Fast Sampling of b-Matchings and b-Edge Covers

Zongchen Chen* Yuzhou Gu†

Abstract

For an integer $b \geq 1$, a b-matching (resp. b-edge cover) of a graph $G = (V, E)$ is a subset $S \subseteq E$ of edges such that every vertex is incident with at most (resp. at least) b edges from S. We prove that for *any* $b \geq 1$ the simple Glauber dynamics for sampling (weighted) b-matchings and b-edge covers mixes in $O(n \log n)$ time on all n-vertex bounded-degree graphs. This significantly improves upon previous results which have worse running time and only work for b-matchings with $b \leq 7$ and for b-edge covers with $b \leq 2$.

More generally, we prove *spectral independence* for a broad class of binary symmetric Holant problems with *log-concave* signatures, including b-matchings, b-edge covers, and antiferromagnetic 2-spin edge models. We hence deduce optimal mixing time of the Glauber dynamics from spectral independence.

The core of our proof is a recursive coupling inspired by [CZ23] which upper bounds the Wasserstein W_1 distance between distributions under different pinnings. Using a similar method, we also obtain the optimal $O(n \log n)$ mixing time of the Glauber dynamics for the hardcore model on n-vertex bounded-degree claw-free graphs, for *any* fugacity λ. This improves over previous works which have at least cubic dependence on n.

1 Introduction

1.1 b-Matchings and b-edge covers Let $G = (V, E)$ be a graph and $b \geq 1$ be an integer. Let $E_v = \{e \in E : e \text{ incident to } v\}$ be the set of all adjacent edges of a vertex $v \in V$. A b-matching of G is a subset $S \subseteq E$ of edges such that $|S \cap E_v| \leq b$ for all $v \in V$. When $b = 1$ this reduces to a usual matching of G. We consider the problem of sampling random weighted b-matchings of a given graph G. Write $\mathcal{M}_b = \mathcal{M}_b(G)$ for the collection of all b-matchings of G. For $\lambda > 0$, consider the Gibbs distribution $\mu = \mu_{G,b,\lambda}$ on $\mathcal{M}_b$ given by

$$\mu(S) := \frac{\lambda^{|S|}}{Z}, \quad \forall S \in \mathcal{M}_b$$

where $Z = Z_{G,b}(\lambda)$ is a normalization constant, known as the partition function, defined as

$$Z := \sum_{S \in \mathcal{M}_b} \lambda^{|S|}.$$

Note that if $\lambda = 1$ then μ is the uniform distribution over $\mathcal{M}_b$ and Z counts the total number of b-matchings in $\mathcal{M}_b$.

For $b = 1$, namely the usual matchings, such a model is called the *monomer-dimer model*. Approximately counting and sampling matchings is a fundamental problem in theoretical computer science and also one of the first successful applications of Markov chain Monte Carlo (MCMC) methods in approximate sampling and counting combinatorial objects. In a classical work [JS89], Jerrum and Sinclair proved rapid mixing of Glauber dynamics for sampling from the monomer-dimer model. The best mixing time result to date is $O(n^2 m \log n)$ on arbitrary graphs where m is the number of edges [Jer03], and only very recently this was improved to $O(n \log n)$ on all bounded-degree graphs [CLV22].

For general $b \geq 1$, [HLZ16] presented a polynomial-time algorithm for approximately sampling b-matchings on all graphs when $b \leq 7$. Their algorithm is based on MCMC and they utilize the notion of *windable functions* introduced in [McQ13] to construct canonical paths and bound the spectral gap of the Markov chain. However, as pointed out in [HLZ16], for 8-matchings the associated constraint function is no longer windable under their characterization and hence their approach cannot work for $b \geq 8$.

*zchen83@buffalo.edu. Department of Computer Science and Engineering, University at Buffalo.

†yuzhougu@ias.edu. School of Mathematics, Institute for Advanced Study. Supported by National Science Foundation under Grant No. DMS-1926686.

Another closely related problem is sampling b-edge covers of a given graph. A subset $S \subseteq E$ of edges is called a b-edge cover if every vertex is incident with at least b edges, i.e., $|S \cap E_v| \geq b$ for every $v \in V$. For $b = 1$, i.e., the usual edge covers, the counting and sampling problems have been extensively studied as well [BR09, LLL14, LLZ14, HLZ16, GLLZ21, BCR21, CLV22]. In particular, [LLL14] presented a deterministic algorithm for counting unweighted edge covers for all graphs using the correlation decay approach with a running time $O(m^{1+\log_2 6} n^2)$, and this was later generalized to weighted edge covers in [LLZ14]. Deterministic algorithms based on the polynomial interpolation approach were also given for all bounded-degree graphs in [GLLZ21, BCR21]. More recently, it was shown in [CLV22] that the Glauber dynamics for sampling edge covers mixes in $O(n \log n)$ time on all bounded-degree graphs.

Meanwhile, the problem of sampling and counting b-edge covers for larger b is far from clear. The MCMC-based algorithm in [HLZ16] can be applied to count b-edge covers for $b \leq 2$, which only slightly extends the classical case of $b = 1$. Similar to b-matchings, the approach from [HLZ16] no longer works for $b \geq 3$ due to the failure of windability.

In this paper we attempt to answer the following question: Are there polynomial-time algorithms for approximately sampling/counting b-matchings and b-edge covers of a given graph for *any* $b \geq 1$? We give a positive answer to this question for all bounded-degree graphs. More specifically, we show that the Glauber dynamics, a simple Markov chain for sampling b-matchings/b-edge covers, converges in $O(n \log n)$ time which is optimal.

One can simultaneously generalize both b-matchings and b-edge covers by assigning a different threshold to each vertex. More specifically, let $\boldsymbol{b} = (b_v)_{v \in V} \in \mathbb{N}^V$ be a vector of thresholds on all vertices. We consider the collection $\mathcal{M}_{\boldsymbol{b}} = \mathcal{M}_{\boldsymbol{b}}(G)$ of generalized $\boldsymbol{b}$-matchings, defined as

$$\mathcal{M}_{\boldsymbol{b}} = \{S \subseteq E : \forall v \in V, |S \cap E_v| \leq b_v\}.$$

For $\lambda > 0$ the Gibbs distribution $\mu = \mu_{G,\boldsymbol{b},\lambda}$ is given by

$$\mu(S) := \frac{\lambda^{|S|}}{Z}, \quad \forall S \in \mathcal{M}_{\boldsymbol{b}}$$

and the partition function $Z = Z_{G,\boldsymbol{b}}(\lambda)$ is defined as

$$Z := \sum_{S \in \mathcal{M}_{\boldsymbol{b}}} \lambda^{|S|}.$$

Thus, for uniform $\boldsymbol{b} = b\mathbf{1}$ where $\mathbf{1}$ is the all-ones vector we obtain b-matchings, and for $b_v = d_v - b$ where d_v is the degree of v we get the complements of b-edge covers.

Our main contribution is to establish rapid mixing of the Glauber dynamics for sampling general $\boldsymbol{b}$-matchings for any $\boldsymbol{b} \in \mathbb{N}^V$ on all bounded-degree graphs. In each step of the Glauber dynamics, one picks an edge $e \in E$ uniformly at random and updates its status, $e \in S$ (occupied) or $e \notin S$ (unoccupied), conditional on the configuration of all other edges; in particular, if including e violates the subset S being a $\boldsymbol{b}$-matching then e must be unoccupied in this update. It is easy to show that the Glauber dynamics is ergodic for sampling $\boldsymbol{b}$-matchings.

THEOREM 1.1. ($\boldsymbol{b}$-MATCHINGS) *Let $\Delta \geq 3$ be an integer and $G = (V, E)$ be an n-vertex graph of maximum degree Δ. Let $\boldsymbol{b} \in \mathbb{N}^V$ be a vector of vertex thresholds. Then for any $\lambda > 0$, the Gibbs distribution $\mu = \mu_{G,\boldsymbol{b},\lambda}$ over $\boldsymbol{b}$-matchings is $O_{\Delta,\lambda}(1)$-spectrally independent. Furthermore, the Glauber dynamics for sampling from μ mixes in $O_{\Delta,\lambda}(n \log n)$ time.*

We prove Theorem 1.1 by the spectral independence method which was introduced recently in [ALO20] and has become a powerful tool for proving optimal mixing time of Glauber dynamics. Our proof of spectral independence is inspired by [CZ23] and uses a recursive coupling to bound the Wasserstein W_1 distance under two distinct pinnings. For uniformly random b-matchings with small b, our bound on spectral independence is $O(\Delta^b)$; see Remark 3.1 for more discussions. We remark that one interesting open problem is to show spectral independence with a constant independent of Δ even just for the usual matchings (monomer-dimer model), since then one would obtain $O(n \log n)$ mixing of Glauber dynamics on all graphs even with unbounded degrees, using new powerful techniques such as the field dynamics [CFYZ21, AJK+22, CFYZ22, CE22].

1.2 Holant problem with log-concave signatures Both b-matchings and b-edge covers belong to a much more general family of models called *Holant problems*, which can be understood as graphical models defined over subsets of edges of a given graph. Examples and applications of Holant problems include also perfect matchings [JSV04], even subgraphs [JS93, GJ18, LSS19, CLV22, CZ23, FGW22], Fibonacci gates [LWZ14], spin systems on line graphs [DHJM21, GLLZ21, BCR21, CLV22], etc.

We consider the following *binary symmetric Holant problem.* Let $G = (V, E)$ be a graph with n vertices. For each vertex v let d_v denote the degree of v. We consider a family of constraint functions on all vertices denoted by $\boldsymbol{f} = (f_v)_{v\in V}$, where each vertex v is associated with a constraint function $f_v : \mathbb{N} \to \mathbb{R}_{\geq 0}$. Also, let $\boldsymbol{\lambda} = (\lambda_e)_{e\in E} \in \mathbb{R}^E_{>0}$ be a vector of edge weights. The Gibbs distribution $\mu = \mu_{G,\boldsymbol{f},\boldsymbol{\lambda}}$ and the partition function $Z = Z_{G,\boldsymbol{f},\boldsymbol{\lambda}}$ of the Holant problem is defined as

$$\mu(S) := \frac{1}{Z}\prod_{v\in V} f_v(|S\cap E_v|)\prod_{e\in S}\lambda_e, \quad \forall S\subseteq E;$$
$$Z := \sum_{S\subseteq E}\prod_{v\in V} f_v(|S\cap E_v|)\prod_{e\in S}\lambda_e.$$

When $f_v(k) = \mathbb{1}\{k \leq b_v\}$ for some $\boldsymbol{b} = (b_v)_{v\in V}$, the Holant problem becomes $\boldsymbol{b}$-matchings.

Holant problems can be defined more generally by allowing each $f_v : 2^{E_v} \to \mathbb{R}_{\geq 0}$ to be a set function over subsets of neighboring edges of v. In this paper we consider only the *symmetric* case, i.e., the value of f_v depends only on $|S \cap E_v|$, the number of adjacent edges in S. Such symmetric constraint function f_v can be equivalently identified by the sequence $f_v = [f_v(0), f_v(1), \ldots, f_v(d_v)]$, which is called the *signature* at v.

Our main result for Holant problems establishes spectral independence and rapid mixing of Glauber dynamics when all the signatures are *log-concave* sequences.

Definition 1.1. (Log-concave signature) *A sequence $f = [f(0), f(1), \ldots, f(d)]$ of non-negative real numbers is called a log-concave signature if it satisfies the following conditions:*

(a) Log-concavity: $f(k)^2 \geq f(k-1)f(k+1)$ for all $1 \leq k \leq d-1$;

(b) No internal zeros: if $f(k_1) > 0$ and $f(k_2) > 0$ for some $0 \leq k_1 < k_2 \leq d$, then $f(k) > 0$ for all $k_1 \leq k \leq k_2$ (i.e., the support of f is consecutive).

For example, the signature $f = [1, \ldots, 1, 0, \ldots, 0]$ for the function $f(k) = \mathbb{1}\{k \leq b\}$ is log-concave.

Theorem 1.2. (Holant problem, informal) *Let $G = (V, E)$ be an n-vertex graph of maximum degree Δ. Suppose that $\boldsymbol{f} = (f_v)_{v\in V}$ is a collection of log-concave signatures with $f_v(0) > 0$ for all $v \in V$. Let $\boldsymbol{\lambda} \in \mathbb{R}^E_{>0}$ be a vector of edge weights. Then the Gibbs distribution $\mu = \mu_{G,\boldsymbol{f},\boldsymbol{\lambda}}$ for the Holant problem $(G, \boldsymbol{f}, \boldsymbol{\lambda})$ is $O_{\Delta,\boldsymbol{f},\boldsymbol{\lambda}}(1)$-spectrally independent. Furthermore, the Glauber dynamics for sampling from μ has modified log-Sobolev constant at least $1/(Cn)$ and mixing time at most $Cn\log n$, where $C = C(\Delta, \boldsymbol{f}, \boldsymbol{\lambda})$ does not depend on n.*

This is informal because technically speaking, vectors $\boldsymbol{f}$ and $\boldsymbol{\lambda}$ are dependent on n in dimensions. For a precise statement, see Theorem 3.1.

We remark that the Gibbs distribution $\mu = \mu_{G,\boldsymbol{f},\boldsymbol{\lambda}}$ in Theorem 1.2 is supported on $\boldsymbol{b}$-matchings where $b_v = \max\{0 \leq k \leq d_v : f_v(k) > 0\}$, and thus the Glauber dynamics for sampling from μ is ergodic. Our assumptions of log-concave signatures in fact generalize previous works [GLLZ21, BCR21, CLV22] which essentially require that the generating polynomial $P(x) = \sum_{k=0}^{d}\binom{d}{k}f(k)x^k$ associated with every signature f is real-rooted, which implies the log-concavity of f by Newton inequalities, see e.g. [Brä15]. Hence, Theorem 1.2 applies to many classes of Holant problems including the antiferromagnetic 2-spin systems on line graphs.

1.3 Hardcore model on claw-free graphs Another contribution of ours is that the Glauber dynamics has the optimal $O(n\log n)$ mixing time for the hardcore model on n-vertex bounded-degree claw-free graphs. In the hardcore model, we are given a graph $G = (V, E)$ and $\boldsymbol{\lambda} = (\lambda_v)_{v\in V} \in \mathbb{R}^V_{>0}$ a vector of vertex weights called fugacity. A set $I \subseteq V$ is called an independent set if $e \not\subseteq I$ for all $e \in E$. Let $\mathcal{I} \subseteq 2^V$ be the set of all

independent sets of G. Define the Gibbs distribution $\mu = \mu_{G,\boldsymbol{\lambda}}$ and the partition function $Z = Z_{G,\boldsymbol{\lambda}}$, also called the (multivariate) independence polynomial, as

$$\mu(I) := \frac{1}{Z}\prod_{v\in I}\lambda_v, \qquad \forall I \in \mathcal{I};$$
$$Z := \sum_{I\in\mathcal{I}}\prod_{v\in I}\lambda_v.$$

Specially, when $\boldsymbol{\lambda} = \lambda\mathbf{1}$, we denote the model as $\mu_{G,\lambda}$.

The Glauber dynamics is a natural Markov chain for sampling from the hardcore model. In each step of the Glauber dynamics, a vertex $v \in V$ is picked uniformly at random, and its status, $v \in I$ (occupied) or $v \notin I$ (unoccupied), is updated according to the configuration on all other vertices. Specifically, if v has at least one neighbor in the current independent set I, then nothing changes; if v has no neighbors in I, then it becomes occupied with probability $\frac{\lambda_v}{1+\lambda_v}$ and unoccupied with probability $\frac{1}{1+\lambda_v}$.

We consider sampling from the hardcore model on a special class of graphs, the *claw-free* graphs. A graph $G = (V, E)$ is claw-free if it does not include an induced $K_{1,3}$. In other words, there do not exist four distinct vertices $a, b, c, d \in V$ such that $(a, b), (a, c), (a, d) \in E$ but $(b, c), (b, d), (c, d) \notin E$. The class of claw-free graphs includes all line graphs by definition, and thus the hardcore model on claw-free graphs includes in particular the monomer-dimer model for matchings as a special case.

It was known that one can sample from the hardcore model on claw-free graphs in polynomial time. Generalizing the approach from [JS89, Jer03] for matchings, Matthews [Mat08] gave a Markov chain which mixes in $O(\Delta n^3)$ time where Δ is the maximum degree. Recently, [DGM21] proved that the Glauber dynamics mixes in $O(n^5 \log n)$ time for claw-free graphs, and more generally in polynomial time for graphs with bounded bipartite pathwidth.

In another direction, Patel and Regts [PR17] gave a polynomial-time deterministic algorithm (FPTAS) for approximating the partition function based on Barvinok's polynomial interpolation method [Bar16] and real-rootedness of the independence polynomial on claw-free graphs [CS07, LR19]. As is common for deterministic approximate counting algorithms, the exponent in n in the running time depends on parameters of the model.

Our main result for the hardcore model on claw-free graphs is that the Glauber dynamics has optimal mixing when the maximum degree is bounded. Again we prove optimal mixing by establishing spectral independence via a recursive coupling procedure.

THEOREM 1.3. (HARDCORE MODEL ON CLAW-FREE GRAPHS) *Let $\Delta \geq 3$ be an integer and $G = (V, E)$ be an n-vertex claw-free graph of maximum degree Δ. Let $\boldsymbol{\lambda} \in \mathbb{R}^V_{>0}$ be a vector of fugacity with $\lambda_{\min} := \min_{v\in V}\lambda_v$ and $\lambda_{\max} := \max_{v\in V}\lambda_v$. Then the Gibbs distribution $\mu = \mu_{G,\boldsymbol{\lambda}}$ of the hardcore model is $2(1 + \Delta\lambda_{\max})$-spectrally independent. Furthermore, the Glauber dynamics for sampling from μ has modified log-Sobolev constant at least $1/(Cn)$ and mixing time at most $Cn \log n$, where $C = C(\Delta, \lambda_{\max}, \lambda_{\min})$ does not depend on n.*

2 Preliminaries

In this section we give definitions and lemmas that are needed. We introduce with the Holant problems in mind, but the definitions and results work for the hardcore model with minor changes (e.g., replacing E with V).

Let 2^E be the collection of all subsets of E. We consider a distribution μ on 2^E. We view a subset $S \subseteq E$ equivalently as a binary indicator vector $\sigma = \mathbf{1}_S \in \{0, 1\}^E$, where $\sigma_e = 1$ for $e \in S$, and $\sigma_e = 0$ for $e \notin S$.

DEFINITION 2.1. (PINNING) *A pinning is a partial configuration $\tau \in \{0, 1\}^\Lambda$ for some $\Lambda \subseteq E$ such that $\mu_\Lambda(\tau) > 0$, where μ_Λ is the marginal distribution on Λ. Let μ^τ denote the conditional distribution on $E\backslash\Lambda$.*

DEFINITION 2.2. (MARGINAL BOUNDEDNESS) *We say μ is b-marginally bounded if for all pinnings τ on $\Lambda \subseteq E$ and all $e \in E\backslash\Lambda$, we have either $b \leq \mu^\tau(\sigma_e = 1) \leq 1 - b$ or $\mu^\tau(\sigma_e = 1) \in \{0, 1\}$.*

DEFINITION 2.3. (INFLUENCE MATRIX) *Let μ be a distribution on 2^E and τ be a pinning on $\Lambda \subseteq E$. The pairwise influence matrix $J^\tau_\mu \in \mathbb{R}^{(E\backslash\Lambda)\times(E\backslash\Lambda)}$ is defined as following: for all $e, f \in E\backslash\Lambda$, let*

$$J^\tau_\mu(e, f) = \mu^\tau(\sigma_f = 1|\sigma_e = 1) - \mu^\tau(\sigma_f = 1|\sigma_e = 0)$$

when $e \neq f$ and $\min\{\mu^\tau(\sigma_e = 1), \mu^\tau(\sigma_e = 0)\} > 0$, and let $J^\tau_\mu(e, f) = 0$ otherwise.

Note that all eigenvalues of the influence matrix J_μ^τ are real (see e.g., [ALO20]).

DEFINITION 2.4. (SPECTRAL INDEPENDENCE [ALO20]) *We say μ is η-spectrally independent if for all pinnings τ we have $\lambda_{\max}(J_\mu^\tau) \le \eta$.*

THEOREM 2.1. ([CLV21, BCC+22, CLV22]) *Let μ be the Gibbs distribution of a Holant problem on an n-vertex graph of maximum degree Δ. If μ is η-spectrally independent and b-marginally bounded, then the Glauber dynamics has modified log-Sobolev constant at least $1/(Cn)$ and mixing time at most $Cn\log n$, where $C = C(\Delta, \eta, b)$ is a constant independent of n.*

For two distributions ν, π on 2^E, the 1-Wasserstein distance between them is defined as

$$W_1(\nu, \pi) = \inf_{\mathcal{C}} \mathbb{E}_{(\sigma,\tau)\sim\mathcal{C}} \left[d_{\mathrm{H}}(\sigma, \tau)\right],$$

where the infimum is over all couplings between ν and π, and $d_{\mathrm{H}}(\cdot,\cdot)$ denotes the Hamming distance between two elements from 2^E. We use the following lemma from [CZ23] to establish spectral independence; see also [CMM23, GGGHP22, CLMM23] which use similar approach.

LEMMA 2.1. ([CZ23]) *Let μ be the Gibbs distribution of a Holant problem. Suppose that for some constant $\eta > 0$, the following is true: For any two pinnings τ, τ' on the same subset $\Lambda \subseteq E$ which differ on exactly one edge, we have*

$$W_1(\mu^\tau, \mu^{\tau'}) \le \eta.$$

Then μ is η-spectrally independent.

3 Fast sampling for Holant problems with log-concave signatures

We first give a precise statement for Theorem 1.2. It is helpful to define the following local generating polynomial associated with each vertex, as introduced in [GLLZ21].

DEFINITION 3.1. (NORMALIZED GENERATING POLYNOMIAL) *For a signature $f = [f(0), f(1), \ldots, f(d)]$ with $f(0) > 0$, define the* normalized generating polynomial *to be*

$$P_f(x) = \frac{1}{f(0)} \sum_{k=0}^{d} \binom{d}{k} f(k) x^k.$$

DEFINITION 3.2. *Let $(G, \boldsymbol{f}, \boldsymbol{\lambda})$ be a Holant problem with log-concave signatures and $f_v(0) > 0$ for all $v \in V$. We define:*

$$r_{\max} := \max_{v\in V} \frac{f_v(1)}{f_v(0)}, \qquad r_{\min} := \min_{v\in V} \min_{k: f_v(k)>0} \frac{f_v(k)}{f_v(k-1)},$$

$$\lambda_{\max} := \max_{e\in E} \lambda_e, \qquad \lambda_{\min} := \min_{e\in E} \lambda_e,$$

$$P_{\max} := \max_{v\in V} P_{f_v}\left(r_{\max}\lambda_{\max}\right), \qquad \Delta := \max_{v\in V} d_v.$$

THEOREM 3.1. (HOLANT PROBLEM) *Let $G = (V, E)$ be an n-vertex graph of maximum degree Δ. Suppose that $\boldsymbol{f} = (f_v)_{v\in V}$ is a collection of log-concave signatures with $f_v(0) > 0$ for all $v \in V$. Let $\boldsymbol{\lambda} \in \mathbb{R}_{>0}^E$ be a vector of edge weights. Then the Gibbs distribution $\mu = \mu_{G,\boldsymbol{f},\boldsymbol{\lambda}}$ for the Holant problem $(G, \boldsymbol{f}, \boldsymbol{\lambda})$ is $O(P_{\max})$-spectrally independent, where $P_{\max}$ is defined in Definition 3.2. Furthermore, the Glauber dynamics for sampling from μ has modified log-Sobolev constant at least $1/(Cn)$ and mixing time at most $Cn\log n$ for some $C = C(\Delta, P_{\max}, r_{\min}, \lambda_{\min})$, where $P_{\max}, r_{\min}, \lambda_{\min}$ are defined in Definition 3.2.*

See also Lemma 3.3 for a simple upper bound of $P_{\max}$ in terms of $\Delta, r_{\max}, \lambda_{\max}$.

By Theorem 2.1, to prove Theorem 3.1, it suffices to establish spectral independence and marginal boundedness. We focus on spectral independence in Section 3.1 whose proof is based on the log-concavity of signatures. We give the marginal bound analysis in Section 3.2. The proofs of main results can be found in Section 3.3.

3.1 Spectral independence In this subsection we derive a constant bound of spectral independence via Lemma 2.1.

PROPOSITION 3.1. (SPECTRAL INDEPENDENCE) *Under the assumptions of Theorem 3.1, the Gibbs distribution μ of the Holant problem $(G, \boldsymbol{f}, \boldsymbol{\lambda})$ is $O(P_{\max})$-spectrally independent.*

For spectral independence we need to consider the conditional distribution μ^τ under an arbitrary pinning τ. We note that a pinning τ on a subset $\Lambda \subseteq E$ of edges induces a Holant problem on the subgraph $G \backslash \Lambda$. To formalize this relationship, it is helpful to define the following notation of downward shifting operator.

DEFINITION 3.3. (DOWNWARD SHIFTING) *For a function $f : \mathbb{N} \to \mathbb{R}_{\geq 0}$, we define the function $\mathcal{D}f : \mathbb{N} \to \mathbb{R}_{\geq 0}$ as*

$$(\mathcal{D}f)(k) = f(k+1), \quad \forall k \in \mathbb{N}.$$

We further define $\mathcal{D}^m f = \mathcal{D}(\mathcal{D}^{m-1} f)$ for integer $m \geq 1$.

Let $(G, \boldsymbol{f}, \boldsymbol{\lambda})$ be a Holant problem. For a subset $U \subseteq V$ of vertices, we define $\mathcal{D}_U \boldsymbol{f}$ as

$$(\mathcal{D}_U \boldsymbol{f})_v = \begin{cases} \mathcal{D}f_v, & v \in U, \\ f_v, & v \notin U. \end{cases}$$

For a pinning τ, we define $\mathcal{D}_\tau \boldsymbol{f}$ as for all $v \in V$,

$$(\mathcal{D}_\tau \boldsymbol{f})_v = \mathcal{D}^{|\tau \cap E_v|} f_v.$$

Observe that if an adjacent edge of a vertex v is pinned to be occupied, it corresponds to changing the signature of v from f_v to $\mathcal{D}f_v$. Hence, the Holant problem $(G, \boldsymbol{f}, \boldsymbol{\lambda})$ with pinning τ on $\Lambda \subseteq E$ induces a smaller instance of Holant problem $(G \backslash \Lambda, \mathcal{D}_\tau \boldsymbol{f}, \boldsymbol{\lambda}_{E \backslash \Lambda})$ on the subgraph $G \backslash \Lambda$.

OBSERVATION 3.1. *Consider a Holant problem $(G, \boldsymbol{f}, \boldsymbol{\lambda})$ satisfying the conditions in Theorem 3.1. Then for any pinning the induced Holant problem also satisfies the conditions. Furthermore, all parameters in Definition 3.2 are "monotone" in pinnings; i.e., we have that $r_{\max}, \lambda_{\max}, P_{\max}, \Delta$ are non-increasing under any pinning and $r_{\min}, \lambda_{\min}$ are non-decreasing.*

Proof. The observation follows from that all signatures are log-concave and hence for all $v \in V$,

$$\frac{f_v(k)}{f_v(0)} \geq \frac{f_v(k+\ell)}{f_v(\ell)}$$

where $0 \leq k \leq d_v$ and $0 \leq \ell \leq d_v - k$ (assuming $0/0 = 0$). ∎

By Observation 3.1, μ^τ corresponds to an induced Holant problem still satisfying the conditions in Theorem 3.1; hence it suffices to focus on the no-pinning case. The following proposition gives the key step for bounding the spectral independence constant via Lemma 2.1. It upper bounds the expected number of discrepancies when one signature f_v is changed to $\mathcal{D}f_v$, i.e., the difference between an adjacent (half-)edge of v is occupied and unoccupied.

PROPOSITION 3.2. *Let (G, f, λ) be a Holant problem satisfying the conditions in Theorem 3.1 with Gibbs distribution $\mu = \mu_{G, \boldsymbol{f}, \boldsymbol{\lambda}}$. Suppose $v \in V$ is a vertex with $f_v(1) > 0$, and let $\mu' = \mu_{G, \mathcal{D}_v \boldsymbol{f}, \boldsymbol{\lambda}}$ be the Gibbs distribution of the Holant problem obtained by changing f_v to $\mathcal{D}f_v$. Then we have*

$$W_1(\mu, \mu') \leq P_{\max} - 1.$$

Our coupling between μ and μ' is inspired by [CZ23] which proves spectral independence for weighted even subgraphs with signatures $[1, a, 1, a, \dots]$ for some $a > 0$. Note that such signatures have period two which is crucial for the coupling arguments in [CZ23]. Our new ingredient is to construct a coupling without periodicity of signatures but incorporating the log-concavity in a suitable way.

Proof. [Proof of Proposition 3.2] We construct a simple coupling between $\mu = \mu_{G, \boldsymbol{f}, \boldsymbol{\lambda}}$ and $\mu' = \mu_{G, \mathcal{D}_v \boldsymbol{f}, \boldsymbol{\lambda}}$ using Algorithm 1.

Algorithm 1 Coupling procedure for Holant problems

1: **procedure** COUPLE$(G, \boldsymbol{f}, \boldsymbol{\lambda}, v)$
2: **Input:** $(G, \boldsymbol{f}, \boldsymbol{\lambda})$ a Holant problem, $v \in V$ a disagreeing vertex
3: **Output:** A pair of random configurations $(\sigma, \sigma') \in 2^E \times 2^E$ drawn from a coupling between $\mu = \mu_{G,\boldsymbol{f},\boldsymbol{\lambda}}$ and $\mu' = \mu_{G,\mathcal{D}_v\boldsymbol{f},\boldsymbol{\lambda}}$
4: **if** v is isolated **then**
5: Sample $\sigma \sim \mu$
6: **return** (σ, σ)
7: Choose $e = \{u, v\} \in E$ such that $\mu_e(1) \geq \mu'_e(1)$. ▷ Lemma 3.1
8: Sample (σ_e, σ'_e) from an optimal coupling of (μ_e, μ'_e).
9: **if** $\sigma_e = \sigma'_e = 0$ **then**
10: $(\sigma_{E\setminus e}, \sigma'_{E\setminus e}) \leftarrow$ COUPLE$(G\setminus e, \boldsymbol{f}, \boldsymbol{\lambda}_{E\setminus e}, v)$
11: **else if** $\sigma_e = \sigma'_e = 1$ **then**
12: $(\sigma_{E\setminus e}, \sigma'_{E\setminus e}) \leftarrow$ COUPLE$(G\setminus e, \mathcal{D}_e\boldsymbol{f}, \boldsymbol{\lambda}_{E\setminus e}, v)$
13: **else** ▷ We must have $\sigma_e = 1$, $\sigma'_e = 0$
14: $(\sigma'_{E\setminus e}, \sigma_{E\setminus e}) \leftarrow$ COUPLE$(G\setminus e, \mathcal{D}_v\boldsymbol{f}, \boldsymbol{\lambda}_{E\setminus e}, u)$
15: **return** $(\sigma_e \cup \sigma_{E\setminus e}, \sigma'_e \cup \sigma'_{E\setminus e})$

Coupling. We prove that Algorithm 1 produces a coupling between μ and μ'. We prove this by induction on the number of edges of G.

Base case is when v is isolated. In this case we have $\mu = \mu'$. The algorithm produces the identity coupling.

Induction step: Suppose v is not isolated. By Lemma 3.1, there exists an edge $e \in E_v$ such that $\mu_e(1) \geq \mu'_e(1)$, so Line 7 runs successfully; we shall prove Lemma 3.1 right after this proof. By Line 8, the marginal distributions on edge e are correct. In particular, note that because in Line 8 we choose an optimal coupling, it is impossible to have $\sigma_e = 0$ and $\sigma'_e = 1$ since $\mu_e(1) \geq \mu'_e(1)$. We need to prove that the recursive calls produce the desired distributions on $E\setminus e$.

- Case 1: $\sigma_e = \sigma'_e = 0$. By induction hypothesis,

$$\sigma_{E\setminus e} \sim \mu_{G\setminus e,\boldsymbol{f},\boldsymbol{\lambda}_{E\setminus e}} = \mu_{E\setminus e}(\cdot \mid \sigma_e = 0),$$
$$\sigma'_{E\setminus e} \sim \mu_{G\setminus e,\mathcal{D}_v\boldsymbol{f},\boldsymbol{\lambda}_{E\setminus e}} = \mu'_{E\setminus e}(\cdot \mid \sigma'_e = 0).$$

- Case 2: $\sigma_e = \sigma'_e = 1$. By induction hypothesis,

$$\sigma_{E\setminus e} \sim \mu_{G\setminus e,\mathcal{D}_e\boldsymbol{f},\boldsymbol{\lambda}_{E\setminus e}} = \mu_{E\setminus e}(\cdot \mid \sigma_e = 1),$$
$$\sigma'_{E\setminus e} \sim \mu_{G\setminus e,\mathcal{D}_v\mathcal{D}_e f,\boldsymbol{\lambda}_{E\setminus e}} = \mu'_{E\setminus e}(\cdot \mid \sigma'_e = 1).$$

- Case 3: $\sigma_e = 1$, $\sigma'_e = 0$. By induction hypothesis,

$$\sigma_{E\setminus e} \sim \mu_{G\setminus e,\mathcal{D}_u\mathcal{D}_v\boldsymbol{f},\boldsymbol{\lambda}_{E\setminus e}} = \mu_{E\setminus e}(\cdot \mid \sigma_e = 1),$$
$$\sigma'_{E\setminus e} \sim \mu_{G\setminus e,\mathcal{D}_v\boldsymbol{f},\boldsymbol{\lambda}_{E\setminus e}} = \mu'_{E\setminus e}(\cdot \mid \sigma'_e = 0).$$

In all three cases, we see that $\sigma_{E\setminus e}$ and $\sigma'_{E\setminus e}$ have the desired distributions. So the algorithm returns correctly.

W_1 distance. Let us bound the expected ℓ_1 distance under the coupling generated by Algorithm 1. We prove by induction on the number of edges that the expected ℓ_1 distance is at most $P_{\max} - 1$, as defined in Definition 3.2.

Base case is when v is isolated. In this case the ℓ_1 distance is 0.

Induction step: Suppose that v is not isolated. We consider recursive calls of Algorithm 1 until: (1) the input vertex becomes some other vertex $u \neq v$, or (2) the algorithm halts. We claim that Case (2) happens with probability at least $\mu_{E_v}(\mathbf{0})$, which is the probability that all edges in E_v are unoccupied under the Gibbs distribution μ.

To see this, let $\mathcal{A}$ be the following event: For all $1 \le i \le d$ where $d = d_v = |E_v|$, in the i-th call of Algorithm 1, the algorithm picks an edge from E_v, denoted by e_i, which has not yet be chosen, and sets its values to be $\sigma_{e_i} = \sigma'_{e_i} = 0$ in both samples. Note that if $\mathcal{A}$ occurs then the algorithm halts in the $(d+1)$-th call and the ℓ_1 distance is 0. For simplicity of the proof we assume that there is a total ordering of edges and in Line 7 we always pick the smallest edge satisfying the requirement. Hence, the sequence of edges $(e_1, \dots, e_d)$ associated with the event $\mathcal{A}$ is fixed and deterministic: e_i is the smallest edge in $E_v \backslash \{e_1, \dots, e_{i-1}\}$ such that

$$\mu(\sigma_e = 1 \mid \sigma_{e_1} = \cdots = \sigma_{e_{i-1}} = 0) \ge \mu'(\sigma'_e = 1 \mid \sigma'_{e_1} = \cdots = \sigma'_{e_{i-1}} = 0). \tag{3.1}$$

It follows that

$$\begin{aligned}
\Pr(\text{Case (2)}) \ge \Pr(\mathcal{A}) &= \prod_{i=1}^{d} \Pr\left(\sigma_{e_i} = \sigma'_{e_i} = 0 \,\middle|\, \forall j < i, \sigma_{e_j} = \sigma'_{e_j} = 0\right) \\
&\overset{(*)}{=} \prod_{i=1}^{d} \mu(\sigma_e = 0 \mid \sigma_{e_1} = \cdots = \sigma_{e_{i-1}} = 0) \\
&= \mu(\sigma_{E_v} = \mathbf{0}),
\end{aligned}$$

where $(*)$ is because of Eq. (3.1) and the optimal coupling of (σ_e, σ'_e).

To summarize, Case (2) happens with probability at least $\mu_{E_v}(\mathbf{0})$ and the ℓ_1 distance is 0 in this case; Case (1) happens with probability at most $1 - \mu_{E_v}(\mathbf{0})$ and the expected ℓ_1 distance is at most $1 + (P_{\max} - 1) = P_{\max}$, where 1 comes from the edge $\{u, v\}$ and $P_{\max} - 1$ comes from induction hypothesis. Therefore, the expected ℓ_1 distance produced by the root call is at most

$$(1 - \mu_{E_v}(\mathbf{0}))\, P_{\max} \le P_{\max} - 1, \tag{3.2}$$

where the inequality follows from $\mu_{E_v}(\mathbf{0}) \ge 1/P_{\max}$ by Lemma 3.2, whose proof is technical and postponed to Section 3.2. This completes the induction. ∎

Note that as long as it holds $\mu_{E_v}(\mathbf{0}) = \Omega(1)$ one can deduce $O(1)$ Wasserstein distance using the inductive argument in Proposition 3.2.

We now present and prove Lemma 3.1 which is crucially used in Line 7 of Algorithm 1.

Lemma 3.1. *Let $(G, \boldsymbol{f}, \boldsymbol{\lambda})$ be a Holant problem and $v \in V$ be a vertex such that $\mu = \mu_{G,\boldsymbol{f},\boldsymbol{\lambda}}$ and $\mu' = \mu_{G,\mathcal{D}_v\boldsymbol{f},\boldsymbol{\lambda}}$ are both well-defined. If f_v is a log-concave signature, then we have*

$$\mathbb{E}_{\sigma\sim\mu}|\sigma \cap E_v| \ge \mathbb{E}_{\sigma'\sim\mu'}|\sigma' \cap E_v|. \tag{3.3}$$

In particular, there exists $e \in E_v$ such that

$$\mu(\sigma_e = 1) \ge \mu'(\sigma'_e = 1). \tag{3.4}$$

Proof. Let $d_v = |E_v|$ be the degree of v in G. For $0 \le k \le d_v$, define

$$\Phi_k := \sum_{\sigma \in 2^E :\, |\sigma \cap E_v| = k} \prod_{u \in V \backslash v} f_u(|\sigma \cap E_u|) \prod_{e \in E :\, \sigma_e = 1} \lambda_e.$$

Then we have that

$$\begin{aligned}
\mathbb{E}_{\sigma\sim\mu}|\sigma \cap E_v| &= \frac{\sum_{k=0}^{d_v} k f_v(k) \Phi_k}{\sum_{k=0}^{d_v} f_v(k) \Phi_k}, \\
\mathbb{E}_{\sigma'\sim\mu'}|\sigma' \cap E_v| &= \frac{\sum_{k=0}^{d_v} k f_v(k+1) \Phi_k}{\sum_{k=0}^{d_v} f_v(k+1) \Phi_k}.
\end{aligned}$$

It follows that

$$\begin{aligned}
&\left(\sum_{k=0}^{d_v} k f_v(k)\Phi_k\right)\left(\sum_{\ell=0}^{d_v} f_v(\ell+1)\Phi_\ell\right) - \left(\sum_{k=0}^{d_v} f_v(k)\Phi_k\right)\left(\sum_{\ell=0}^{d_v} \ell f_v(\ell+1)\Phi_\ell\right)\\
&= \sum_{k,\ell=0}^{d_v} (k-\ell) f_v(k) f_v(\ell+1)\Phi_k\Phi_\ell\\
&= \frac{1}{2}\sum_{k,\ell=0}^{d_v} (k-\ell)\big(f_v(k)f_v(\ell+1) - f_v(k+1)f_v(\ell)\big)\Phi_k\Phi_\ell \geq 0,
\end{aligned}$$

where the last step is because of the log-concavity assumption of f_v. Therefore Eq. (3.3) holds as desired, and Eq. (3.4) is an immediate consequence of Eq. (3.3). ∎

We are now ready to prove Proposition 3.1 on spectral independence.

Proof. [Proof of Proposition 3.1] We prove spectral independence via Lemma 2.1. Let τ be a pinning on $\Lambda \subseteq E$ and let $e = \{u, v\} \in E\backslash\Lambda$ such that $0 < \mu_e^\tau(1) < 1$. For two pinnings $\tau' = \tau \cup \{e \leftarrow 0\}$, $\tau'' = \tau \cup \{e \leftarrow 1\}$ on $\Lambda \cup \{e\}$ which differ only at e, we consider the two conditional distributions

$$\begin{aligned}
\mu^{\tau'} &= \mu_{G,\boldsymbol{f},\boldsymbol{\lambda}}(\cdot \mid \sigma_\Lambda = \tau, \sigma_e = 0) = \mu_{\widetilde{G},\widetilde{\boldsymbol{f}},\widetilde{\boldsymbol{\lambda}}},\\
\mu^{\tau''} &= \mu_{G,\boldsymbol{f},\boldsymbol{\lambda}}(\cdot \mid \sigma_\Lambda = \tau, \sigma_e = 1) = \mu_{\widetilde{G},\mathcal{D}_v\mathcal{D}_u\widetilde{\boldsymbol{f}},\widetilde{\boldsymbol{\lambda}}},
\end{aligned}$$

where $\widetilde{G} = G\backslash\Lambda\backslash\{e\}$, $\widetilde{\boldsymbol{f}} = \mathcal{D}_\tau \boldsymbol{f}$, and $\widetilde{\boldsymbol{\lambda}} = \boldsymbol{\lambda}_{E\backslash\Lambda\backslash\{e\}}$. By Proposition 3.2 and the triangle inequality, we deduce that

$$W_1(\mu^{\tau'}, \mu^{\tau''}) \leq W_1\left(\mu_{\widetilde{G},\widetilde{\boldsymbol{f}},\widetilde{\boldsymbol{\lambda}}}, \mu_{\widetilde{G},\mathcal{D}_u\widetilde{\boldsymbol{f}},\widetilde{\boldsymbol{\lambda}}}\right) + W_1\left(\mu_{\widetilde{G},\mathcal{D}_u\widetilde{\boldsymbol{f}},\widetilde{\boldsymbol{\lambda}}}, \mu_{\widetilde{G},\mathcal{D}_v\mathcal{D}_u\widetilde{\boldsymbol{f}},\widetilde{\boldsymbol{\lambda}}}\right) \leq 2(P_{\max} - 1).$$

Note that all these Holant sub-problems under downward shifting operators still satisfy the assumptions of Theorem 3.1 by Observation 3.1. Therefore, by Lemma 2.1 $\mu_{G,\boldsymbol{f},\boldsymbol{\lambda}}$ is $2(P_{\max} - 1)$-spectrally independent. ∎

3.2 Marginal bounds We first lower bound the probability that all adjacent edges of a vertex are unoccupied, which justifies Eq. (3.2).

LEMMA 3.2. *Let $(G, \boldsymbol{f}, \boldsymbol{\lambda})$ be a Holant problem satisfying the assumptions of Theorem 3.1. Then for any vertex v, we have*

$$\mu_{E_v}(\mathbf{0}) \geq \frac{1}{P_v(r_{\max}\lambda_{\max})}.$$

Proof. For any $\sigma \in 2^{E\backslash E_v}$ and $\tau \in 2^{E_v}$ viewed as subsets, suppose $\tau = \{e_1, \ldots, e_k\} \subseteq E_v$ where $k = |\tau|$ and $e_i = \{u_i, v\}$ for each i, and we have that

$$\begin{aligned}
\frac{\mu(\sigma \cup \tau)}{\mu(\sigma)} &\leq \frac{f_v(k)}{f_v(0)}\left(\prod_{i=1}^{k} \max_{1\leq \ell \leq d_{u_i}} \frac{f_{u_i}(\ell)}{f_{u_i}(\ell-1)}\right)\left(\prod_{i=1}^{k}\lambda_{e_i}\right)\\
&= \frac{f_v(k)}{f_v(0)}\prod_{i=1}^{k}\left(\frac{f_{u_i}(1)}{f_{u_i}(0)}\lambda_{e_i}\right)\\
&\leq \frac{f_v(k)}{f_v(0)}(r_{\max}\lambda_{\max})^k.
\end{aligned}$$

Summing over σ and τ, we obtain that

$$\begin{aligned}
1 &= \sum_{\sigma \in 2^{E \setminus E_v}} \sum_{\tau \in 2^{E_v}} \mu(\sigma \cup \tau) \\
&\leq \sum_{\sigma \in 2^{E \setminus E_v}} \sum_{\tau \in 2^{E_v}} \frac{f_v(|\tau|)}{f_v(0)} (r_{\max}\lambda_{\max})^{|\tau|} \mu(\sigma) \\
&= \sum_{\sigma \in 2^{E \setminus E_v}} \mu(\sigma) \sum_{\tau \in 2^{E_v}} \frac{f_v(|\tau|)}{f_v(0)} (r_{\max}\lambda_{\max})^{|\tau|} \\
&\leq \mu_{E_v}(\mathbf{0}) P_v(r_{\max}\lambda_{\max}),
\end{aligned}$$

as claimed. ∎

We now give marginal lower bounds.

PROPOSITION 3.3. (MARGINAL BOUNDEDNESS) *Let $(G, \boldsymbol{f}, \boldsymbol{\lambda})$ be a Holant problem satisfying the assumptions of Theorem 3.1. Then for any pinning τ on a subset $\Lambda \subseteq E$ and for any edge $e \in E \setminus \Lambda$ that can be occupied, we have*

$$\mu^\tau(\sigma_e = 0) \geq \frac{1}{P_{\max}} \quad \text{and} \quad \mu^\tau(\sigma_e = 1) \geq \frac{r_{\min}^2 \lambda_{\min}}{P_{\max}^2}.$$

Proof. By the monotonicity of pinning given in Observation 3.1 it suffices to focus on the no-pinning case. Let $e = \{u, v\}$. The first part follows from Lemma 3.2. Let us prove the second part.

By applying Lemma 3.2 twice, we have

$$\mu(\sigma_{E_u \cup E_v} = \mathbf{0}) = \mu(\sigma_{E_u} = \mathbf{0}) \cdot \mu(\sigma_{E_v \setminus E_u} = \mathbf{0} \mid \sigma_{E_u} = \mathbf{0}) \geq \frac{1}{P_{\max}^2},$$

where we note that $\mu(\cdot \mid \sigma_{E_u} = \mathbf{0})$ corresponds to an induced Holant problem whose $P_{\max}$ is no bigger by Observation 3.1. Then we have

$$\begin{aligned}
\mu(\sigma_e = 1) &\geq \mu(\sigma_e = 1, \sigma_{E_u \cup E_v \setminus \{e\}} = \mathbf{0}) \\
&= \frac{f_u(1) f_v(1)}{f_u(0) f_v(0)} \lambda_e \cdot \mu(\sigma_{E_u \cup E_v} = \mathbf{0}) \\
&\geq \frac{r_{\min}^2 \lambda_{\min}}{P_{\max}^2}.
\end{aligned}$$

This verifies the second part. ∎

The following lemma provides an upper bound on $P_{\max}$ which in turn gives simple and clean constant bounds on spectral independence and marginal boundedness.

LEMMA 3.3. *For a Holant problem $(G, \boldsymbol{f}, \boldsymbol{\lambda})$ with log-concave signatures and $f_v(0) > 0$ for all $v \in V$, we have*

$$P_{\max} \leq \left(r_{\max}^2 \lambda_{\max} + 1\right)^{\Delta}.$$

Proof. We note that for any vertex v,

$$\frac{f_v(k)}{f_v(0)} = \prod_{i=1}^{k} \frac{f_v(i)}{f_v(i-1)} \leq \left(\frac{f_v(1)}{f_v(0)}\right)^k \leq r_{\max}^k$$

by the log-concavity of f_v. Hence, we deduce that

$$P_{f_v}(x) = \frac{1}{f_v(0)} \sum_{k=0}^{d_v} \binom{d_v}{k} f_v(k) x^k \leq \sum_{k=0}^{d_v} \binom{d_v}{k} r_{\max}^k x^k = (1 + r_{\max} x)^{d_v}.$$

Letting $x = r_{\max}\lambda_{\max}$ and taking maximum over v gives the desired bound. ∎

3.3 Proofs of main results

Proof. [Proof of Theorem 3.1] By Proposition 3.1, μ is $O(P_{\max})$-spectrally independent. By Proposition 3.3, μ is b-marginally bounded for $b = b(P_{\max}, r_{\min}, \lambda_{\min}) > 0$. Using Theorem 2.1, we finish the proof. ∎

Proof. [Proof of Theorem 1.1] Because $[1, \dots, 1, 0, \dots, 0]$ is a log-concave signature, we can apply Theorem 3.1. For the b-matching problem $(G, \boldsymbol{b}, \lambda)$, we have $r_{\max} = r_{\min} = 1$ and $\lambda_{\max} = \lambda_{\min} = \lambda$. Therefore $P_{\max}$ is upper bounded by a function of Δ and λ. So Theorem 1.1 is a corollary of Theorem 3.1. ∎

Remark 3.1. *Note that for uniformly random b-matchings on a graph of maximum degree Δ where $1 \le b \le \Delta$, we have $r_{\max} = r_{\min} = 1$, $\lambda_{\max} = \lambda_{\min} = 1$, and*

$$P_{\max} \le \sum_{k=0}^{b} \binom{\Delta}{k}.$$

Hence, for small $b \ll \Delta$ we have $O(\Delta^b)$-spectral independence, and for all b we have $O(2^\Delta)$-spectral independence.

3.4 Further discussions Theorem 3.1 establishes spectral independence for a broad class of Holant problems with log-concave signatures. The support of the corresponding Gibbs distribution is over all $\boldsymbol{b}$-matchings for some $\boldsymbol{b} \in \mathbb{N}^V$. Note that $O(1)$-spectral independence fails for the signature $[0, 1, 0, \dots, 0]$ which corresponds to perfect matchings. Moreover, it also fails for the signature $[0, 1, \dots, 1, 0]$ in the ℓ_∞ sense (maximum absolute row sum of influence matrices), as illustrated by the examples below.

- Consider a path $v_0 \leftrightarrow v_1 \leftrightarrow \cdots \leftrightarrow v_n$. Suppose the signature at v_{2i-1} is $[1, 1, 0]$ for all $i \ge 1$, i.e., requiring at least one adjacent edge unoccupied, and the signature at v_{2i} is $[0, 1, 1]$ for all $i \ge 1$, i.e., requiring at least one adjacent edge occupied. If $\sigma(v_0 v_1) = 1$, then there is only one feasible configuration where $\sigma(v_{2i-1} v_{2i}) = 0$ and $\sigma(v_{2i} v_{2i+1}) = 1$. If $\sigma(v_0 v_1) = 0$, then one can check that there are n feasible configurations. The absolute sum of influences of edge $v_0 v_1$ on all other edges is $\Omega(n)$.
- Consider a path $v_0 \leftrightarrow v_1 \leftrightarrow \cdots \leftrightarrow v_n$ together with edges $u_i v_i$ where $1 \le i < n$. Suppose every vertex v_i has signature $[0, 1, 1, 0]$ for $i \ge 1$. Fix $\sigma(u_{2i-1} v_{2i-1}) = 1$ and $\sigma(u_{2i} v_{2i}) = 0$ for all $i \ge 1$ as a pinning, then the resulted Holant problem is equivalent to the previous example. So the absolute sum of influences of $v_0 v_1$ is $\Omega(n)$.

4 Hardcore model on claw-free graphs

In this section we prove Theorem 1.3 from introduction. The proof idea is similar to Theorem 3.1 on a high level. We establish an upper bound on the Wasserstein W_1 distance between distributions under different pinnings, which implies spectral independence and fast mixing via Lemma 2.1 and Theorem 2.1. The W_1 upper bound is proved using a recursive coupling (Algorithm 2), which shares some similarities with Algorithm 1 but requires new ideas to utilize claw-freeness of the graph.

While the hardcore model is a distribution on 2^V rather than 2^E, notations in Section 2 can be applied here with simple changes. For example, we view a subset $I \subseteq V$ equivalently as a binary indicator vector $\sigma = \mathbf{1}_I \in \{0, 1\}^V$. We similarly define pinnings, marginal boundedness, influence matrices, and spectral independence for the hardcore model; importantly, Theorem 2.1 and Lemma 2.1 still hold. We refer to [ALO20, CLV21] for formal definitions and precise statements.

4.1 Preliminaries on claw-free graphs A graph is called *claw-free* if it does not contain $K_{1,3}$ (that is, a star graph comprising three edges, three leaves, and a central vertex) as an induced subgraph. Note that any induced subgraph of a claw-free graph is still claw-free by definition. The class of claw-free graphs includes in particular all line graphs.

For a vertex v of a graph $G = (V, E)$, let $N_v = \{u \in V : (u, v) \in E\}$ denote the neighborhood of v, and let $N_v^* = v \cup N_v$ denote the closed neighborhood including v itself. We say a vertex v is *simplicial* if its neighborhood N_v (or, equivalently, its closed neighborhood N_v^*) forms a clique, i.e., every two neighbors of v are adjacent. The following simple lemma is very helpful to us.

CLAIM 4.1. *Suppose $G = (V, E)$ is a claw-free graph and $v \in V$ is a vertex. Let $u \in N_v$ be a neighbor of v. Then the subgraph $G \backslash (N_v^* \backslash u)$ is claw-free and u is a simplicial vertex of $G \backslash (N_v^* \backslash u)$.*

Proof. The subgraph $G \backslash (N_v^* \backslash u)$ is claw-free since it is an induced subgraph on $V \backslash (N_v^* \backslash u)$. Suppose for contradiction that u is not simplicial in $G \backslash (N_v^* \backslash u)$. By definition there exist two neighbors w_1, w_2 of u in $G \backslash (N_v^* \backslash u)$ that are not adjacent. Then, $\{u, v, w_1, w_2\}$ forms a claw centered at u in G: u is adjacent to all of v, w_1, w_2 by our choice, and w_1, w_2 are not adjacent to v since $w_1, w_2 \notin N_v^*$. This is a contradiction and hence u must be simplicial. ∎

4.2 Spectral independence We establish spectral independence in this subsection. For a vertex v and a spin $i \in \{0, 1\}$, we use the notation $v \leftarrow i$ to represent the pinning $\sigma_v = i$. Hence, if $\Lambda \subseteq V$ is a subset of vertices, τ is a pinning on Λ, and $v \in V \backslash \Lambda$, then $\mu^{\tau, v \leftarrow i}$ represents the conditional Gibbs distribution on the subset $V \backslash \Lambda \backslash v$ conditioned on both $\sigma_\Lambda = \tau$ and $\sigma_v = i$.

PROPOSITION 4.1. (SPECTRAL INDEPENDENCE) *Work under the setting of Theorem 1.3. For any pinning τ on a subset $\Lambda \subseteq V$ and any vertex $v \in V \backslash \Lambda$, we have*

$$W_1\left(\mu^{\tau, v \leftarrow 0}, \mu^{\tau, v \leftarrow 1}\right) \leq 2(1 + \Delta\lambda_{\max}).$$

In particular, μ is $2(1 + \Delta\lambda_{\max})$-spectrally independent.

As a standard trick for the hardcore model, we can view the the conditional Gibbs distribution under a pinning τ on $\Lambda \subseteq V$ as a hardcore model on an induced subgraph by removing all pinned vertices in Λ together with all neighbors of those fixed to be occupied (pinned to spin 1). Note that the maximum degree of the resulting subgraph does not increase. Thus, for simplicity we consider the case without pinnings except at v.

We first argue that for claw-free graphs, to prove Proposition 4.1 it suffices to consider only when v is a simplicial vertex. Sample $\xi \sim \mu_{N_v}^{v \leftarrow 0}$ and $\xi' \sim \mu_{N_v}^{v \leftarrow 1}$ two configurations on the neighborhood of v. Observe that we must have $\xi' = \mathbf{0}$ since v is occupied. Suppose $\|\xi\|_1 = m$ and denote the occupied vertices in ξ by $\{u_1, \ldots, u_m\}$ under any ordering. We must have $m \leq 2$; otherwise, since there can be no edge between any pair of occupied vertices, the set $\{v, u_1, u_2, u_3\}$ forms a claw. For $0 \leq i \leq m$, let ξ^i be the configuration on N_v with $\{u_1, \ldots, u_i\}$ occupied and all other vertices in N_v unoccupied, so $\xi^0 = \xi' = \mathbf{0}$ and $\xi^m = \xi$. We note that each ξ^i is feasible because ξ is feasible. We then deduce from the triangle inequality that

$$\begin{aligned} W_1\left(\mu^{v \leftarrow 0}, \mu^{v \leftarrow 1}\right) &\leq \mathbb{E}[\|\xi - \xi'\|_1] + \mathbb{E}\left[W_1\left(\mu^{v \leftarrow 0, \xi}, \mu^{v \leftarrow 1, \xi'}\right)\right] \\ &= \mathbb{E}[\|\xi - \xi'\|_1] + \mathbb{E}\left[W_1\left(\mu^{v \leftarrow 0, \xi^m}, \mu^{v \leftarrow 0, \xi^0}\right)\right] \\ &\leq 2 + \mathbb{E}\left[\sum_{i=1}^{m} W_1\left(\mu^{v \leftarrow 0, \xi^{i-1}}, \mu^{v \leftarrow 0, \xi^i}\right)\right]. \end{aligned} \tag{4.5}$$

For each i, the Wasserstein distance $W_1\left(\mu^{v \leftarrow 0, \xi^{i-1}}, \mu^{v \leftarrow 0, \xi^i}\right)$ corresponds to a hardcore model on the subgraph $G \backslash (N_v^* \backslash u_i) \backslash \bigcup_{j=1}^{i-1} N_{u_j}$ under two pinnings $u_i \leftarrow 0$ and $u_i \leftarrow 1$; the pinning on v does not matter since the configuration on N_v is fixed. In particular, u_i is a simplicial vertex in this subgraph since it is already simplicial in $G \backslash (N_v^* \backslash u_i)$ by Claim 4.1. Thus, Eq. (4.5) shows that it suffices to consider the case where the disagreeing vertex is simplicial.

We show in the next proposition that the Wasserstein distance is constant between Gibbs distributions under different pinnings on a simplicial vertex.

PROPOSITION 4.2. *Work under the setting of Theorem 1.3. If v is a simplicial vertex, then we have*

$$W_1\left(\mu^{v \leftarrow 0}, \mu^{v \leftarrow 1}\right) \leq \Delta\lambda_{\max}.$$

We can then deduce Proposition 4.1 from Proposition 4.2 by the arguments above.

Proof. [Proof of Proposition 4.1] Combining Eq. (4.5) and Proposition 4.2, we deduce that

$$W_1(\mu^{v \leftarrow 0}, \mu^{v \leftarrow 1}) \leq 2 + 2\Delta\lambda_{\max}.$$

Spectral independence then follows from Lemma 2.1 and the bound on W_1. ∎

Algorithm 2 Coupling procedure for the hardcore model on claw-free graphs

1: **procedure** COUPLE$(G, \boldsymbol{\lambda}, v)$
2: **Input:** $G = (V, E)$ a claw-free graph, $\boldsymbol{\lambda} \in \mathbb{R}^V_{>0}$ a vector of fugacity, $v \in V$ a simplicial vertex which is disagreeing
3: **Output:** A pair of random configurations $(\sigma, \sigma') \in 2^{V \setminus v} \times 2^{V \setminus v}$ drawn from a coupling between $\mu^{v \leftarrow 0}_{G,\boldsymbol{\lambda}}$ and $\mu^{v \leftarrow 1}_{G,\boldsymbol{\lambda}}$
4: **if** v is isolated **then**
5: Sample $\sigma \sim \mu^{v \leftarrow 0}_{G,\boldsymbol{\lambda}}$
6: **return** (σ, σ)
7: Sample $\sigma_{N_v} \sim \left(\mu^{v \leftarrow 0}_{G,\boldsymbol{\lambda}}\right)_{N_v}$ and $\sigma'_{N_v} \leftarrow \mathbf{0}$
8: **if** $\sigma_{N_v} = \mathbf{0}$ **then**
9: Sample $\sigma_{V \setminus N^*_v} \sim \mu^{v \leftarrow 0, N_v \leftarrow \mathbf{0}}_{G,\boldsymbol{\lambda}}$
10: **return** (σ, σ)
11: **else** ▷ We must have $\|\sigma_{N_v}\|_1 = 1$
12: Let $u \in N_v$ such that $\sigma_u = 1$
13: $\left(\sigma'_{V \setminus (N^*_v \setminus u)}, \sigma_{V \setminus (N^*_v \setminus u)}\right) \leftarrow$ COUPLE$\left(G \setminus (N^*_v \setminus u), \boldsymbol{\lambda}_{V \setminus (N^*_v \setminus u)}, u\right)$ ▷ Claim 4.1
14: **return** (σ, σ')

It remains to prove Proposition 4.2, which is again proved by a recursive coupling.

Proof. [Proof of Proposition 4.2]

We construct a coupling between $\mu^{v \leftarrow 0}_{G,\boldsymbol{\lambda}}$ and $\mu^{v \leftarrow 1}_{G,\boldsymbol{\lambda}}$ for configurations on $V \setminus v$ using Algorithm 2, and upper bound the W_1 distance via claw-freeness of the graph.

Coupling. We prove that Algorithm 2 produces a valid coupling by induction on the number of vertices in G. If v is isolated then the configuration on $V \setminus v$ is independent of σ_v, and hence $\mu^{v \leftarrow 0}_{G,\boldsymbol{\lambda}} = \mu^{v \leftarrow 1}_{G,\boldsymbol{\lambda}}$; this justifies Line 6 and also the base case for our induction. For non-isolated v, we sample $\sigma_{N_v} \sim \left(\mu^{v \leftarrow 0}_{G,\boldsymbol{\lambda}}\right)_{N_v}$ and $\sigma'_{N_v} \sim \left(\mu^{v \leftarrow 1}_{G,\boldsymbol{\lambda}}\right)_{N_v}$; note that we must have $\sigma'_{N_v} = \mathbf{0}$ for the latter since v is occupied. If $\sigma_{N_v} = \mathbf{0} = \sigma'_{N_v}$, then the configuration on the remaining graph is independent of σ_v, namely $\mu^{v \leftarrow 0, N_v \leftarrow \mathbf{0}}_{G,\boldsymbol{\lambda}} = \mu^{v \leftarrow 1, N_v \leftarrow \mathbf{0}}_{G,\boldsymbol{\lambda}}$, which justifies Line 10. Otherwise, we have $\sigma_{N_v} \neq \mathbf{0}$. Since v is simplicial, N_v is a clique and hence there is exactly one vertex in N_v that is occupied under σ_{N_v}, which we denote by u. Let $A = N^*_v \setminus u$ be the closed neighborhood at v excluding u, and we have $\sigma_A = \mathbf{0}$. By the induction hypothesis we have in Line 13 that

$$\sigma_{V \setminus A} \sim \mu^{u \leftarrow 1}_{G \setminus A, \boldsymbol{\lambda}_{V \setminus A}} = \mu^{v \leftarrow 0, N_v \leftarrow \sigma_{N_v}}_{G,\boldsymbol{\lambda}};$$
$$\sigma'_{V \setminus A} \sim \mu^{u \leftarrow 0}_{G \setminus A, \boldsymbol{\lambda}_{V \setminus A}} = \mu^{v \leftarrow 1, N_v \leftarrow \sigma'_{N_v}}_{G,\boldsymbol{\lambda}}.$$

Notice that we can recursively call Algorithm 2 on the input $(G \setminus A, \boldsymbol{\lambda}_{V \setminus A}, u)$ because $G \setminus A$ is a claw-free graph and u is a simplicial vertex by Claim 4.1. Thus, $(\sigma_{V \setminus A}, \sigma'_{V \setminus A})$ comes from the desired conditional distributions and therefore the output of Algorithm 2 is from a coupling of $\mu^{v \leftarrow 0}_{G,\boldsymbol{\lambda}}$ and $\mu^{v \leftarrow 1}_{G,\boldsymbol{\lambda}}$ by induction.

W_1 distance. Next, we bound the expected ℓ_1 distance under the coupling generated by Algorithm 2 by induction on the number of vertices. Base case is when v is isolated, in which case the ℓ_1 distance is 0. Now suppose that v is not isolated. Observe that in one run of Algorithm 2: either $\sigma_{N_v} = \mathbf{0}$ and the ℓ_1 distance is 0, or $\|\sigma_{N_v}\| = 1$ and it recursively calls Algorithm 2 in Line 13 on a smaller instance so that the combined expected ℓ_1 distance is at most $1 + \Delta\lambda_{\max}$ by our induction hypothesis (1 for the discrepancy at u and $\Delta\lambda_{\max}$ for the recursive call). We note that the first case happens with probability exactly $\mu^{v \leftarrow 0}_{G,\boldsymbol{\lambda}}(\sigma_{N_v} = \mathbf{0})$, the probability that all neighbors of v are unoccupied. Therefore, the expected ℓ_1 distance produced by the root call is at most

$$\left(1 - \mu^{v \leftarrow 0}_{G,\boldsymbol{\lambda}}(\sigma_{N_v} = \mathbf{0})\right)(1 + \Delta\lambda_{\max}). \tag{4.6}$$

It remains to lower bound $\mu_{G,\boldsymbol{\lambda}}^{v\leftarrow 0}(\sigma_{N_v}=\mathbf{0})$. Since v is simplicial, every feasible configuration ξ on N_v satisfies $\xi=\mathbf{0}$ or $\|\xi\|_1=1$. We have

$$1=\mu_{G,\boldsymbol{\lambda}}^{v\leftarrow 0}(\sigma_{N_v}=\mathbf{0})+\sum_{\xi\in 2^{N_v}:\,\|\xi\|_1=1}\mu_{G,\boldsymbol{\lambda}}^{v\leftarrow 0}(\sigma_{N_v}=\xi)\le(1+\Delta\lambda_{\max})\cdot\mu_{G,\boldsymbol{\lambda}}^{v\leftarrow 0}(\sigma_{N_v}=\mathbf{0}),$$

and thus $\mu_{G,\boldsymbol{\lambda}}^{v\leftarrow 0}(\sigma_{N_v}=\mathbf{0})\ge(1+\Delta\lambda_{\max})^{-1}$. Plugging into Eq. (4.6) finishes the proof. ∎

4.3 Marginal bounds

In this subsection we give marginal bounds that are needed.

PROPOSITION 4.3. (MARGINAL BOUNDEDNESS) *Work under the setting of Theorem 1.3. For any pinning τ on a subset $\Lambda\subseteq V$ and any vertex $v\in V\backslash\Lambda$ that can be occupied, we have*

$$\mu^\tau(\sigma_v=0)\ge\frac{1}{1+\lambda_{\max}}\quad\textit{and}\quad\mu^\tau(\sigma_v=1)\ge\frac{\lambda_{\min}}{(1+\lambda_{\min})\left(1+\Delta\lambda_{\max}+\binom{\Delta}{2}\lambda_{\max}^2\right)}.$$

Proof. For any configuration $\sigma_1\in 2^{V\backslash(\Lambda\cup v)}$ with non-zero probability under μ^τ, we have

$$\mu^\tau(\sigma_1,\sigma_v=1)\le\lambda_v\mu^\tau(\sigma_1,\sigma_v=0).$$

Therefore

$$\mu^\tau(\sigma_v=0)=\sum_{\sigma_1\in 2^{V\backslash(\Lambda\cup v)}}\mu^\tau(\sigma_1,\sigma_v=0)\ge\frac{1}{1+\lambda_v}\sum_{\sigma_1\in 2^{V\backslash(\Lambda\cup v)}}\mu^\tau(\sigma_1)=\frac{1}{1+\lambda_v}\ge\frac{1}{1+\lambda_{\max}}.$$

This proves the first part.

For the second part, note that for claw-free graphs, at most two neighboring vertices of v can be occupied at the same time, and hence we have

$$1=\sum_{\xi\in 2^{N_v}:\,\|\xi\|_1\le 2}\mu^\tau(\sigma_{N_v}=\xi)\le\left(1+\Delta\lambda_{\max}+\binom{\Delta}{2}\lambda_{\max}^2\right)\mu^\tau(\sigma_{N_v}=\mathbf{0}).$$

So $\mu^\tau(\sigma_{N_v}=\mathbf{0})\ge\left(1+\Delta\lambda_{\max}+\binom{\Delta}{2}\lambda_{\max}^2\right)^{-1}$ and it follows that

$$\begin{aligned}\mu^\tau(\sigma_v=1)&\ge\mu^\tau(\sigma_v=1,\sigma_{N_v}=\mathbf{0})\\&=\frac{\lambda_v}{1+\lambda_v}\cdot\mu^\tau(\sigma_{N_v}=\mathbf{0})\\&\ge\frac{\lambda_{\min}}{(1+\lambda_{\min})\left(1+\Delta\lambda_{\max}+\binom{\Delta}{2}\lambda_{\max}^2\right)}.\end{aligned}$$

This proves the second part. ∎

We are now ready to prove Theorem 1.3.

Proof. [Proof of Theorem 1.3] Follows from Propositions 4.1 and 4.3 and Theorem 2.1. ∎

References

[AJK+22] Nima Anari, Vishesh Jain, Frederic Koehler, Huy Tuan Pham, and Thuy-Duong Vuong. Entropic independence: optimal mixing of down-up random walks. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1418–1430, 2022.

[ALO20] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1319–1330, 2020.

[Bar16] Alexander Barvinok. *Combinatorics and Complexity of Partition Functions*, volume 30. Springer Algorithms and Combinatorics, 2016.

[BCC+22] Antonio Blanca, Pietro Caputo, Zongchen Chen, Daniel Parisi, Daniel Štefankovič, and Eric Vigoda. On mixing of Markov chains: coupling, spectral independence, and entropy factorization. *Electronic Journal of Probability*, 27:1–42, 2022.

[BCR21] Ferenc Bencs, Péter Csikvári, and Guus Regts. Some applications of Wagner's weighted subgraph counting polynomial. *The Electronic Journal of Combinatorics*, 28(4), 2021.

[BR09] Ivona Bezáková and William A. Rummler. Sampling edge covers in 3-regular graphs. In *Proceedings of the 34th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pages 137–148, 2009.

[Brä15] Petter Brändén. Unimodality, log-concavity, real-rootedness and beyond. *Handbook of enumerative combinatorics*, 87:437, 2015.

[CE22] Yuansi Chen and Ronen Eldan. Localization schemes: A framework for proving mixing bounds for Markov chains. In *Proceedings of the 63rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 110–122. IEEE, 2022.

[CFYZ21] Xiaoyu Chen, Weiming Feng, Yitong Yin, and Xinyuan Zhang. Rapid mixing of Glauber dynamics via spectral independence for all degrees. In *Proceedings of the 62nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 137–148. IEEE, 2021.

[CFYZ22] Xiaoyu Chen, Weiming Feng, Yitong Yin, and Xinyuan Zhang. Optimal mixing for two-state anti-ferromagnetic spin systems. In *Proceedings of the 63rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 588–599. IEEE, 2022.

[CLMM23] Zongchen Chen, Kuikui Liu, Nitya Mani, and Ankur Moitra. Strong spatial mixing for colorings on trees and its algorithmic applications. *arXiv preprint arXiv:2304.01954*, 2023.

[CLV21] Zongchen Chen, Kuikui Liu, and Eric Vigoda. Optimal mixing of Glauber dynamics: Entropy factorization via high-dimensional expansion. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1537–1550, 2021.

[CLV22] Zongchen Chen, Kuikui Liu, and Eric Vigoda. Spectral independence via stability and applications to Holant-type problems. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 149–160. IEEE, 2022.

[CMM23] Zongchen Chen, Nitya Mani, and Ankur Moitra. From algorithms to connectivity and back: finding a giant component in random k-SAT. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3437–3470. SIAM, 2023.

[CS07] Maria Chudnovsky and Paul Seymour. The roots of the independence polynomial of a clawfree graph. *Journal of Combinatorial Theory, Series B*, 97(3):350–357, 2007.

[CZ23] Xiaoyu Chen and Xinyuan Zhang. A near-linear time sampler for the Ising model with external field. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4478–4503. SIAM, 2023.

[DGM21] Martin Dyer, Catherine Greenhill, and Haiko Müller. Counting independent sets in graphs with bounded bipartite pathwidth. *Random Structures & Algorithms*, 59(2):204–237, 2021.

[DHJM21] Martin Dyer, Marc Heinrich, Mark Jerrum, and Haiko Müller. Polynomial-time approximation algorithms for the antiferromagnetic Ising model on line graphs. *Combinatorics, Probability and Computing*, pages 1–17, 2021.

[FGW22] Weiming Feng, Heng Guo, and Jiaheng Wang. Swendsen-Wang dynamics for the ferromagnetic Ising model with external fields. *arXiv preprint arXiv:2205.01985*, 2022.

[GGGHP22] Andreas Galanis, Leslie Ann Goldberg, Heng Guo, and Andrés Herrera-Poyatos. Fast sampling of satisfying assignments from random k-SAT. *arXiv preprint arXiv:2206.15308*, 2022.

[GJ18] Heng Guo and Mark Jerrum. Random cluster dynamics for the Ising model is rapidly mixing. *Annals of Applied Probability*, 28(2):1292–1313, 2018.

[GLLZ21] Heng Guo, Chao Liao, Pinyan Lu, and Chihao Zhang. Zeros of Holant problems: Locations and algorithms. *ACM Transactions on Algorithms*, 17(1):1–25, 2021.

[HLZ16] Lingxiao Huang, Pinyan Lu, and Chihao Zhang. Canonical paths for MCMC: From art to science. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on discrete algorithms (SODA)*, pages 514–527. SIAM, 2016.

[Jer03] Mark Jerrum. *Counting, Sampling and Integrating: Algorithms and Complexity.* Lectures in Mathematics, ETH Zürich. Birkhäuser Basel, 2003.

[JS89] Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM journal on computing*, 18(6):1149–1178, 1989.

[JS93] Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on Computing*, 22(5):1087–1116, 1993.

[JSV04] Mark Jerrum, Alistair Sinclair, and Eric Vigoda. A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries. *Journal of the ACM (JACM)*, 51(4):671–697, 2004.

[LLL14] Chengyu Lin, Jingcheng Liu, and Pinyan Lu. A simple FPTAS for counting edge covers. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 341–348, 2014.

[LLZ14] Jingcheng Liu, Pinyan Lu, and Chihao Zhang. FPTAS for counting weighted edge covers. In *Proceedings of the 22nd Annual European Symposium on Algorithms (ESA)*, pages 654–665, 2014.

[LR19] Jonathan D Leake and Nick R Ryder. Generalizations of the matching polynomial to the multivariate independence polynomial. *Algebraic Combinatorics*, 2(5):781–802, 2019.

[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.

[LWZ14] Pinyan Lu, Menghui Wang, and Chihao Zhang. FPTAS for weighted Fibonacci gates and its applications. In *Proceedings of the 41st International Colloquium on Automata, Languages and Programming (ICALP)*, pages 787–799. Springer, 2014.

[Mat08] James Matthews. *Markov Chains for Sampling Matchings*. PhD thesis, University of Edinburgh, 2008.

[McQ13] Colin McQuillan. Approximating Holant problems by winding. *ArXiv preprint, arXiv:1301.2880*, 2013.

[PR17] Viresh Patel and Guus Regts. Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials. *SIAM Journal on Computing*, 46(6):1893–1919, 2017.