u_i)}(y)}{1 - p_{\tau(u_i)}(y)}.$$

and set $\underline{\ell}_{i, \sigma, \tau} = \text{round}_P^-(\ell_{i, \sigma, \tau}^*)$, $\bar{u}_{i, \sigma, \tau} = \text{round}_P^+(u_{i, \sigma, \tau}^*)$. Setting $L_{\sigma, \tau}^t = \prod_{i=1}^m \underline{\ell}_{i, \sigma, \tau}$ and $U_{\sigma, \tau}^t = \prod_{i=1}^m \bar{u}_{i, \sigma, \tau}$, the polytope then contains

$$(5.9) \quad L_{\sigma, \tau}^t \cdot z_\tau \leq z_\sigma \leq U_{\sigma, \tau}^t \cdot z_\tau.$$

Fact 5.5 (Well-definedness and soundness). *At every depth the recursive construction is well-defined, the true marginal vector on B belongs to $\mathcal{P}_t(\Phi, B)$ and the bit complexity to encode $\mathcal{P}_t(\Phi, B)$ is at most $O(q^{4|B|} \cdot (|B| + |\partial_E B| \cdot P))$.*

Proof. The proof is similar to Fact 4.2. When the polytope contains only an exact point (the case $\partial_E B = \emptyset$), the fact is trivial. Otherwise, at depth 0, it is direct to see the true marginal vector p_B on B satisfies constraints (5.5), (5.6) and (5.7). By Lemma 5.3, it is easy to verify p_B satisfies (5.8) since every denominator is at least $1/2$ and every enumerator is at most 1. The bit complexity to encode $\mathcal{P}_0(\Phi, B)$ is at most $O(q^{4|B|}(|B| + |\partial_E B|))$ since we have $q^{|B|}$ variables and $O(q^{2|B|})$ constraints with each constraint encoded with at most $O(q^{2|B|}(|B| + |\partial_E B|))$ bits.

Suppose that the assertion holds through depth $t - 1$. It is trivial to see the true marginal vector p_B satisfies (5.5), (5.6) and (5.7). By induction and Lemma 5.3, it is not hard to see p_B also satisfies the constraint (5.9) even after applying rounding oracles. Therefore we conclude that $p_B \in \mathcal{P}_t(\Phi, B)$. To compute the number of bits to encode $\mathcal{P}_t(\Phi, B)$, note that we have $q^{|B|}$ variables and $O(q^{2|B|})$ linear constraints. For a constraint (5.9), note that for each $i = 1, \dots, m$, it holds that

$$\frac{1}{2} \leq \ell_{i,\sigma,\tau}^* \leq u_{i,\sigma,\tau}^* \leq 2.$$

Hence, it takes $O(P)$ bits to encode the rounded rational number $\ell_{i,\sigma,\tau}, \bar{u}_{i,\sigma,\tau}$ and thus $O(|\partial_E B| \cdot P)$ bits to encode $L_{\sigma,\tau}^t$ and $U_{\sigma,\tau}^t$. Therefore, the bit complexity to encode $\mathcal{P}_t(\Phi, B)$ is at most $O(q^{4|B|}(|B| + |\partial_E B| \cdot P))$. $\square$

Recall that in Section 4, for an admissible rational polytope $\mathcal{P}$, we use the Hilbert distance to denote the diameter of $\mathcal{P}$:

$$\text{diam}_H(\mathcal{P}) = \max_{x,y \in \mathcal{P}} d_H(x,y).$$

The following lemma for the contraction of Hilbert distance is the key ingredient which is implicitly stated in [GK12, GKM15, Yin14].

Lemma 5.6. *For two positive probability vectors p, p' supported on a same space Ω such that every coordinate of both is at most $1/s$ for $s \geq 2$, then*

$$\left| \ln \frac{1-p_a}{1-p'_a} - \ln \frac{1-p_b}{1-p'_b} \right| \leq \frac{d_H(p, p')}{s-1}.$$

A clear proof of Lemma 5.6 is included in Appendix A for completeness. With Lemma 5.6, the following recursion controls the accuracy of $\mathcal{P}_t(\Phi, B)$ by its child polytopes.

Proposition 5.7. *Fix an admissible feasible query (Φ, B, t) . When $\partial_E B = \emptyset$, $\text{diam}_H(\mathcal{P}_t(\Phi, B)) = 0$. Otherwise, list $\partial_E B$ in a fixed order $(u_1, v_1), \dots, (u_m, v_m)$. It holds that*

$$\text{diam}_H(\mathcal{P}_t(\Phi, B)) \leq \begin{cases} 2m \ln 2, & t = 0 \\ \max_{\sigma, \tau \in \Omega_B} \sum_{i=1}^m (\kappa_{v_i} \cdot \text{diam}_H(\mathcal{P}_{t-1}(\Phi_{i,\sigma,\tau}, B_{\Phi_{i,\sigma,\tau}}(v_i))) + 2\xi_P), & t > 0 \end{cases}.$$

Proof. When $\partial_E B = \emptyset$, the rational polytope $\mathcal{P}_t(\Phi, B)$ contains exactly one feasible point representing the uniform distribution on Ω_B . Therefore, $\text{diam}_H(\mathcal{P}_t(\Phi, B)) = 0$.

Suppose that $\partial_E B = \{(u_1, v_1), \dots, (u_m, v_m)\}$ for some $m \geq 1$. For $t = 0$, by (5.8), for every $\sigma, \tau \in \Omega_B$,

$$-m \ln 2 \leq \ln \frac{z_\sigma}{z_\tau} \leq m \ln 2.$$

Thus $\text{diam}_H(\mathcal{P}_0(\Phi, B)) \leq 2m \ln 2$.

At positive depth $t > 0$, fix a pair of $\sigma, \tau \in \Omega_B$. For $i = 1, \dots, m$, consider the child rational polytope $\mathcal{P}_i = \mathcal{P}_{t-1}(\Phi_{i,\sigma,\tau}, B_{\Phi_{i,\sigma,\tau}}(v_i))$. Since $v_i \notin X$ by the permissiveness of B , for every $y \in \mathcal{P}_i$,

by (5.7),

$$p_c(y) \leq \frac{1}{s_{v_i}}.$$

Hence by Lemma 5.6, for every $y, y' \in \mathcal{P}_i$,

$$\left| \ln \frac{1 - p_{\sigma(u_i)}(y)}{1 - p_{\tau(u_i)}(y)} - \ln \frac{1 - p_{\sigma(u_i)}(y')}{1 - p_{\tau(u_i)}(y')} \right| \leq \frac{d_H(p(y), p(y'))}{s_{v_i} - 1} \leq \kappa_{v_i} d_H(y, y')$$

where the last inequality holds since $p(\cdot)$ is a weighted average operator. Together with the rounding accuracy (2.2),

$$\ln \frac{U_{\sigma, \tau}^t}{L_{\sigma, \tau}^t} \leq \sum_{i=1}^m (\kappa_{v_i} \text{diam}_H(\mathcal{P}_i) + 2\xi_P).$$

Thus a similar argument to the proof of Proposition 4.3 shows that

$$\text{diam}_H(\mathcal{P}_t(\Phi, B)) \leq \max_{\sigma, \tau \in \Omega_B} \ln \frac{U_{\sigma, \tau}^t}{L_{\sigma, \tau}^t},$$

which concludes the proposition. $\square$

5.3. Random graph properties. The random graph estimates below serve two purposes. Decay of weighted self-avoiding walks controls the accumulated truncation error, while bounds on exceptional components and weighted block walks control the block sizes and the total recursive cost.

To choose a degree cutoff D , we balance the rarity of vertices above the cutoff against contraction below it. For a vertex of degree at most D , the contraction coefficient κ_v is at most $1/(q - D - 1)$. Ignoring integer rounding and additive constants, the two requirements are

$$D > d \quad \text{and} \quad \frac{d}{q - D} < 1,$$

or equivalently $d < D < q - d$. This balance explains the role of $q > 2d$ in the present analysis. We choose

$$(5.10) \quad D := \left\lfloor \frac{d + q}{3} \right\rfloor, \quad \gamma := \frac{1}{q - D - 1},$$

and define the exceptional set for the probabilistic estimates by

$$(5.11) \quad H_D := \{v \in V \mid \deg_G(v) > D\}.$$

For the degree tail estimates, write

$$h(x) := (1 + x) \ln(1 + x) - x \quad (x > 0),$$

and set

$$(5.12) \quad \begin{aligned} \delta &:= \frac{q/d - 2}{6}, \quad r := e^{-dh(\delta)}, \\ \lambda_0 &:= \frac{3 + \eta/2}{3 + \eta}, \quad \rho := \frac{1 + \lambda_0}{2}. \end{aligned}$$

Here r bounds the conditional probability that a vertex is exceptional when at most three incident edges have been prescribed. The constants λ_0 and ρ are the decay rates used in the expectation and high-probability bounds, respectively.

By Lemma B.1, these parameters satisfy

$$(5.13) \quad D - 3 > (1 + \delta)d, \quad D < q - 1,$$

$$(5.14) \quad d(\gamma + \sqrt{r}) \leq \lambda_0 < \rho < 1,$$

$$(5.15) \quad edr^{1/4} \leq \frac{1}{4}.$$

The first estimate leaves room for up to three prescribed incident edges in the degree tail bounds. The second yields decay of path weights, and the third controls the sizes of exceptional components. In particular, $X \subseteq H_D$. The set X determines the permissive blocks, whereas H_D is used to bound their sizes and to control the propagation of error.

5.3.1. *Path weights.* For a graph G , we define the weight function $w : V(G) \rightarrow \mathbb{R}$ as

$$\forall v \in V(G), \quad w(v) = \begin{cases} \gamma & \deg_G(v) \leq D \\ 1 & \deg_G(v) > D \end{cases}.$$

For $\ell \geq 1$, define

$$W_\ell(G) := \sum_{v_0, \dots, v_\ell} \prod_{i=1}^{\ell} w(v_i)$$

where the sum takes over all simple paths of length ℓ in G . We have the following exponential decay of $W_\ell(G)$ with high probability over the random graph G .

Lemma 5.8. *For every positive integer $n \geq 2 + 12/\eta$, it holds that*

$$\mathbb{P}_{G \sim \mathcal{G}(n, d/n)} \left(\exists \ell < n : W_\ell > n^6 \rho^\ell \right) \leq n^{-4}.$$

Proof. Fix $\ell \geq 1$ and a simple path $P = v_0, \dots, v_\ell$ of length ℓ . Conditional on the event $\mathcal{E}_P = \{P \text{ is a simple path in } G\}$, for every $v \in V$, write $\deg_G(v) = f_v + X_v$ where $f_v = 1$ if $v = v_0$ or $v = v_\ell$, $f_v = 2$ if $v \in \{v_1, \dots, v_{\ell-1}\}$ and $f_v = 0$ otherwise. Note that X_v is a random variable dominated by a binomial random variable $Y \sim \text{Bin}(n, d/n)$. For the event $\deg_G(v) > D$, it must hold that

$$X_v = \deg_G(v) - f_v > D - 2 > (1 + \delta)d$$

where the last inequality follows from (5.13). Therefore by the Chernoff bound,

$$\mathbb{P}(v \in H_D \mid \mathcal{E}_P) \leq \mathbb{P}(Y \geq (1 + \delta)d) \leq e^{-dh(\delta)} = r.$$

Now we turn to the random variables $\{\mathbb{1}[v \in H_D]\}_{v \in V}$. Observe that, in spite of dependence, the random variables depend on independent random variables $\{\mathbb{1}[(u, v) \in E]\}_{(u, v) \in \binom{V}{2}}$. By Finner's generalized Hölder inequality for a read-two family,

$$\mathbf{E} \left[\prod_{v \in S} \mathbb{1}[v \in H_D] \right] \leq \prod_{v \in S} \mathbf{E} \left[\mathbb{1}[v \in H_D]^2 \right]^{1/2} \leq r^{|S|/2}$$

Since $w(v) \leq \gamma + \mathbb{1}[v \in H_D]$,

$$\begin{aligned} \mathbf{E} \left[\prod_{i=1}^{\ell} w(v_i) \mid \mathcal{E}_P \right] &\leq \sum_{S \subseteq [\ell]} \gamma^{\ell - |S|} \mathbb{P}(\forall i \in S, v_i \in H_D \mid \mathcal{E}_P) \\ &\leq \sum_{i=0}^{\ell} \binom{\ell}{i} \gamma^{\ell-i} r^{i/2} \\ &\leq (\gamma + \sqrt{r})^{\ell}. \end{aligned}$$

By the law of total expectation and (5.14),

$$\mathbf{E}[W_{\ell}(G)] \leq nd^{\ell}(\gamma + \sqrt{r})^{\ell} \leq n\lambda_0^{\ell}.$$

Then by the Markov inequality and the union bound,

$$\mathbb{P}\left(\exists \ell < n, W_{\ell}(G) > n^6 \rho^{\ell}\right) \leq n \cdot n^{-6} \frac{\lambda_0/\rho}{1 - \lambda_0/\rho} \leq n^{-4}$$

for every $n \geq \frac{\lambda_0}{\rho - \lambda_0} = 2 + 12/\eta$. □

5.3.2. Exceptional component size and incident volume. For the induced subgraph $G[H_D]$, we first consider the size of each component in it.

Lemma 5.9. *With probability at least $1 - \frac{1}{2}n^{-4}$ over $G \sim \mathcal{G}(n, d/n)$, every connected component in $G[H_D]$ has size at most $4 \ln n$.*

Proof. For a component C of $G[H_D]$ of size t , take a spanning tree T of C . Conditional on T in G , at least half of $V(T)$ has tree degree at most 3. For each of these vertices, (5.13) gives the same conditional degree tail bound r . Applying the conditional read-two inequality as in the proof of Lemma 5.8 gives

$$\mathbb{P}(T \subseteq H_D \mid T \text{ in } G) \leq r^{t/4}.$$

Then by the union bound and the Cayley's formula,

$$\begin{aligned} \mathbb{P}(\exists \text{ a component } C \text{ in } H_D \text{ of size } \geq t) &\leq \binom{n}{t} t^{t-2} \left(\frac{d}{n}\right)^{t-1} r^{t/4} \\ &\leq \frac{e^t n^t}{t^t} t^{t-2} \cdot \frac{d^{t-1}}{n^{t-1}} \cdot r^{t/4} \\ &\leq \frac{n}{d} \left(edr^{1/4}\right)^t. \end{aligned}$$

Using (5.15) and taking $t = \lceil 4 \ln n \rceil$ gives the claimed failure probability. □

The second result is about the incident volume of a component in $G[H_D]$.

Lemma 5.10. *Assume the event of Lemma 5.9. With probability at least $1 - \frac{1}{2}n^{-4}$, for every component C in $G[H_D]$,*

$$\sum_{v \in C} \deg_G(v) \leq C_{\text{vol}} \ln n$$

where $C_d = 2 + d(e - 1) + d(e^2 - 1)/4$ and $C_{\text{vol}} = 4C_d + 4 \ln d + 14$.

Proof. Fix a vertex subset S of size t and let T be a tree on S . Define a random variable $Y_S := \sum_{v \in S} \deg_G(v)$. Conditional on the event $\mathcal{E}_T$ that T is in G , it holds that

$$Y_S = 2(t-1) + \sum_{e \in S \times (V \setminus S)} X_e + 2 \sum_{e \in \binom{S}{2} \setminus E(T)} X_e$$

where $\{X_e = \mathbb{1}[e \in E(G)]\}_{e \in \binom{V}{2}}$ are independent identical binomial random variables with rate $p = d/n$. For every $t \leq n/2$,

$$\begin{aligned} \mathbf{E}[e^{Y_S} \mid \mathcal{E}_T] &= e^{2(t-1)} \cdot (1 + p(e-1))^{t(n-t)} \cdot (1 + p(e^2-1))^{\binom{t}{2} - (t-1)} \\ &\leq \exp\left(2(t-1) + p(e-1)t(n-t) + p(e^2-1)\left(\binom{t}{2} - (t-1)\right)\right) \\ &\leq \exp(2t + t \cdot d(e-1) + t \cdot d(e^2-1)/4) = e^{C_d \cdot t}. \end{aligned}$$

Conditional on Lemma 5.9, by the Markov inequality and the union bound,

$$\mathbb{P}(\exists \text{ a component } C \text{ in } G[H_D] \text{ of size } t, Y_C \geq C_{\text{vol}} \ln n) \leq \binom{n}{t} t^{t-2} \left(\frac{d}{n}\right)^{t-1} e^{C_d t - C_{\text{vol}} \ln n}.$$

The crude inequality $\binom{n}{t} \leq (en/t)^t$ gives that this failure probability is at most

$$\frac{e^t n^t}{t^t} t^{t-2} \frac{d^{t-1}}{n^{t-1}} e^{C_d t - C_{\text{vol}} \ln n} \leq \frac{n}{d} e^{-C_{\text{vol}} \ln n} (de^{C_d+1})^t.$$

Therefore,

$$\begin{aligned} \mathbb{P}(\exists \text{ a maximal component } C \text{ in } G[H_D], Y_C \geq C_{\text{vol}} \ln n) &\leq \frac{n}{d} \cdot e^{-C_{\text{vol}} \ln n} \sum_{t=1}^{4 \ln n} (de^{C_d+1})^t \\ &\leq \frac{n}{d} \cdot e^{-C_{\text{vol}} \ln n} (4 \ln n) (de^{C_d+1})^{4 \ln n} \\ &\leq n^{-5}. \end{aligned}$$

Thus we conclude the lemma. $\square$

5.3.3. Permissive block walks. We next bound the total weight of permissive block walks. The fixed exponent K below absorbs the color-pair choices and the local optimization cost associated with each block; see (5.19). Set

$$(5.16) \quad K := 200, \quad J := (q-1) \ln 2 - d, \quad y := 4dq^K e^{-J/2}.$$

By Lemma B.1,

$$(5.17) \quad ey \leq \frac{1}{4}.$$

This estimate controls the sum over additional vertices in the witness enumeration below.

For an instance Φ on G , a permissive block walk of length ℓ is a sequence of permissive blocks generated by the following process. Choose a root vertex $v_0 \notin X$ and set $B_0 = B_{\Phi_0}(v_0)$ where $\Phi_0 = \Phi$. At step $j \geq 0$ to generate B_{j+1} , let Φ_{j+1} be the instance generated by removing B_j from Φ_j . Choose an edge $(u, v) \in \partial_E B_j$, set the next root $v_{j+1} = v$ and let $B_{j+1} = B_{\Phi_{j+1}}(v_{j+1})$. Thus the blocks are disjoint and we record the blocks, roots and transition edges for a block walk. The weight of this block walk is defined by $q^{K \sum_{j=0}^{\ell} |B_j|}$ and the sum of total weights of all block walks of length ℓ is denoted by $\mathcal{W}_\ell(G)$.

Lemma 5.11. *For all sufficiently large positive integer n , with probability at least $1 - n^{-4}$ over $G \sim \mathcal{G}(n, d/n)$, for every $0 \leq \ell < n$,*

$$\mathcal{W}_\ell(G) \leq n^7 (4dq^K e^{2y})^\ell.$$

Proof. At first, we define the witness of a permissive block walk. For a permissive block walk $B_0, \dots, B_\ell$ of length ℓ with roots $v_0, \dots, v_\ell$, assume that the union of it contains t vertices in X and set $t_j = |B_j \cap X|$. The witness is defined as: for every $j = 0, \dots, \ell$,

- pick a spanning tree every connected component in $B_j \cap X$;
- for every spanning tree in $B_j \cap X$, pick an edge connecting v_j and it.

Note that the witness has exactly t edges.

For every $j < \ell$, there are at most $t_j + 1$ possible transitions between B_j and v_{j+1} . Hence there are at most $\prod_{j=0}^{\ell-1} (t_j + 1) \leq 2^t$ possible recorded transitions. Meanwhile, by the Cayley's formula, for a fixed collection containing $\ell + 1$ roots and additional t vertices in blocks, the number of witness is at most

$$(\ell + 1)(\ell + 1 + t)^{t-1}.$$

Then for every $t = 0, \dots, n$, the total number of witnesses with recorded transitions containing $\ell + 1$ roots and additional t vertices in blocks is at most

$$n(n-1) \dots (n-\ell) \binom{n-\ell-1}{t} (\ell+1)(\ell+1+t)^{t-1} 2^t \leq \frac{n^{\ell+1+t}}{t!} (\ell+1)(\ell+1+t)^{t-1} 2^t.$$

Fix a fixed witness with recorded transitions containing $\ell + 1$ roots and additional t vertices in blocks (denoted by $\mathbf{W}$). For a vertex $v \in \mathbf{W}$ such that it is prescribed to be in X , set $f_v = \deg_{\mathbf{W}}(v)$. Conditional on the event $\mathcal{E}_{\mathbf{W}}$ that $\mathbf{W}$ is in G , it holds that

$$\deg_G(v) = f_v + Y_v, \quad \text{where } Y_v \sim \text{Bin}(n-1-f_v, d/n).$$

By the Markov inequality,

$$\mathbb{P}(v \in X \mid \mathcal{E}_{\mathbf{W}}) \leq 2^{-(q-1-f_v)} \mathbf{E}[2^{Y_v}] \leq e^{-J+f_v \ln 2}.$$

Finner's inequality for a read-two family tells us that

$$\begin{aligned} \mathbb{P}(\text{All } t \text{ prescribed vertices are in } X \mid \mathcal{E}_{\mathbf{W}}) &\leq \prod_v \mathbb{P}(v \in X \mid \mathcal{E}_{\mathbf{W}})^{1/2} \\ &\leq \exp\left(-\frac{Jt}{2} + \frac{\ln 2}{2} \sum_v f_v\right) \\ &\leq e^{-Jt/2} 2^{\ell+t}. \end{aligned}$$

since $\sum_v f_v \leq 2(\ell + t)$. Therefore,

$$\begin{aligned} \mathbf{E}[\mathcal{W}_\ell(G)] &\leq \sum_{t \geq 0} \frac{n^{\ell+1+t}}{t!} (\ell+1)(\ell+1+t)^{t-1} 2^t \frac{d^{\ell+t}}{n^{\ell+t}} q^{K(\ell+1+t)} e^{-Jt/2} 2^{\ell+t} \\ &= nq^K (2dq^K)^\ell \sum_{t \geq 0} \frac{(\ell+1)(\ell+1+t)^{t-1}}{t!} \left(4dq^K e^{-J/2}\right)^t. \end{aligned}$$

To compute the series, we consider the exponential rooted-tree generating function T :

$$T(x) = \sum_{m \geq 1} m^{m-1} \frac{x^m}{m!}.$$

Then $T(x) = xe^{T(x)}$ and $T(x) \in (0, 1)$. By Lagrange inversion $[x^t]F(u) = \frac{1}{t}[u^{t-1}]F'(u)\phi(u)^t$ for $u = x\phi(u)$ with choice $\phi(u) = e^u$, $F(u) = e^{(\ell+1)u}$ and $u = T(x)$, it holds that

$$\sum_{t \geq 0} \frac{(\ell+1)(\ell+1+t)^{t-1}}{t!} x^t = e^{(\ell+1)T(x)}.$$

Hence,

$$\mathbf{E}[\mathcal{W}_\ell(G)] \leq nq^K (2dq^K)^\ell e^{(\ell+1)T(y)}.$$

Consider the fix-point iteration

$$T_0 = 0, \quad \forall k \geq 1, \quad T_k = ye^{T_{k-1}}.$$

Since $T_1 = y \geq T_0$ and $x \mapsto ye^x$ is increasing, by induction it holds that $T_0 \leq T_1 \leq T_2 \leq \dots$. For all $0 \leq x \leq 2y \leq (1/2e)$, it holds that $ye^x \leq ye^{2y} \leq 2y$. Therefore $\lim_{k \rightarrow \infty} T_k = T(y)$ and $0 \leq T_k \leq 2y$ for every $k \geq 0$. This means $T(y) \leq 2y$. Then by the Markov inequality at the threshold $n^7(4dq^K e^{2y})^\ell$, as long as $n \geq q^K e^{2y}$, it holds that

$$\mathbb{P}\left(\exists \ell < n, \mathcal{W}_\ell(G) \geq n^7(4dq^K e^{2y})^\ell\right) \leq \sum_{\ell=0}^{n-1} n^2 n^{-7} \cdot 2^{-\ell} \leq n^{-4}$$

concluding the lemma. $\square$

5.4. Approximation scheme to proper colorings. To design our algorithm, we show that with high probability, the rational polytope $\mathcal{P}_t(\Phi, B)$ can induce a good approximation to the real marginal vector $\mu_\Phi(B \leftarrow \cdot)$.

Lemma 5.12. *On the event of Lemmas 5.8 to 5.10, for every list-coloring instance Φ induced by pinning components in $G[X]$ from the original q -coloring instance, an unassigned connected component B in X and every $R = 0, \dots, n$, it holds that*

$$\text{diam}_H(\mathcal{P}_R(\Phi, B)) \leq \frac{C_{\text{leaf}} n^6 \ln n}{1 - \rho} \rho^R + \frac{2(q-1)n^6}{1 - \rho} \xi_P$$

where $C_{\text{leaf}} = 2 \ln 2 \cdot (q-2) \cdot (1 + C_{\text{vol}})$. Moreover, if $\text{diam}_H(\mathcal{P}_R(\Phi, B)) \leq \varepsilon$, then for every $z \in \mathcal{P}_R(\Phi, B)$ and $\sigma \in \Omega_B$, $e^{-\varepsilon} \mu_\Phi(B \leftarrow \sigma) \leq z_\sigma \leq e^\varepsilon \mu_\Phi(B \leftarrow \sigma)$.

Remark 5.13. We remark here that after X is pinned, the residual instance Φ satisfies that $|L_\Phi(v)| - \deg_{G_\Phi}(v) \geq q - \deg_G(v) \geq 2$. Therefore the same accuracy bound in Lemma 5.12 holds for every admissible feasible list-coloring instance induced by a feasible partial pinning in Φ .

Proof. We only need to consider the non-trivial case. By Proposition 5.7, at depth 0,

$$\text{diam}_H(\mathcal{P}_0(\Phi, B)) \leq 2m \ln 2.$$

Observe that B is either a maximal connected component in $G[X]$ or a union of at most $q-2$ connected components in $G[X]$ incident to a vertex $v \notin X$. Therefore, by Lemma 5.10,

$$m \leq (q-2) + (q-2)C_{\text{vol}} \ln n \leq (q-2)(1 + C_{\text{vol}}) \ln n.$$

Thus $\text{diam}_H(\mathcal{P}_0(\Phi, B)) \leq C_{\text{leaf}} \ln n$.

At depth $R > 0$, we first show that how to encode a recursive branch with a self-avoiding walk. For a branch

$$B_0 \rightarrow B_1 \rightarrow \dots \rightarrow B_R$$

with root $x_0, x_1, \dots, x_R$ where $x_0 \in B_0$ is the smallest vertex in B_0 . Assume that x_i is accessed by $u_{i-1} \in B_{i-1}$. Choose a minimal simple path from x_i to u_i in B_i . Then there is a self-avoiding

walk $P = x_0 = v_0, \dots, v_\ell = x_R$ of length at least R containing $x_0, x_1, \dots, x_R$ with other vertices of degree $> D$. Furthermore,

$$\prod_{i=1}^R \kappa_{x_i} \leq \prod_{i=1}^{\ell} w(v_i).$$

Conversely, it can be easily verified that different branches do not use the same self-avoiding walk. Unfold the recursion in Proposition 5.7. Thus by Lemma 5.8, it holds that

$$\begin{aligned} \text{diam}_H(\mathcal{P}_R(\Phi, B)) &\leq C_{\text{leaf}} \ln n \sum_{\ell=R}^{n-1} W_\ell(G) + 2(q-1)\xi_P \sum_{\ell=1}^{n-1} W_\ell(G) \\ &\leq C_{\text{leaf}} n^6 \ln n \sum_{\ell \geq R} \rho^\ell + 2(q-1)n^6 \cdot \xi_P \sum_{\ell \geq 1} \rho^\ell \\ &\leq \frac{C_{\text{leaf}} n^6 \ln n}{1-\rho} \rho^R + \frac{2(q-1)n^6}{1-\rho} \xi_P. \end{aligned}$$

For the approximation to $p_B(\cdot) = \mu_\Phi(B \leftarrow \cdot)$, note that $d_H(z, p_B) \leq \varepsilon$. Then it is not hard to see $-\varepsilon \leq \ln(z_\sigma/p_B(\sigma)) \leq \varepsilon$ and hence $e^{-\varepsilon} \mu_\Phi(B \leftarrow \sigma) \leq z_\sigma \leq e^\varepsilon \mu_\Phi(B \leftarrow \sigma)$ for every $\sigma \in \Omega$. $\square$

For a rational number $\theta \in (0, 1)$, set parameters

$$(5.18) \quad R_\theta = \left\lceil \frac{\ln(4C_{\text{leaf}} n^6 \ln n / ((1-\rho)\theta))}{\ln(1/\rho)} \right\rceil, \quad P_\theta = \max \left\{ 3, \left\lceil \log_2 \frac{32(q-1)n^6}{\theta(1-\rho)} \right\rceil \right\}.$$

Now we are ready to prove Theorem 1.6.

Proof of Theorem 1.6. Assume the joint event of Lemmas 2.4 and 5.8 to 5.11. The event occurs with probability at least $1 - o_n(1) - 5n^{-4} = 1 - o_n(1)$.

Set $R = R_{\varepsilon/(32n)}$ and $P = P_{\varepsilon/(32n)}$ like (5.18). It is clear to verify that $R = O(\ln(n/\varepsilon))$ and $P = O(\ln(n/\varepsilon))$ where the hidden constants depend only on η, d and q . If $R \geq n$, it must hold that $\varepsilon = e^{-\Omega(n)}$ and we enumerate all q^n assignments and check whether it is a proper coloring. This is a deterministic algorithm to exactly compute the number of proper colorings in $\text{poly}(n, \varepsilon^{-1})$ bit operations.

Otherwise, using the algorithm in Lemma 2.4, we obtain a proper q -coloring σ of G . The approximation scheme $\mathbf{A}_{\eta, d, q}$ runs as following:

- (1) initialize $\Phi = (G_\Phi = G, (L_\Phi(v) = [q])_{v \in V})$;
- (2) process the component of $G[X]$ in a preset order:
 - (a) for the current component B , construct the rational polytope $\mathcal{P} = \mathcal{P}_R(\Phi, B)$ and record $\hat{p}_{\sigma(B)} = \min_{z \in \mathcal{P}} z_{\sigma(B)}$;
 - (b) pin B to the partial coloring $\sigma(B)$ in Φ ;
- (3) process the remaining vertex of G in a preset order:
 - (a) for the current vertex v , construct the rational polytope $\mathcal{P} = \mathcal{P}_R(\Phi, v)$ and pick $\hat{p}_{\sigma(v)} = \min_{z \in \mathcal{P}} z_{\sigma(v)}$;
 - (b) pin v to color $\sigma(v)$ in Φ ;
- (4) output the rational product of reciprocals of recorded rational numbers.

We first show the accuracy of the algorithm. Suppose that the instances in the algorithm are $\Phi = \Phi_0, \dots, \Phi_M = \emptyset$, the true values of the marginal probabilities are $p_0, \dots, p_{M-1}$ and the estimations by the algorithm are $\hat{p}_0, \dots, \hat{p}_{M-1}$. Note that all needed analysis holds through the

running of $A_{\eta,d,q}$. By Lemma 5.12, it holds that $e^{-\varepsilon/(16n)} \leq \widehat{p}_j/p_j \leq 1$. Hence

$$1 \leq \frac{\prod_{j=0}^{M-1} \widehat{p}_j^{-1}}{|\Omega(G, q)|} = \frac{\prod_{j=0}^{M-1} \widehat{p}_j^{-1}}{\prod_{j=0}^{M-1} p_j^{-1}} \leq e^{\varepsilon/16} \leq 1 + \varepsilon.$$

Now we show the bit complexity. For a linear-fractional programming with rounding process in the recursive construction of rational polytopes, it contains $O(q^{|B|})$ variables and $O(q^{2|B|})$ linear constraints. Using a crude bound for solving the linear-fractional programming, it cost at most

$$q^{50|B|} (nP)^{40}$$

bit operations to finish the linear-fractional programming task in one call of a child polytope.

We turn to the total bit complexity to estimate p_i . Note that a recursive branch of the computation is encoded by a starting block B , a block permissive block walk of length at most R and pairs of partial colorings on blocks. Therefore, the cost to solve all linear-fractional programmings is at most

$$(5.19) \quad q^{40|B|} \cdot \left(\sum_{v \in B} \deg_G(v) \right) (nP)^{40} \cdot (nP)^{40} \sum_{\substack{\text{a permissive block walk} \\ B_0, \dots, B_\ell: \ell \leq R}} \prod_{j=0}^{\ell} q^{2|B_j|} \cdot q^{50|B_j|}.$$

By Lemmas 5.9 to 5.11, $|B|$ is of size at most $4 \ln n$ and there are at most $C_{\text{vol}} \ln n$ roots to start permissive block walks. Hence by the choice of R and P , the bit complexity to estimate a marginal is at most

$$O \left(n^{40 \ln q + K \cdot C_{\text{vol}}} (nP)^{80} \sum_{\ell=0}^R \mathcal{W}_\ell(G) \right) = \left(\frac{n + \text{bit}(\varepsilon)}{\varepsilon} \right)^C$$

where C depends only on η, d and q . Then we conclude the desired bit complexity of $A_{\eta,d,q}$ $\square$

ACKNOWLEDGEMENT

We thank Guoliang Qiu and Chihao Zhang for helpful discussion on counting proper colorings of random graphs.

REFERENCES

- [BBR26] Ferenc Bencs, Khallil Berrekkal, and Guus Regts. Deterministic approximate counting of colorings with fewer than 2Δ colors via absence of zeros. *TheoretCS*, Volume 5, Jan 2026. 6
- [BGK⁺07] Mohsen Bayati, David Gamarnik, Dimitriy Katz, Chandra Nair, and Prasad Tetali. Simple deterministic approximation algorithms for counting matchings. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*, STOC '07, pages 122–127, New York, NY, USA, 2007. Association for Computing Machinery. 1
- [Bir57] Garrett Birkhoff. Extensions of Jentzsch's theorem. *Transactions of the American Mathematical Society*, 85:219–227, 1957. 8
- [CC62] Abraham Charnes and William W. Cooper. Programming with linear fractional functionals. *Naval Research Logistics Quarterly*, 9(3-4):181–186, 1962. 8
- [CC19] Jin-Yi Cai and Xi Chen. A decidable dichotomy theorem on directed graph homomorphisms with non-negative weights. *computational complexity*, 28(3):345–408, 2019. 1
- [CCL13] Jin-Yi Cai, Xi Chen, and Pinyan Lu. Graph homomorphisms with complex values: A dichotomy theorem. *SIAM Journal on Computing*, 42(3):924–1029, 2013. 1

- [CFG⁺25] Xiaoyu Chen, Weiming Feng, Heng Guo, Xinyuan Zhang, and Zongrui Zou. Deterministic counting from coupling independence. In *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 667–683, 2025. 5, 6
- [DCS12] Hugo Duminil-Copin and Stanislav Smirnov. The connective constant of the honeycomb lattice equals $\sqrt{2} + \sqrt{2}$. *Annals of Mathematics*, 175(3):1653–1665, 2012. 4
- [EHŠV18] Charilaos Efthymiou, Thomas P. Hayes, Daniel Štefankovič, and Eric Vigoda. Sampling random colorings of sparse random graphs. In *Proceedings of the 2018 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, SODA '18, page 1759–1771, USA, 2018. Society for Industrial and Applied Mathematics. 6
- [FGW⁺25] Weiming Feng, Heng Guo, Chunyang Wang, Jiaheng Wang, and Yitong Yin. Toward derandomizing markov chain monte carlo. *SIAM Journal on Computing*, 54(3):775–813, 2025. 6
- [GK12] David Gamarnik and Dmitriy Katz. Correlation decay and deterministic FPTAS for counting colorings of a graph. *Journal of Discrete Algorithms*, 12:29–47, 2012. Selected papers from the 4th Algorithms and Complexity in Durham Workshop ACiD 2010. 1, 7, 9, 19
- [GKM15] David Gamarnik, Dmitriy Katz, and Sidhant Misra. Strong spatial mixing of list coloring of graphs. *Random Structures & Algorithms*, 46(4):599–613, 2015. 7, 9, 19
- [HPR20] Tyler Helmuth, Will Perkins, and Guus Regts. Algorithmic Pirogov–Sinai theory. *Probability Theory and Related Fields*, 176(3):851–895, 2020. 6
- [JKP20] Matthew Jenssen, Peter Keevash, and Will Perkins. Algorithms for #BIS-Hard Problems on Expander Graphs. *SIAM Journal on Computing*, 49(4):681–710, January 2020. 6
- [JVV86] Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986. 8
- [LSS19] Jingcheng Liu, Alistair Sinclair, and Piyush Srivastava. The Ising partition function: Zeros and deterministic approximation. *Journal of Statistical Physics*, 174(2):287–315, 2019. 6
- [LSS25] Jingcheng Liu, Alistair Sinclair, and Piyush Srivastava. Correlation decay and partition function zeros: Algorithms and phase transitions. *SIAM Journal on Computing*, 54(4):FOCS19–200–FOCS19–252, 2025. 6
- [LWZ14] Pinyan Lu, Menghui Wang, and Chihao Zhang. Fptas for weighted fibonacci gates and its applications. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming*, pages 787–799, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg. 1
- [LY13] Pinyan Lu and Yitong Yin. Improved FPTAS for multi-spin systems. In Prasad Raghavendra, Sofya Raskhodnikova, Klaus Jansen, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 16th International Workshop, APPROX 2013, and 17th International Workshop, RANDOM 2013, Berkeley, CA, USA, August 21-23, 2013. Proceedings*, pages 639–654, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. 1, 10
- [PSW96] Boris Pittel, Joel Spencer, and Nicholas Wormald. Sudden emergence of a giantk-core in a random graph. *Journal of Combinatorial Theory, Series B*, 67(1):111–151, 1996. 9
- [SSŠY17] Alistair Sinclair, Piyush Srivastava, Daniel Štefankovič, and Yitong Yin. Spatial mixing and the connective constant: optimal bounds. *Probability Theory and Related Fields*, 168(1):153–197, 2017. 2
- [SSY13] Alistair Sinclair, Piyush Srivastava, and Yitong Yin. Spatial mixing and approximation algorithms for graphs with bounded connective constant. In *2013 IEEE 54th Annual*

- Symposium on Foundations of Computer Science*, pages 300–309, 2013. [2](#)
- [Wei06] Dror Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '06, pages 140–149, New York, NY, USA, 2006. Association for Computing Machinery. [1](#)
- [Yin14] Yitong Yin. Spatial mixing of coloring random graphs. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming*, pages 1075–1086, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg. [2](#), [5](#), [7](#), [16](#), [17](#), [19](#)
- [YZ13] Yitong Yin and Chihao Zhang. Approximate counting via correlation decay on planar graphs. In *Proceedings of the 2013 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, SODA '13, pages 47–66, USA, 2013. Society for Industrial and Applied Mathematics. [1](#)
- [YZ16] Yitong Yin and Chihao Zhang. Sampling in Potts model on sparse random graphs. In Klaus Jansen, Claire Mathieu, José D. P. Rolim, and Chris Umans, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2016)*, volume 60 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 47:1–47:22, Dagstuhl, Germany, 2016. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. [2](#), [5](#), [6](#), [7](#), [16](#), [17](#)

APPENDIX A. PROOF OF CONTRACTIONS

A.1. Birkhoff contraction of Hilbert distance. Now we show Lemma [2.2](#). By Lemma [2.1](#), for two distinct probability vectors p, q ,

$$\|p - q\|_{\text{TV}} \leq \tanh(\delta/4)$$

where $\delta = d_H(p, q)$. Define $r_k = p(k) - q(k)$. For any real vector h , set $m = \min_k h_k$ and $M = \max_k h_k$. It holds that $\sum_{r_k > 0} r_k = -\sum_{r_k < 0} r_k = \|p - q\|_{\text{TV}}$. Therefore,

$$(A.1) \quad \sum_k r_k h_k \leq M \sum_{r_k > 0} r_k + m \sum_{r_k < 0} r_k = (M - m) \|p - q\|_{\text{TV}}.$$

Since [\(A.1\)](#) also holds for $-h$, it holds that

$$(A.2) \quad \left| \sum_k (p(k) - q(k)) h_k \right| \leq \tanh(\delta/4) \left(\max_k h_k - \min_k h_k \right).$$

Set $h_k = \ln(x(k)/y(k))$ and $z(t)_k = y(k)e^{th_k}$. For every a , define

$$F_a(t) = \ln(Az(t))_a, \quad p^{a,t}(k) = \frac{A(a, k)z(t)_k}{(Az(t))_a}.$$

It holds that $p^{a,t}$ is a positive probability vector and $F'_a(t) = \sum_k p^{a,t}(k) h_k$. Note that for a, b ,

$$\frac{\max_i p^{a,t}(i)/p^{b,t}(i)}{\min_j p^{a,t}(j)/p^{b,t}(j)} = \max_{i,j} \frac{A(a, i)A(b, j)}{A(a, j)A(b, i)} \leq e^{\Delta_B(A)}.$$

Then applying [\(A.2\)](#) to $p^{a,t}$ and $p^{b,t}$ and integrating from 0 to 1, it holds that

$$\begin{aligned} |(F_a(1) - F_a(0)) - (F_b(1) - F_b(0))| &\leq \int_0^1 |F'_a(t) - F'_b(t)| dt \\ &\leq \tanh(\Delta_B(A)/4) \left(\max_k h_k - \min_k h_k \right) \\ &= c_B(A) d_H(x, y). \end{aligned}$$

Observe that $F_a(1) - F_a(0) = \ln \frac{(Ax)_a}{(Ay)_a}$. Thus we conclude the lemma.

A.2. Contraction of Hilbert distance for list colorings. In this part, we prove Lemma 5.6. When $p = p'$, the upper bound is trivial.

Assume that $p \neq p'$. Let $p(t) = (1-t)p' + tp$ for $t \in [0, 1]$. It holds that $p(t)$ is a positive probability vector on Ω and for every $c \in \Omega$, $p_c(t) \leq 1/s$. Let $d_c = p_c - p'_c$ and $x_c(t) = d_c/p_c(t)$. Then it holds that $\sum_c p_c(t)x_c(t) = 0$ and therefore $\min_c x_c(t) \leq 0 \leq \max_c x_c(t)$. Let

$$D_{a,b}(t) := \frac{d}{dt} \left(\ln \frac{1 - p_a(t)}{1 - p_b(t)} \right) = -\frac{p_a(t)}{1 - p_a(t)} x_a(t) + \frac{p_b(t)}{1 - p_b(t)} x_b(t).$$

Put $m(t) = \min_c x_c(t)$ and $M(t) = \max_c x_c(t)$. Moreover, set

$$\alpha = \frac{p_a(t)}{1 - p_a(t)}, \quad \beta = \frac{p_b(t)}{1 - p_b(t)}.$$

It is clear that $\alpha, \beta \in [0, 1/(s-1)]$ and therefore $\alpha x_a(t), \beta x_b(t) \in [m(t)/(s-1), M(t)/(s-1)]$ since $m(t) \leq 0 \leq M(t)$. Then for every $a, b \in \Omega$ and $t \in [0, 1]$,

$$|D_{a,b}(t)| \leq \frac{M(t) - m(t)}{s-1}.$$

Write $R_c := p_c/p'_c$ for every $c \in \Omega$. It holds that

$$x_c(t) = \frac{d_c}{p_c(t)} = \frac{p_c - p'_c}{(1-t)p'_c + tp_c} = \frac{R_c - 1}{tR_c + (1-t)}.$$

With the observation that $x_c(t)$ is increasing with R_c , we have

$$\begin{aligned} \left| \ln \frac{1 - p_a}{1 - p'_a} - \ln \frac{1 - p_b}{1 - p'_b} \right| &\leq \frac{1}{s-1} \int_0^1 (M(t) - m(t)) dt \\ &= \frac{1}{s-1} \int_0^1 \left(\frac{\max_c R_c - 1}{t \max_c R_c + (1-t)} - \frac{\min_c R_c - 1}{t \min_c R_c + (1-t)} \right) dt \\ &= \frac{1}{s-1} \left(\ln \max_c R_c - \ln \min_c R_c \right) \\ &= \frac{d_H(p, p')}{s-1}. \end{aligned}$$

where the last equality comes directly from the definition of $d_H(\cdot, \cdot)$.

APPENDIX B. PARAMETER CHOICES FOR RANDOM GRAPH COLORINGS

We verify that all parameter estimates used in Section 5 hold for a common threshold $d_0(\eta)$, uniformly over $q \geq (2 + \eta)d$.

Lemma B.1 (Parameter choice). *For every fixed $\eta \in (0, 1)$, there exists $d_0 = d_0(\eta) \geq 8$ such that, for every $d \geq d_0$ and every integer $q \geq (2 + \eta)d$, the parameters defined in (5.10), (5.12), and (5.16) satisfy (5.13), (5.14), (5.15), and (5.17).*

Proof. Recall that $K = 200$ is an absolute constant. Set $j_\eta = (2 + \eta) \ln 2 - 1 > 0$. Choose $d_0 \geq 8$ so that, for every $d \geq d_0$,

$$\begin{aligned}
& \eta d \geq 24, \\
& d e^{-dh(\eta/6)/2} \leq \frac{\eta}{2(3 + \eta)}, \\
& 4e d e^{-dh(\eta/6)/4} \leq 1, \\
& d \geq \frac{2K}{(2 + \eta) \ln 2}, \\
& j_\eta d - \ln 2 \geq 2 \ln(16e) + 2(K + 1) \ln d + 2K \ln(2 + \eta).
\end{aligned}
\tag{B.1}$$

Such a choice exists because $h(\eta/6) > 0$ and $j_\eta > 0$: the exponential terms dominate the polynomial factors, and the linear term in the last inequality dominates $\ln d$. All these conditions depend only on η and the fixed constant K .

First, $\delta \geq \eta/6$ and $\delta d \geq 4$. Since $(d + q)/3 = (1 + 2\delta)d$, we have

$$D - 3 > (1 + 2\delta)d - 4 \geq (1 + \delta)d.$$

Moreover,

$$q - D - 1 \geq \frac{2q - d}{3} - 1 \geq \frac{(3 + 2\eta)d}{3} - 1 \geq \frac{(3 + \eta)d}{3}.$$

In particular, $D < q - 1$ and

$$d\gamma \leq \frac{3}{3 + \eta}.$$

This proves (5.13).

Next, $h'(x) = \ln(1 + x) > 0$ for $x > 0$, so $r \leq e^{-dh(\eta/6)}$. Consequently,

$$d(\gamma + \sqrt{r}) \leq \frac{3}{3 + \eta} + \frac{\eta}{2(3 + \eta)} = \lambda_0 < \rho < 1,$$

and

$$e d r^{1/4} \leq e d e^{-dh(\eta/6)/4} \leq \frac{1}{4}.$$

These are (5.14) and (5.15).

Finally, put $q_0 = (2 + \eta)d$. By (B.1), $q_0 \geq 2K/\ln 2$. The function $u \mapsto u^K e^{-u \ln 2/2}$ is nonincreasing for $u \geq q_0$, since its logarithmic derivative is $K/u - \ln 2/2 \leq 0$. It follows that

$$\begin{aligned}
e y &= 4e d e^{(d + \ln 2)/2} q^K e^{-q \ln 2/2} \\
&\leq 4e d ((2 + \eta)d)^K e^{-(j_\eta d - \ln 2)/2} \\
&\leq \frac{1}{4},
\end{aligned}$$

where the last step uses the final inequality in (B.1). This proves (5.17) with a threshold d_0 independent of q . $\square$