

Spectral Independence via Stability and Applications to Holant-Type Problems

Received Sep 18, 2022

Revised Feb 4, 2024

Accepted Mar 15, 2024

Published Jul 12, 2024

Key words and phrases

Glauber dynamics, spectral independence, stability of polynomials, partition function, Holant problem

Zongchen Chen^a ✉

Kuikui Liu^b ✉

Eric Vigoda^c ✉

^a Georgia Institute of Technology, Atlanta, GA, USA

^b Massachusetts Institute of Technology, Cambridge, MA, USA

^c University of California, Santa Barbara, Santa Barbara, CA, USA

ABSTRACT. This paper formalizes connections between stability of polynomials and convergence rates of Markov Chain Monte Carlo (MCMC) algorithms. We prove that if a (multivariate) partition function is nonzero in a region around a real point λ then spectral independence holds at λ . As a consequence, for Holant-type problems (e.g., spin systems) on bounded-degree graphs, we obtain optimal $O(n \log n)$ mixing time bounds for the single-site update Markov chain known as the Glauber dynamics. Our result significantly improves the running time guarantees obtained via the polynomial interpolation method of Barvinok (2017), refined by Patel and Regts (2017).

There are a variety of applications of our results. In this paper, we focus on Holant-type (i.e., edge-coloring) problems, including weighted edge covers and weighted even subgraphs. For the weighted edge cover problem (and several natural generalizations) we obtain an $O(n \log n)$ sampling algorithm on bounded-degree graphs. The even subgraphs problem corresponds to the high-temperature expansion of the ferromagnetic Ising model. We obtain an $O(n \log n)$ sampling algorithm for the ferromagnetic Ising model with a nonzero external field on bounded-degree graphs, which improves upon the classical result of Jerrum and Sinclair (1993) for this class of graphs. We obtain further applications to antiferromagnetic two-spin models on line graphs, weighted graph homomorphisms, tensor networks, and more.

An extended abstract of this work appeared at the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS 2021) [26]. Kuikui Liu is supported by NSF grant CCF-1907845 and ONR-YIP grant N00014-17-1-2429. Eric Vigoda is supported by NSF grant CCF-2147094.

1. Introduction

A fundamental problem in a variety of settings, such as the study of spin systems in statistical physics or Bayesian inference in undirected graphical models, is the counting problem of estimating a partition function or the sampling problem of generating a random sample from the associated Gibbs distribution. The classical tool for approximate counting/sampling problems is the Markov Chain Monte Carlo (MCMC) method. There are now several alternative algorithmic approaches, most notably the correlation decay approach presented by Weitz [54] and the polynomial interpolation method presented by Barvinok [7] and refined by Patel and Regts [48].

Anari, Liu, and Oveis Gharan [3] presented a powerful new tool for analyzing MCMC methods known as spectral independence. Spectral independence yields optimal mixing time bounds for the Glauber dynamics (which updates a randomly chosen vertex in each step) [24], and more generally yields optimal mixing time bounds for any block dynamics and the Swendsen-Wang dynamics [17].

Spectral independence serves as an important technical bridge: for many results established with different approaches, their proofs also imply spectral independence and as a consequence we obtain much stronger running time guarantees. In particular, in [25] it was shown that potential function techniques used for analyzing correlation decay algorithms also imply spectral independence; see also [22, 30]. More recent works [17, 43] prove that path coupling arguments for a broad class of chains, including any local chain, yields spectral independence.

In this paper we prove, in an almost black-box fashion, that methods for establishing large zero-free regions needed for the polynomial interpolation method also yield spectral independence. The polynomial interpolation method is a mathematically elegant approach which works in the following manner. To approximate a partition function at a positive real value λ , one needs to prove there is a zero-free region around λ in the complex plane which means that the partition function has no roots in an open connected region (in the complex plane) containing λ . If such a zero-free region also contains an “easy” point (usually 0) at which the partition function and its derivatives can be efficiently evaluated, then this implies that one can approximate the Taylor series of a simple transformation of the partition function using a logarithmic number of terms, which yields a polynomial-time algorithm to approximate the partition function at λ .

We prove that a zero-free region implies spectral independence. This immediately yields several new rapid mixing results for MCMC methods. We also obtain significantly improved running times in many instances. For a spin system on a graph with n vertices and constant maximum degree Δ , the polynomial interpolation method [48] yields a running time of $O(n^C)$ where the constant C depends on Δ and parameters of the model. In contrast, spectral indepen-

dence implies an optimal mixing time bound of $O(n \log n)$ for the Glauber dynamics [24] (and more generally optimal mixing for the block dynamics [17]).

The study of zero-free regions has a long and rich history in the analysis of phase transitions in statistical physics models. The well-known work of Lee and Yang [39] utilizes zeros of the partition function to study phase transitions for the ferromagnetic Ising model; see Remark 1.10 for a discussion of the Lee-Yang Theorem.

Before stating our results we formally define the Glauber dynamics, also known as the Gibbs sampler. Let V be a finite set and let Q be a finite label/spin set. Consider a distribution μ over Q^V containing all labelings of elements in V . Note that this includes spin systems and Holant problems as special cases: for spin systems μ is the distribution over all spin assignments of vertices of the underlying graph, and for Holant problems μ is the distribution over all subsets of edges ($\{0, 1\}$ -labelings of edges). A transition of the Glauber dynamics chooses a random element $x \in V$, and then resamples the label of x from the marginal distribution at x conditional on the configuration outside x . The mixing time is the number of steps, from the worst initial state, which is guaranteed to be close (in total variation distance) to the desired distribution μ .

1.1 Applications

We state here three sample applications of our techniques; further applications are stated later in the paper.

For a graph $G = (V, E)$, we say a vertex v is covered by a subset $S \subseteq E$ of edges if v is incident to at least one edge in S . The subset $S \subseteq E$ is called an *edge cover* if all vertices are covered by S . Note there is always a trivial edge cover by setting $S = E$. An FPRAS (fully polynomial randomized approximation scheme) was presented for counting the number of edge covers for 3-regular graphs [16]. In [40] an FPTAS (deterministic analog of an FPRAS) for counting edge covers was presented for all graphs using the correlation decay approach, and the running time was $O(m^{1+\log_2 6} n^2)$ where m is the number of edges and n is the number of vertices. An FPRAS for all graphs using MCMC was presented in [35].

The correlation decay algorithm of [40] was extended to weighted (partial) edge covers (with worse running time guarantees) in [41]. In the weighted version, each edge has a weight $\lambda > 0$ and each vertex receives a penalty $\rho \in [0, 1]$ for being uncovered. Every subset $S \subseteq E$ is associated with the weight $w(S) = \rho^{|\text{unc}(S)|} \lambda^{|S|}$, where $\text{unc}(S)$ denotes the set of vertices that are not covered by S . The Gibbs distribution over all subsets of edges is given by $\mu(S) \propto w(S)$. Note, the case $\lambda = 1$ and $\rho = 0$ corresponds to uniformly random exact edge covers.

Finally, an FPTAS using the polynomial interpolation algorithm was presented for graphs with constant maximum degree [34], see also [15]. Using the zero-free results in [34] with our new technical contributions we immediately obtain an FPRAS using a simple MCMC algorithm and with significantly faster running time guarantees.

THEOREM 1.1 (Weighted Edge Covers). *Let $\Delta \geq 3$ be an integer and let $\lambda > 0$, $\rho \in [0, 1]$ be reals. Then for any n -vertex graph $G = (V, E)$ of maximum degree Δ , the Glauber dynamics for sampling random weighted edge covers of G with parameters λ, ρ mixes in $Cn \log n$ steps where $C = C(\Delta, \lambda, \rho)$ is a constant independent of n .*

One of the seminal results in the field of approximate counting is the work of Jerrum and Sinclair [37] presenting an FPRAS for the partition function of the ferromagnetic Ising model on any graph. The Ising model on a graph $G = (V, E)$ is described by two parameters, the edge activity $\beta_{\text{Ising}} > 0$ and the vertex activity $\lambda_{\text{Ising}} > 0$. The Gibbs distribution of the Ising model is over all $\{+, -\}$ spin assignments to vertices. Every configuration $\sigma : V \rightarrow \{+, -\}$ has density $\mu_{\text{Ising}}(\sigma) \propto \beta_{\text{Ising}}^{m(\sigma)} \lambda_{\text{Ising}}^{|\sigma^{-1}(+)|}$ where $m(\sigma)$ denotes the number of monochromatic edges in σ and $\sigma^{-1}(+)$ is the set of vertices assigned spin $+$. The model is ferromagnetic when $\beta_{\text{Ising}} > 1$, in which case neighboring vertices are more likely to have the same spin.

The central task of the Jerrum-Sinclair algorithm is sampling from the Gibbs distribution for the high-temperature expansion of the Ising model which is defined on all subsets of edges weighted to prefer subgraphs with more even degree vertices. For a graph $G = (V, E)$, an edge weight $\lambda > 0$, and a vertex penalty $\rho \in [0, 1]$, the Gibbs distribution μ for weighted (partial) even subgraphs is defined on all subsets of edges; a subset $S \subseteq E$ has weight $w(S) = \rho^{|\text{odd}(S)|} \lambda^{|S|}$ where $\text{odd}(S)$ is the set of odd-degree vertices in the subgraph (V, S) , and $\mu(S) \propto w(S)$. The weighted even subgraphs model where $\lambda \in (0, 1)$ is closely related to the ferromagnetic Ising model by $\beta_{\text{Ising}} = \frac{1+\lambda}{1-\lambda}$ and $\lambda_{\text{Ising}} = \frac{1+\rho}{1-\rho}$. Note that if $\rho = 0$ then μ is the distribution over all weighted exact even subgraphs, corresponding to the ferromagnetic Ising model without external fields (i.e., $\lambda_{\text{Ising}} = 1$). Using the Jerrum-Sinclair algorithm, Randall and Wilson [49] first gave an approximate sampler for the ferromagnetic Ising model via an equivalence with the random cluster model [31] and using self-reducibility. Grimmett and Janson discovered a direct coupling between weighted even subgraphs and the random cluster model [33, Theorem 3.5], which together with the coupling between random cluster and Ising models [28] yields a simpler and more efficient sampler for Ising configurations. In a recent work [29], a grand coupling among the ferromagnetic Ising model, random cluster model, and weighted even subgraphs is presented when external fields exist.

In [37], an MCMC algorithm is presented to sample weighted even subgraphs of an arbitrary (unbounded-degree) graph in time $O(m^3 \text{poly}(1/\rho))$ where m is the number of edges. In another direction, [42] presents an FPTAS for approximating the partition function of the ferromagnetic Ising model with nonzero fields on bounded-degree graphs, using Barvinok's polynomial interpolation method and the Lee-Yang theory. As is common for this type of approach, the running time of [42] is n^C for a constant C depending on the maximum degree of the graph and the parameters of the Ising model.

Here we use our results relating zero-free regions and spectral independence to obtain a faster MCMC algorithm for bounded-degree graphs when $\rho > 0$.

THEOREM 1.2 (Weighted Even Subgraphs). *Let $\Delta \geq 3$ be an integer and let $\lambda > 0$, $\rho \in (0, 1]$ be reals. Then for any n -vertex graph $G = (V, E)$ of maximum degree Δ , the Glauber dynamics for sampling random weighted even subgraphs of G with parameters λ, ρ mixes in $Cn \log n$ steps where $C = C(\Delta, \lambda, \rho)$ is a constant independent of n .*

In particular, this gives an approximate sampling algorithm with running time $O(n \log n)$ for the ferromagnetic Ising model with edge activity $\beta_{\text{Ising}} = \frac{1+\lambda}{1-\lambda}$ and vertex activity $\lambda_{\text{Ising}} = \frac{1+\rho}{1-\rho}$.

REMARK 1.3. In [37], the MCMC method can actually be used to obtain a sampler for $\rho = 0$ corresponding to weighted exact even subgraphs. This is achieved by taking $\rho = 1/n$ and using rejection sampling. Notice that the running time of [37] is polynomial in $1/\rho$, and therefore this gives a $\text{poly}(n)$ time algorithm for sampling weighted exact even subgraphs and hence for the ferromagnetic Ising model without fields. Unfortunately, Theorem 1.2 cannot be used to obtain a sampler for $\rho = 0$, since our bound on the mixing time of the Glauber dynamics (the constant C from Theorem 1.2) depends exponentially on $1/\rho$.

Finally, we simultaneously generalize [36, 27, 15] to all antiferromagnetic two-spin edge models, i.e., antiferromagnetic two-spin models on the class of line graphs. Again, in the bounded-degree regime we obtain optimal mixing times. Before we state the result, let us define the model more precisely. For a graph $G = (V, E)$ and fixed parameters $\beta \geq 0$, $\gamma > 0$, $\lambda > 0$, the Gibbs distribution of the corresponding two-spin edge model on G is given by

$$\mu(\sigma) \propto \beta^{m_1(\sigma)} \gamma^{m_0(\sigma)} \lambda^{|\sigma^{-1}(1)|}, \quad \forall \sigma \in \{0, 1\}^E \quad (1)$$

where $m_i(\sigma)$ denotes the number of pairs of edges e, f sharing a single endpoint such that $\sigma(e) = \sigma(f) = i$, for each $i = 0, 1$. We say the system is antiferromagnetic if $\beta\gamma < 1$ and ferromagnetic if $\beta\gamma > 1$ (note that $\beta\gamma = 1$ corresponds to a trivial product measure). The case $\beta = 0$ and $\gamma = 1$ recovers the monomer-dimer model for matchings weighted by λ , and the case $\beta = \gamma$ recovers the Ising model on the line graph of G .

THEOREM 1.4 (Antiferromagnetic Two-Spin Edge Models). *Let $\Delta \geq 3$ be an integer and let $\beta \geq 0$, $\gamma > 0$, $\lambda > 0$ be reals such that $\beta\gamma < 1$. Then for any n -vertex graph $G = (V, E)$ of maximum degree Δ , the Glauber dynamics for sampling from the antiferromagnetic two-spin edge model on G with parameters β, γ, λ mixes in $Cn \log n$ steps where $C = C(\Delta, \beta, \gamma, \lambda)$ is a constant independent of n .*

We present further applications of our methods in Sections 5 and 6.

1.2 Spectral Independence via Stability of the Partition Function

We need a few preliminary definitions before formally stating our technical results. Our results hold for an arbitrary distribution on a discrete product space; this general setup contains spin

systems as a special case. Let V be a finite set and we refer to the elements in V as vertices. For an integer $q \geq 2$, the set of spins is $Q = \{0\} \cup Q_1$ where $Q_1 = \{1, \dots, q-1\}$. The state space is $\Omega = Q^V$, the collection of all spin assignments of vertices. Finally, let $w : \Omega \rightarrow \mathbb{R}_{\geq 0}$ be a nonnegative weight function that is not always zero; i.e., $w(\sigma) > 0$ for at least one $\sigma \in \Omega$.

Let $\lambda : V \times Q_1 \rightarrow \mathbb{C}$ be a vector of (complex) external fields; each $\lambda_{v,k}$ represents the weight of vertex v receiving spin k . Without loss of generality, we assume spin 0 does not have associated external fields (this can be obtained by normalizing). Given w , the *partition function* is a multivariate polynomial of λ defined as:

$$Z_w(\lambda) = \sum_{\sigma \in \Omega} w(\sigma) \lambda^\sigma, \quad \text{where } \lambda^\sigma = \prod_{v \in V: \sigma_v \neq 0} \lambda_{v, \sigma_v}. \quad (2)$$

If λ is real and positive (i.e., every $\lambda_{v,k} \in \mathbb{R}_+$), then the *Gibbs distribution* $\mu = \mu_{w, \lambda}$ is given by:

$$\mu(\sigma) = \frac{w(\sigma) \lambda^\sigma}{Z_w(\lambda)}, \quad \forall \sigma \in \Omega. \quad (3)$$

Note that $Z_w(\lambda) > 0$ since w is not identically zero.

To establish spectral independence we need to consider the model with an arbitrary “pinning” which is a fixed configuration on an arbitrary subset of vertices. We formally define pinnings and the associated notions in Section 2.1, and introduce the relevant notation here.

A configuration $\sigma \in \Omega$ is said to be valid or feasible if $w(\sigma) > 0$. For $\Lambda \subseteq V$, let Ω_Λ denote the set of pinnings on Λ ; this is the set of configurations on Λ which have a valid extension to the remaining vertices $V \setminus \Lambda$. For $\Lambda \subseteq V$ and $\tau \in \Omega_\Lambda$, let $V^\tau = V \setminus \Lambda$ denote the set of unpinned vertices, let $Z_w^\tau(\lambda)$ be the multivariate conditional partition function under the pinning τ , and let μ^τ be the corresponding conditional distribution.

We can now define the notion of spectral independence. Let $\mathcal{T} = \bigcup_{\Lambda \subseteq V} \Omega_\Lambda$ be the collection of all pinnings. For $\tau \in \mathcal{T}$ let $\mathcal{P}^\tau = \{(v, k) \in V \times Q : v \in V^\tau, k \in \Omega_v^\tau\}$ be the collection of feasible vertex-spin pairs under τ , where Ω_v^τ represents the set of feasible spins at v conditioned on τ . The following definition is taken from [22]; see also [3, 30].

DEFINITION 1.5 (Influence Matrix). Let $\tau \in \mathcal{T}$ be an arbitrary pinning. For every $(u, j), (v, k) \in \mathcal{P}^\tau$, the (pairwise) influence of (u, j) on (v, k) under the pinning τ is given by $\Psi_\mu^\tau(u, j; v, k) = 0$ for $u = v$ and

$$\Psi_\mu^\tau(u, j; v, k) = \mu(\sigma_v = k \mid \sigma_u = j, \sigma_\Lambda = \tau) - \mu(\sigma_v = k \mid \sigma_\Lambda = \tau) \quad \text{for } u \neq v.$$

The (pairwise) influence matrix Ψ_μ^τ is a $|\mathcal{P}^\tau| \times |\mathcal{P}^\tau|$ matrix with entries given above.

All eigenvalues of the influence matrix Ψ_μ^τ are real since Ψ_μ^τ becomes symmetric after left-multiplication by a suitable diagonal matrix; see [3, 22, 30]. For a square matrix M with real eigenvalues, let $\text{EigMax}(M)$ denote the maximum eigenvalue of M .

DEFINITION 1.6 (Spectral Independence). We say μ is spectrally independent with constant η if for every pinning $\tau \in \mathcal{T}$ one has

$$\text{EigMax}(\Psi_\mu^\tau) \leq \eta.$$

For a non-empty region Γ of the complex plane, we say that a multivariate polynomial $P(z_1, \dots, z_n)$ is Γ -stable if $P(z_1, \dots, z_n) \neq 0$ whenever $z_j \in \Gamma$ for all j , see Definition 2.3. We present a sequence of results connecting spectral independence of the distribution with stability of the partition function. Our first result holds when the zero-free region of the partition function is sufficiently “large”, e.g., containing the whole positive real axis. Below for $\Gamma \subseteq \mathbb{C}$ let $\bar{\Gamma}$ denote the closure of Γ and let $\partial\Gamma$ be the boundary of Γ ; for $\lambda \in \mathbb{C}$ let $\text{dist}(\lambda, \partial\Gamma) = \inf_{z \in \partial\Gamma} |z - \lambda|$; see Section 2.4.

THEOREM 1.7. *Let $\Gamma \subseteq \mathbb{C}$ be a non-empty open connected region such that Γ is unbounded and $0 \in \bar{\Gamma}$. If the multivariate partition function Z_w is Γ -stable, then for any $\lambda \in \mathbb{R}_+ \cap \Gamma$ the Gibbs distribution $\mu = \mu_{w,\lambda}$ with the uniform external field λ is spectrally independent with constant*

$$\eta = \frac{8}{\delta}$$

where $\delta = \frac{1}{\lambda} \text{dist}(\lambda, \partial\Gamma)$.

In particular, the statement is true when Γ is a non-empty open connected region containing the positive real axis; i.e., $\mathbb{R}_+ \subseteq \Gamma$.

We can also obtain bounds when Γ is bounded or $0 \in \bar{\Gamma}$, under the assumption that Γ contains a part of the positive real axis. In this case we also need to further assume that all conditional partition functions under pinnings are stable. Moreover, our bound on spectral independence depends on the marginal bound of the distribution μ , which is defined as

$$b = \min_{\substack{\tau \in \mathcal{T} \\ (v,k) \in \mathcal{P}^\tau}} \mu^\tau(\sigma_v = k).$$

Note that $b > 0$ since $\mathcal{P}^\tau$ contains only feasible vertex-spin pairs.

THEOREM 1.8. *Let $\lambda^* \in \mathbb{R}_+$ and let $\Gamma \subseteq \mathbb{C}$ be a non-empty open connected region such that $(0, \lambda^*) \subseteq \Gamma$ (respectively, $(\lambda^*, \infty) \subseteq \Gamma$). If for every pinning $\tau \in \mathcal{T}$ the multivariate conditional partition function Z_w^τ is Γ -stable, then for any $\lambda \in (0, \lambda^*)$ (respectively, $\lambda \in (\lambda^*, \infty)$) the Gibbs distribution $\mu = \mu_{w,\lambda}$ with the uniform external field λ is spectrally independent with constant*

$$\eta = \frac{8}{\delta} \min \left\{ \frac{1-b}{b}, \frac{\lambda}{b(\lambda^* - \lambda)} + 1 \right\}$$

$$\left(\text{respectively, } \eta = \frac{8}{\delta} \min \left\{ \frac{1-b}{b}, \frac{\lambda^*}{b(\lambda - \lambda^*)} + 1 \right\} \right)$$

where b is the marginal bound for μ and $\delta = \frac{1}{\lambda} \text{dist}(\lambda, \partial\Gamma)$.

REMARK 1.9. The first term $(1 - b)/b$ is better when λ is close to λ^* , while the second term is better when λ is close to 0 (respectively, ∞), because usually b/λ is bounded from below when $\lambda \rightarrow 0$ (respectively, $b\lambda$ is bounded from below when $\lambda \rightarrow \infty$).

REMARK 1.10. We point out here that Theorem 1.8 does not apply to the ferromagnetic Ising model. The celebrated Lee-Yang theorem states that the partition function for the ferromagnetic Ising model is $\mathbb{D}(0, 1)$ -stable and $\overline{\mathbb{D}}(0, 1)^c$ -stable where $\mathbb{D}(0, 1)$ denotes the open unit ball centered at 0 on the complex plane and $\overline{\mathbb{D}}(0, 1)^c$ denotes the exterior of $\mathbb{D}(0, 1)$. However, when a pinning is applied, particularly when some vertices are pinned to + and some are −, we do not have either $\mathbb{D}(0, 1)$ -stability or $\overline{\mathbb{D}}(0, 1)^c$ -stability for the conditional partition function. To see this, notice that such a pinning can result in inconsistent external fields; some fields are < 1 (hence in $\mathbb{D}(0, 1)$) while others are > 1 (hence in $\overline{\mathbb{D}}(0, 1)^c$), and the Lee-Yang theorem does not apply.

Meanwhile, one should not expect spectral independence to hold for the ferromagnetic Ising model at all temperatures and for all external fields, since the Glauber dynamics is slow mixing when the parameters lie in the tree non-uniqueness region (see, for instance, [32]).

If limited information about the zero-free region is given, then spectral independence still holds with a worse bound.

THEOREM 1.11. *Let $\Gamma \subseteq \mathbb{C}$ be a non-empty open connected region. If for every pinning $\tau \in \mathcal{T}$ the multivariate conditional partition function Z_w^τ is Γ -stable, then for any $\lambda \in \mathbb{R}_+ \cap \Gamma$ the Gibbs distribution $\mu = \mu_{w,\lambda}$ with the uniform external field λ is spectrally independent with constant*

$$\eta = \frac{2}{b\delta^2}$$

where b is the marginal bound for μ and $\delta = \frac{1}{\lambda} \text{dist}(\lambda, \partial\Gamma)$.

It is unclear if the marginal bound b is needed or not. Our results Theorems 1.7, 1.8 and 1.11 also hold for non-uniform external fields, i.e., each pair (v, k) is assigned a distinct field $\lambda_{v,k}$, and the zero-free regions are allowed to be distinct for different pairs. See Theorem 3.2 for a formal statement which implies Theorems 1.7, 1.8 and 1.11 as special cases. Our proof of Theorem 3.2 uses tools from Complex Analysis, such as the Riemann Mapping Theorem (see Theorem 2.6). We remark that for the simple but slightly worse bound Theorem 1.11 the proof does not rely on the Riemann Mapping Theorem.

1.3 Relations with Previous Works

Our work builds upon the recent work of Alimohammadi, Anari, Shiragur, and Vuong [2]. Theorem 16 of [2] established spectral independence for any distribution over $\{0, 1\}^V$ assuming that the generating polynomial is sector-stable (that is, Γ -stable where $\Gamma = \{z \in \mathbb{C} : |\text{Arg}(z)| \leq \theta\}$ is a sector for some $\theta \in (0, \pi/2)$). Our results Theorems 1.7, 1.8 and 1.11 strengthen theirs in

the sense that we do not have any restriction on the zero-free region Γ , and the results hold for any open connected region. This allows us to apply our results in a much broader setting. See Section 3 for more details.

To establish zero-free regions for our main applications, we utilize the approach in [34], which reduces the problem via Asano-Ruelle contractions [5, 52] to showing a sufficiently large zero-free region for a collection of bounded-degree univariate polynomials, one for each vertex of the input graph. These univariate polynomials are referred to as the local polynomials, since they only depend on the configuration restricted to edges incident to the given vertex. We note a very similar idea was also used in [53, 15] to establish zero-free regions, although their methods do not go through Asano-Ruelle contractions. See Section 4 for more details.

It was also shown in a sequence of papers [10, 14, 13, 6, 9, 50] that one can establish large zero-free regions via an inductive approach based on conditioning the distribution. This method of establishing zero-free regions also works nicely for us, as spectral independence requires a bound on the pairwise influences for all conditional distributions. We show that one can deduce rapid mixing of the Glauber dynamics in a nearly black-box fashion from these zero-free methods for several problems in Sections 5 and 6.

Algorithmically, our results have several advantages over prior works utilizing zero-free regions. In particular, the polynomial interpolation method pioneered by Barvinok [7] typically only yields quasi-polynomial time algorithms in general, and polynomial time algorithms with exponent depending on the maximum degree for problems arising from graphs [48]. In contrast, we obtain fast algorithms for sampling and counting. Another feature of our approach is that we only need the zero-free region to be sufficiently large. This is in contrast to the polynomial interpolation technique, which needs the zero-free region to also contain a point at which the partition function is easily computable. On the other hand, our approach is fundamentally restricted to nonnegative real parameters whereas Barvinok's approach can be extended to complex parameters, see, for instance [7].

Outline of Paper In Section 3 we prove our general technical results Theorems 1.7, 1.8 and 1.11 connecting zero-free regions with spectral independence. We prove Theorems 1.1, 1.2 and 1.4 in Section 4 regarding binary symmetric Holant problems. In Section 5 we prove results for weighted graph homomorphisms and tensor networks. Finally, in Section 6 we study arbitrary measures on the discrete cube as studied in the analysis of Boolean functions.

2. Preliminaries

2.1 Pinnings

Let $q \geq 2$ be an integer and write $q_1 = q - 1$. Let V be a finite set of vertices and let $Q = \{0\} \cup Q_1$ be the set of spins where $Q_1 = \{1, \dots, q_1\}$. Every spin assignment $\sigma : V \rightarrow Q$ is called a

configuration. The state space $\Omega = Q^V$ is the collection of all configurations and let $w : \Omega \rightarrow \mathbb{R}_{\geq 0}$ be a nonnegative weight function that is not identically zero. A configuration $\sigma \in \Omega$ is said to be *valid* or *feasible* if $w(\sigma) > 0$. For $\Lambda \subseteq V$, define set of *pinning*s on Λ by

$$\Omega_\Lambda = \{\tau \in Q^\Lambda : \exists \text{ valid } \sigma \in \Omega \text{ s.t. } \sigma_\Lambda = \tau\}.$$

Note Ω_V is the set of all valid configurations. Let $\mathcal{T} = \bigcup_{\Lambda \subseteq V} \Omega_\Lambda$ be the collection of all pinning.

For a pinning $\tau \in \mathcal{T}$, let V^τ denote the set of unpinning vertices; so if $\tau \in \Omega_\Lambda$ then $V^\tau = V \setminus \Lambda$. For $v \in V^\tau$, let Ω_v^τ be the set of valid spins at v under τ :

$$\Omega_v^\tau = \{k \in Q : \exists \text{ valid } \sigma \in \Omega \text{ s.t. } \sigma_\Lambda = \tau \text{ and } \sigma_v = k\}.$$

Define the collection of feasible vertex-spin pairs under τ by $\mathcal{P}^\tau = \{(v, k) \in V \times Q : v \in V^\tau, k \in \Omega_v^\tau\}$, and the collection of pairs with nonzero spins by $\mathcal{P}_1^\tau = \{(v, k) \in \mathcal{P}^\tau : k \neq 0\}$. We write Ω_v , $\mathcal{P}$, and $\mathcal{P}_1$ when no pinning is applied.

Let $\lambda : \mathcal{P}_1 \rightarrow \mathbb{C}$ be a vector of complex external fields. Given a pinning $\tau \in \mathcal{T}$ on $\Lambda \subseteq V$ with $U = V^\tau = V \setminus \Lambda$, the *conditional partition function* Z_w^τ under τ is a multivariate polynomial of λ defined as

$$Z_w^\tau(\lambda) = \sum_{\sigma \in \Omega : \sigma_\Lambda = \tau} w(\sigma) \lambda^{\sigma_U}, \quad \text{where } \lambda^{\sigma_U} = \prod_{v \in U : \sigma_v \neq 0} \lambda_{v, \sigma_v}.$$

When there is no pinning, this matches Equation (2) from the introduction. Observe that Z_w^τ depends only on the variables $\{\lambda_{v,k} : (v, k) \in \mathcal{P}_1^\tau\}$, and that Z_w^τ is not identically zero since τ is a pinning. If λ is real and positive, then $Z_w^\tau(\lambda) > 0$ and we obtain the *conditional Gibbs distribution*:

$$\mu^\tau(\sigma) = \mu(\sigma \mid \sigma_\Lambda = \tau) = \frac{w(\sigma) \lambda^{\sigma_U}}{Z_w^\tau(\lambda)}, \quad \forall \sigma \in \Omega \text{ s.t. } \sigma_\Lambda = \tau.$$

Again this matches Equation (3) when there is no pinning.

2.2 Spectral Independence

Let P be the transition matrix of the Glauber dynamics with stationary distribution μ , and let $P^t(X_0, \cdot)$ denote the distribution of the chain after t steps starting from $X_0 \in \Omega_V$. The *mixing time* of the Glauber dynamics is defined as

$$T_{\text{mix}}(P) = \max_{X_0 \in \Omega_V} \min\{t \geq 0 : \|P^t(X_0, \cdot) - \mu\|_{\text{TV}} \leq 1/4\},$$

where $\|\cdot\|_{\text{TV}}$ denotes the total variation distance. Throughout this paper, we assume that the Glauber dynamics is ergodic for every conditional distribution μ^τ where $\tau \in \mathcal{T}$ is a pinning (such distribution μ is called totally-connected in [24, 17]).

The notion of *spectral independence* (see Definitions 1.5 and 1.6) was introduced in [3] and immediately becomes a powerful tool for establishing rapid mixing of the Glauber dynamics.

THEOREM 2.1 ([1, 3]). *Let μ be an arbitrary distribution over Q^V where $|V| = n$. If μ is spectrally independent with constant η , then the Glauber dynamics mixes in $O(n^{\eta+1} \log(1/\mu_{\min}))$ steps where $\mu_{\min} = \min_{\sigma \in \Omega_V} \mu(\sigma)$.*

The following mixing result for the Glauber dynamics is known for spin systems with nearest-neighbor interactions; it holds for Holant problems and tensor network contractions (see Sections 4 and 5 for definitions) as well since one can view these as spin systems defined on hypergraphs (also known as Markov random fields) and the proof approach of [24, 17] still works when the underlying hypergraph has bounded maximum degree.

THEOREM 2.2 ([24, 17]). *Let μ be the Gibbs distribution of a spin system or a Holant problem or a tensor network contraction defined on an n -vertex graph of maximum degree Δ . If μ is spectrally independent with constant η and the marginal bound for μ is b , then the Glauber dynamics mixes in $Cn \log n$ steps where $C = C(\Delta, \eta, b)$ is a constant independent of n .*

2.3 Stability

For n sets $\Gamma_1, \dots, \Gamma_n$, let $\prod_{\ell=1}^n \Gamma_\ell = \Gamma_1 \times \dots \times \Gamma_n$ denote the Cartesian product of them.

DEFINITION 2.3 (Stability). For an integer $n \geq 1$ and $\mathcal{K} \subseteq \mathbb{C}^n$, we say a multivariate polynomial $P \in \mathbb{C}[z_1, \dots, z_n]$ is $\mathcal{K}$ -stable if $P(z_1, \dots, z_n) \neq 0$ whenever $(z_1, \dots, z_n) \in \mathcal{K}$. In particular, if $\mathcal{K} = \prod_{\ell=1}^n \Gamma$ for some $\Gamma \subseteq \mathbb{C}$, then we simply say P is Γ -stable.

THEOREM 2.4 (Hurwitz' Theorem). *Let $n \geq 1$ be an integer and $\mathcal{K} \subseteq \mathbb{C}^n$ be an open connected set. Suppose that $\{f_m\}_{m=1}^\infty$ is a sequence of non-vanishing analytic functions on $\mathcal{K}$ that converges to f uniformly on compact subsets of $\mathcal{K}$. Then f is either non-vanishing on $\mathcal{K}$ or else identically zero.*

2.4 Complex Plane

We refer to subsets of the complex plane as *regions*. Note that this is slightly nonstandard in Complex Analysis, where a region (or domain) is more commonly defined as a non-empty open connected subset of $\mathbb{C}$.

Let $\Gamma \subseteq \mathbb{C}$ be a region. Denote the *complement* of Γ by $\Gamma^c = \mathbb{C} \setminus \Gamma$, its *interior* by Γ° , its *closure* by $\bar{\Gamma}$, and its *boundary* by $\partial\Gamma$. We say Γ is *unbounded* if for any $M \in \mathbb{R}_+$ there exists $z \in \Gamma$ such that $|z| > M$; otherwise it is called *bounded*. For $z \in \mathbb{C}$ let $\text{dist}(z, \Gamma) = \inf_{w \in \Gamma} |w - z|$ be the distance from z to Γ on the complex plane.

For a region $\Gamma \subseteq \mathbb{C}$ and $z \in \mathbb{C}$, we define $\Gamma + z = \{w + z : w \in \Gamma\}$, $z\Gamma = \{zw : w \in \Gamma\}$, and $\Gamma^{-1} = (\Gamma \setminus \{0\})^{-1} = \{w^{-1} : w \in \Gamma \setminus \{0\}\}$. For $\Gamma_1, \Gamma_2 \subseteq \mathbb{C}$, let $\Gamma_1 \cdot \Gamma_2 = \{zw : z \in \Gamma_1, w \in \Gamma_2\}$ denote their Minkowski product; in particular, for $\Gamma \subseteq \mathbb{C}$ let $\Gamma^2 = \Gamma \cdot \Gamma = \{zw : z, w \in \Gamma\}$ (meanwhile we shall write $\prod_{\ell=1}^2 \Gamma = \Gamma \times \Gamma = \{(z, w) : z, w \in \Gamma\}$ for the Cartesian product).

For $z \in \mathbb{C}$ and $r \in \mathbb{R}_+$, let $\mathbb{D}(z, r) = \{w \in \mathbb{C} : |w - z| < r\}$ denote the *open disk* centered at z with radius r , and let $\overline{\mathbb{D}}(z, r) = \{w \in \mathbb{C} : |w - z| \leq r\}$ denote the *closed disk*. An (open or closed) *polydisk* is a Cartesian product of (open or closed) disks. For $\varepsilon \in \mathbb{R}_+$, let $\mathbb{H}_\varepsilon = \{x + iy : x < -\varepsilon\}$ and $\overline{\mathbb{H}}_\varepsilon = \{x + iy : x \leq -\varepsilon\}$ be open and closed *half-planes*.

Let $\Gamma \subseteq \mathbb{C}$ be a non-empty open region. We say $w, z \in \Gamma$ are *(path-)connected* in Γ if there exists a continuous map $\gamma : [0, 1] \rightarrow \Gamma$ such that $\gamma(0) = w$ and $\gamma(1) = z$. Observe that connectivity in Γ is an equivalence relation, and we call each equivalence class a *(path-)connected component* of Γ . The region Γ is said to be *(path-)connected* if every two points from Γ are connected in Γ ; namely, Γ has a unique connected component which is itself. If Γ is open then every connected component of Γ is also open.

A non-empty open connected region $\Gamma \subseteq \mathbb{C}$ is called *simply connected* if its complement in the Riemann sphere $(\mathbb{C} \cup \infty)$ is also connected. A *Jordan curve* (simple closed curve) is a continuous map $\gamma : [0, 1] \rightarrow \mathbb{C}$ such that $\gamma(0) = \gamma(1)$ and the restriction of γ to $[0, 1)$ is injective. The Jordan curve theorem states that for a Jordan curve γ , the complement of its image on the complex plane consists of exactly two open connected components; one of these components is bounded called the *interior* and the other is unbounded called the *exterior*. A non-empty open connected region $\Gamma \subseteq \mathbb{C}$ is simply connected if and only if for every Jordan curve γ whose image is contained in Γ , the interior of γ is also contained in Γ .

2.5 Complex Analysis

We present here a few useful theorems in Complex Analysis which can be found in many textbooks, see e.g. [51]. Throughout, we select the principal branch for the complex functions $z \mapsto \log z$ and $z \mapsto z^{1/d}$.

THEOREM 2.5 (Schwarz-Pick Theorem). *Let $f : \mathbb{D}(0, 1) \rightarrow \mathbb{D}(0, 1)$ be a holomorphic function. Then*

$$|f'(0)| \leq 1 - |f(0)|^2 \leq 1.$$

For open regions $\Gamma_1, \Gamma_2 \subseteq \mathbb{C}$, a function $f : \Gamma_1 \rightarrow \Gamma_2$ is said to be *biholomorphic* if f is a bijective holomorphic function whose inverse is also holomorphic.

THEOREM 2.6 (Riemann Mapping Theorem). *Let $\Gamma \subseteq \mathbb{C}$ be a non-empty open simply connected region that is not $\mathbb{C}$. Then for any $z \in \Gamma$ there exists a unique biholomorphic mapping $f : \Gamma \rightarrow \mathbb{D}(0, 1)$ such that*

$$f(z) = 0 \quad \text{and} \quad f'(z) \in \mathbb{R}_+.$$

THEOREM 2.7 (Koebe's One-Quarter Theorem). *Let $\Gamma \subseteq \mathbb{C}$ and let $f : \mathbb{D}(0, 1) \rightarrow \Gamma$ be an injective holomorphic function. Then*

$$\mathbb{D}\left(f(0), \frac{1}{4}|f'(0)|\right) \subseteq \Gamma.$$

THEOREM 2.8 (Multivariate Open Mapping Theorem, [38, Theorem 1.8.1]). *Let $n \geq 1$ be an integer and let $\mathcal{K} \subseteq \mathbb{C}^n$ be a non-empty open connected subset of $\mathbb{C}^n$. Let $f : \mathcal{K} \rightarrow \mathbb{C}$ be a non-constant holomorphic function. Then the image of f is an open connected region.*

3. Establishing Spectral Independence via Stability

In this section, we deduce spectral independence of a distribution from the stability of the associated partition function, and thus prove Theorems 1.7, 1.8 and 1.11.

We first show that pinning preserves stability of the partition function if the zero-free region is unbounded and contains 0 in its closure. Intuitively, the pinning $\sigma_v = k$ for $k \neq 0$ corresponds to taking $\lambda_{v,k} = \infty$ (which is achieved by taking derivative with respect to $\lambda_{v,k}$), and the pinning $\sigma_v = 0$ corresponds to taking $\lambda_{v,k} = 0$ for all $k \in Q_1$. Hence, under an arbitrary pinning the conditional partition function is just the original partition function specialized at ∞ and 0 for specific external fields, and the fact that the closure of the zero-free region contains ∞ (i.e., unboundedness) and 0 guarantees that after specialization the resulted partition function is still stable. This is formalized by the following lemma.

LEMMA 3.1. *Let $\{\Gamma_{v,k} \subseteq \mathbb{C} : (v,k) \in \mathcal{P}_1\}$ be a collection of non-empty open connected regions such that for every $(v,k) \in \mathcal{P}_1$ the region $\Gamma_{v,k}$ is unbounded and $0 \in \overline{\Gamma_{v,k}}$. If the multivariate partition function Z_w is $(\prod_{(v,k) \in \mathcal{P}_1} \Gamma_{v,k})$ -stable, then for every pinning $\tau \in \mathcal{T}$ the multivariate conditional partition function Z_w^τ is $(\prod_{(v,k) \in \mathcal{P}_1^\tau} \Gamma_{v,k})$ -stable.*

Next, we show the following theorem for deriving spectral independence under various assumptions in the multivariate setting, from which one can deduce the bounds on spectral independence given in Theorems 1.7, 1.8 and 1.11.

THEOREM 3.2. *Let $\{\Gamma_{v,k} \subseteq \mathbb{C} : (v,k) \in \mathcal{P}_1\}$ be a collection of non-empty open connected regions, and let $\lambda : \mathcal{P}_1 \rightarrow \mathbb{R}_+$ such that $\lambda_{v,k} \in \mathbb{R}_+ \cap \Gamma_{v,k}$ for each $(v,k) \in \mathcal{P}_1$. Suppose that for every pinning $\tau \in \mathcal{T}$ the multivariate conditional partition function Z_w^τ is $(\prod_{(v,k) \in \mathcal{P}_1^\tau} \Gamma_{v,k})$ -stable. Then the Gibbs distribution $\mu = \mu_{w,\lambda}$ with external fields λ is spectrally independent with constant*

$$\eta = \frac{2}{b\delta^2}, \quad (4)$$

where b is the marginal bound for μ and

$$\delta = \min_{(v,k) \in \mathcal{P}_1} \frac{1}{\lambda_{v,k}} \text{dist}(\lambda_{v,k}, \partial\Gamma_{v,k}).$$

Furthermore:

1. For each $v \in V$ let $\Gamma_v \subseteq \mathbb{C}$ be the connected component of the intersection $\bigcap_{k \in Q_1} \frac{1}{\lambda_{v,k}} \Gamma_{v,k}$ that contains 1 (note that $1 \in \frac{1}{\lambda_{v,k}} \Gamma_{v,k}$ for all (v,k)). If for every $v \in V$ the region Γ_v is unbounded

and $0 \in \overline{\Gamma_v}$, then spectral independence holds with constant

$$\eta = \frac{8}{\delta}.$$

In particular, the statement is true if $\mathbb{R}_+ \subseteq \Gamma_{v,k}$ for each $(v, k) \in \mathcal{P}_1$.

2. If there exists $\lambda^* \in \mathbb{R}_+$ such that $\lambda_{v,k} \in (0, \lambda^*) \subseteq \Gamma_{v,k}$ for every $(v, k) \in \mathcal{P}_1$, then spectral independence holds with constant

$$\eta = \frac{8}{\delta} \min \left\{ \frac{1-b}{b}, \frac{\lambda_{\max}}{b(\lambda^* - \lambda_{\max})} + 1 \right\},$$

where $\lambda_{\max} = \max_{(v,k) \in \mathcal{P}_1} \lambda_{v,k}$.

3. If there exists $\lambda^* \in \mathbb{R}_+$ such that $\lambda_{v,k} \in (\lambda^*, \infty) \subseteq \Gamma_{v,k}$ for every $(v, k) \in \mathcal{P}_1$, then spectral independence holds with constant

$$\eta = \frac{8}{\delta} \min \left\{ \frac{1-b}{b}, \frac{\lambda^*}{b(\lambda_{\min} - \lambda^*)} + 1 \right\},$$

where $\lambda_{\min} = \min_{(v,k) \in \mathcal{P}_1} \lambda_{v,k}$.

We prove Theorem 3.2 by upper bounding the absolute row sum of the influence matrix Ψ_μ^τ for any pinning $\tau \in \mathcal{T}$; namely, for each $(u, j) \in \mathcal{P}^\tau$ we bound the sum of absolute values of the influences from (u, j) to all other pairs $(v, k) \in \mathcal{P}^\tau$, see Lemma 3.4. We accomplish this by strengthening and generalizing the proof strategy in [2].

At a high level, the work [2] views the sum of absolute influences as the derivative of some function f produced by the conditional partition functions. The variables of f are just the external fields of the partition function which lie in some zero-free region Γ and the stability of the conditional partition functions guarantees that the image of f is contained in some nice region Γ' . In [2], the authors study sector-stability of the partition function for the binary state space $\{0, 1\}^V$; in particular, both the zero-free region Γ and the region Γ' containing the image are sectors for their choice of f . Then, by applying conformal mappings between the sector and the unit disk, the derivative of f can be upper bounded using the Schwarz-Pick Theorem (Theorem 2.5).

However, here we are facing a more challenging situation since we are trying to establish spectral independence from an arbitrary zero-free region Γ for any discrete product space Q^V . In fact, for us the regions Γ and Γ' are in abstract form and to apply the Schwarz-Pick Theorem we need to design good mappings from Γ and Γ' to the unit disk. This is achieved by utilizing tools from Complex Analysis, especially the Riemann Mapping Theorem (Theorem 2.6). See Sections 3.3.1 and 3.3.2 for details of this part.

We now provide the proofs of Theorems 1.7, 1.8 and 1.11 from the introduction.

PROOF OF THEOREM 1.7. Follows from Lemma 3.1 and Item 1 of Theorem 3.2. ■

PROOF OF THEOREM 1.8. Follows from Items 2 and 3 of Theorem 3.2. ■

PROOF OF THEOREM 1.11. Follows from Equation (4) of Theorem 3.2. ■

After proving Lemma 3.1 in Section 3.1, we establish Theorem 3.2 in Section 3.2 and prove Lemma 3.4, a central lemma for bounding the absolute sum of influences, in Section 3.3.

3.1 Preservation of Stability under Pinnings

In this subsection we present the proof of Lemma 3.1.

Let $\tau \in \mathcal{T}$ be an arbitrary pinning on $\Lambda \subseteq V$ and let $U = V^\tau = V \setminus \Lambda$ be the set of unpinned vertices. We consider the conditional partition function Z_w^τ under the pinning τ . As discussed earlier, one can view Z_w^τ as obtained from the original partition function Z_w by specializing at 0 and taking derivatives for certain variables provided by the pinning τ . To be more precise, we define $\Lambda_0 = \{v \in \Lambda : \tau_v = 0\}$ to be the set of vertices pinned to spin 0, and let $\Lambda_1 = \{v \in \Lambda : \tau_v \neq 0\}$ be those pinned to nonzero spins. We also define $\tau_1 = \tau_{\Lambda_1}$ to be the pinning restricted to vertices with nonzero spins. The key observation here is that

$$Z_w^\tau(\lambda) = \left(\frac{\partial}{\partial \lambda_{\Lambda_1, \tau_1}} Z_w(\lambda) \right) \Big|_{\lambda_{\Lambda_0} = 0} \quad (5)$$

where $\lambda_{\Lambda_0} = 0$ represents plugging in $\lambda_{v,k} = 0$ for all $v \in \Lambda_0$ and $k \in Q_1$, and $\frac{\partial}{\partial \lambda_{\Lambda_1, \tau_1}}$ represents taking derivatives $\frac{\partial}{\partial \lambda_{v, \tau_v}}$ for all $v \in \Lambda_1$. Hence, to establish Lemma 3.1 it suffices to show that specialization at 0 and differentiation preserves Γ -stability if the zero-free region Γ is unbounded and $0 \in \bar{\Gamma}$. This is actually true for any multi-affine polynomial, which is a polynomial whose monomials are all square-free.

LEMMA 3.3. *Let $n \geq 1$ be an integer and let $\Gamma_1, \dots, \Gamma_n \subseteq \mathbb{C}$ be non-empty open connected regions. Let $P \in \mathbb{C}[z_1, \dots, z_n]$ be a multi-affine polynomial and assume that P is $(\prod_{\ell=1}^n \Gamma_\ell)$ -stable. Then:*

1. (Inversion) *The polynomial $P_1(z_1, z_2, \dots, z_n) = z_1 P(\frac{1}{z_1}, z_2, \dots, z_n)$ is $(\Gamma_1^{-1} \times \prod_{\ell=2}^n \Gamma_\ell)$ -stable;*
2. (Specialization) *If $0 \in \bar{\Gamma}_1$, then the polynomial $P_2(z_2, \dots, z_n) = P(0, z_2, \dots, z_n)$ is either $(\prod_{\ell=2}^n \Gamma_\ell)$ -stable or identically zero;*
3. (Differentiation) *If Γ_1 is unbounded, then the polynomial $P_3(z_2, \dots, z_n) = \frac{\partial}{\partial z_1} P(z_1, z_2, \dots, z_n)$ is either $(\prod_{\ell=2}^n \Gamma_\ell)$ -stable or identically zero.*

PROOF. Consider first the inversion property. Suppose for sake of contradiction that P_1 is not $(\Gamma_1^{-1} \times \prod_{\ell=2}^n \Gamma_\ell)$ -stable. Then there exists $w_1 \in \Gamma_1^{-1}$ and $z_\ell \in \Gamma_\ell$ for $2 \leq \ell \leq n$ such that $P_1(w_1, z_2, \dots, z_n) = 0$. Note that $w_1 = 1/z_1$ for some $z_1 \in \Gamma_1 \setminus \{0\}$. It follows that

$$0 = z_1 P_1(w_1, z_2, \dots, z_n) = z_1 w_1 P\left(\frac{1}{w_1}, z_2, \dots, z_n\right) = P(z_1, z_2, \dots, z_n),$$

contradicting to the stability of P . Hence, we have the desired stability for P_1 .

Next consider specialization. Since Γ_1 is open and $0 \in \bar{\Gamma}_1$, there exists a sequence of complex numbers $\{\zeta_m\}_{m=1}^\infty$ such that $\zeta_m \in \Gamma_1$ and $\lim_{m \rightarrow \infty} \zeta_m = 0$. Let $f_m(z_2, \dots, z_n) = P(\zeta_m, z_2, \dots, z_n)$

be a polynomial of degree $\leq \deg(P)$ for each m . Then f_m is $(\prod_{\ell=2}^n \Gamma_\ell)$ -stable by the stability assumption of P . Furthermore, the sequence $\{f_m\}_{m=1}^\infty$ converges to P_2 coefficient-wise, and hence uniformly on compact subsets; see, e.g., Lemma 33 in [2]. Hurwitz' Theorem (Theorem 2.4) then implies that P_2 is either $(\prod_{\ell=2}^n \Gamma_\ell)$ -stable or identically zero, as claimed.

Last we consider differentiation. Since Γ_1 is open and unbounded, we deduce that the region $\Gamma_1^{-1} = \{1/z : z \in \Gamma_1 \setminus \{0\}\}$ is open and satisfies $0 \in \overline{\Gamma_1^{-1}}$. Recall that we have shown the inversion $P_1(z_1, z_2, \dots, z_n) = z_1 P(\frac{1}{z_1}, z_2, \dots, z_n)$ is $(\Gamma_1^{-1} \times \prod_{\ell=2}^n \Gamma_\ell)$ -stable. Now observe that, for a multi-affine polynomial P , the derivative P_3 of P with respect to z_1 is the same as the specialization of P_1 at $z_1 = 0$:

$$P_3(z_2, \dots, z_n) = \frac{\partial}{\partial z_1} P(z_1, z_2, \dots, z_n) = P_1(0, z_2, \dots, z_n).$$

Hence, we immediately conclude from previous results that P_3 is either $(\prod_{\ell=2}^n \Gamma_\ell)$ -stable or identically zero. ■

Lemma 3.1 is an immediate consequence of Lemma 3.3.

PROOF OF LEMMA 3.1. Observe that the partition function Equation (2) is multi-affine. The lemma then follows from Equation (5) and Lemma 3.3. Notice that the conditional partition functions are never identically zero since pinnings are extendable to valid full configurations. ■

3.2 Bounds on Spectral Independence

In this subsection we prove Theorem 3.2.

An important observation is that it is sufficient to assume $1 \in \Gamma_{v,k}$ for every $(v, k) \in \mathcal{P}_1$ and consider the Gibbs distribution with the all-one external fields $\mathbf{1}$. In general, given the external field $\lambda_{v,k} \in \Gamma_{v,k}$ for each (v, k) , we may reweight the configurations by

$$\tilde{w}(\sigma) = w(\sigma) \prod_{v \in V: \sigma_v \neq 0} \lambda_{v, \sigma_v}, \quad \forall \sigma \in \Omega \quad (6)$$

and the new partition function is

$$Z_{\tilde{w}}(\lambda) = \sum_{\sigma \in \Omega} \tilde{w}(\sigma) \lambda^\sigma. \quad (7)$$

In particular, $Z_w(\lambda) = Z_{\tilde{w}}(\mathbf{1})$ and $\mu_{w, \lambda} = \mu_{\tilde{w}, \mathbf{1}}$ for the given $\lambda = (\lambda_{v,k})$. In other words, we hide the external fields into the weight of configurations and under the new weights we are interested in the all-one external fields. This will simplify the notations. Note that, if for $\tau \in \mathcal{T}$ the multivariate conditional partition function Z_w^τ is $(\prod_{(v,k) \in \mathcal{P}_1^\tau} \Gamma_{v,k})$ -stable, then the reweighted conditional partition function $Z_{\tilde{w}}^\tau$ is $(\prod_{(v,k) \in \mathcal{P}_1^\tau} \tilde{\Gamma}_{v,k})$ -stable where $\tilde{\Gamma}_{v,k} = \frac{1}{\lambda_{v,k}} \Gamma_{v,k}$ for each $(v, k) \in \mathcal{P}_1$.

In the rest of this section, we assume that $1 \in \Gamma_{v,k}$ and consider the case of all-one external fields. The following lemma is an important step towards deducing Theorem 3.2; it builds upon the proof strategy of [2] while generalizing their result.

LEMMA 3.4. *Consider the Gibbs distribution $\mu = \mu_{w,1}$ with the all-one external fields $\mathbf{1}$. Let $\tau \in \mathcal{T}$ be a fixed pinning and let $\{\Gamma_{v,k} \subseteq \mathbb{C} : (v,k) \in \mathcal{P}_1^\tau\}$ be a collection of non-empty open connected regions such that $1 \in \Gamma_{v,k}$ for each $(v,k) \in \mathcal{P}_1^\tau$. For every $v \in V^\tau$ let $\Gamma_v \subseteq \mathbb{C}$ be the connected component of the intersection $\bigcap_{0 \neq k \in \Omega_v^\tau} \Gamma_{v,k}$ that contains 1. If the multivariate conditional partition function Z_w^τ is $(\prod_{(v,k) \in \mathcal{P}_1^\tau} \Gamma_{v,k})$ -stable, then the influence matrix Ψ_μ^τ under the pinning τ satisfies*

$$\text{EigMax}(\Psi_\mu^\tau) \leq \|\Psi_\mu^\tau\|_\infty \leq \min \left\{ \frac{2}{b\delta^2}, \frac{8}{\delta} \max_{\substack{\tau \in \mathcal{T} \\ (v,k) \in \mathcal{P}^\tau}} \text{dist}(1, C_{v,k}^\tau) \right\},$$

where: b is the marginal bound for μ ;

$$\delta := \min_{(v,k) \in \mathcal{P}_1^\tau} \text{dist}(1, \partial\Gamma_{v,k}); \quad p_{v,k}^\tau := \mu^\tau(\sigma_v = k); \quad (8)$$

$$C_{v,k}^\tau := -\frac{1}{p_{v,k}^\tau} (\Gamma_{v,k} - 1)^{-1} = \left\{ -\frac{1}{p_{v,k}^\tau(z-1)} : z \in \Gamma_{v,k} \setminus \{1\} \right\}, \quad \text{for } k \neq 0; \quad (9)$$

$$C_{v,0}^\tau := \frac{1}{p_{v,0}^\tau} ((\Gamma_v - 1)^{-1} + 1) = \left\{ \frac{z}{p_{v,0}^\tau(z-1)} : z \in \Gamma_v \setminus \{1\} \right\}. \quad (10)$$

The following two lemmas are helpful for bounding the distance $\text{dist}(1, C_{v,k}^\tau)$ in Lemma 3.4.

LEMMA 3.5. *Let $\tau \in \mathcal{T}$ and $(v,k) \in \mathcal{P}_1^\tau$.*

1. *If $\Gamma_{v,k}$ is unbounded, then*

$$\text{dist}(1, C_{v,k}^\tau) \leq 1.$$

2. *Let $\alpha_{v,k} = \inf(\Gamma_{v,k} \cap \mathbb{R}_+)$ and $\beta_{v,k} = \sup(\Gamma_{v,k} \cap \mathbb{R}_+)$. Then*

$$\text{dist}(1, C_{v,k}^\tau) \leq \min \left\{ \frac{\alpha_{v,k}}{p_{v,k}^\tau(1 - \alpha_{v,k})} + \frac{1 - p_{v,k}^\tau}{p_{v,k}^\tau}, \frac{1}{p_{v,k}^\tau(\beta_{v,k} - 1)} + 1 \right\},$$

with the convention that $\frac{1}{\infty} = 0$ if $\beta_{v,k} = \infty$.

LEMMA 3.6. *Let $\tau \in \mathcal{T}$ and $(v,0) \in \mathcal{P}^\tau$.*

1. *If $0 \in \overline{\Gamma_v}$, then*

$$\text{dist}(1, C_{v,0}^\tau) \leq 1.$$

2. *Let $\alpha_v = \inf(\Gamma_v \cap \mathbb{R}_+)$ and $\beta_v = \sup(\Gamma_v \cap \mathbb{R}_+)$. Then*

$$\text{dist}(1, C_{v,0}^\tau) \leq \min \left\{ \frac{\alpha_v}{p_{v,0}^\tau(1 - \alpha_v)} + 1, \frac{1}{p_{v,0}^\tau(\beta_v - 1)} + \frac{1 - p_{v,0}^\tau}{p_{v,0}^\tau} \right\},$$

with the convention that $\frac{1}{\infty} = 0$ if $\beta_v = \infty$.

Combining Lemmas 3.4 to 3.6, we are able to establish Theorem 3.2. The proofs of Lemmas 3.5 and 3.6 are technical and we postpone them to Appendix B. The proof of Lemma 3.4 is presented in Section 3.3.

PROOF OF THEOREM 3.2. As discussed at the beginning of this subsection, we can reweight the configurations by Equation (6) and consider the all-one external fields under the new weights, so that Lemma 3.4 applies. In particular, for an arbitrary pinning $\tau \in \mathcal{T}$ the reweighted conditional partition function Z_w^τ given by Equation (7) is $(\prod_{(v,k) \in \mathcal{P}_1^\tau} \tilde{\Gamma}_{v,k})$ -stable where $\tilde{\Gamma}_{v,k} = \frac{1}{\lambda_{v,k}} \Gamma_{v,k}$ for each $(v,k) \in \mathcal{P}_1$. Thus, the first upper bound in Lemma 3.4 implies that the Gibbs distribution $\mu = \mu_{w,\lambda} = \mu_{\tilde{w},1}$ is spectrally independent with constant $\eta = 2/(b\delta^2)$ where b is the marginal bound for μ and

$$\delta = \min_{(v,k) \in \mathcal{P}_1} \text{dist}\left(1, \partial \tilde{\Gamma}_{v,k}\right) = \min_{(v,k) \in \mathcal{P}_1} \frac{1}{\lambda_{v,k}} \text{dist}\left(\lambda_{v,k}, \partial \Gamma_{v,k}\right)$$

as claimed.

For each $v \in V$, let $\tilde{\Gamma}_v = \Gamma_v \subseteq \mathbb{C}$ be the connected component of the intersection $\bigcap_{k \in Q_1} \tilde{\Gamma}_{v,k} = \bigcap_{k \in Q_1} \frac{1}{\lambda_{v,k}} \Gamma_{v,k}$ that contains 1. Then we further have the following.

1. If for every $v \in V$ the region $\tilde{\Gamma}_v$ is unbounded and 0 is contained in the closure of $\tilde{\Gamma}_v$, then the first part of Lemmas 3.5 and 3.6 implies that $\text{dist}(1, C_{v,k}^\tau) \leq 1$ for all $\tau \in \mathcal{T}$ and $(v,k) \in \mathcal{P}^\tau$. Hence, by the second upper bound in Lemma 3.4 spectral independence holds with constant $\eta = 8/\delta$.

If $\mathbb{R}_+ \subseteq \Gamma_{v,k}$ for each $(v,k) \in \mathcal{P}_1$, then by definition $\tilde{\Gamma}_v$ is unbounded and 0 is contained in the closure of $\tilde{\Gamma}_v$ for every $v \in V$; therefore, spectral independence holds with $\eta = 8/\delta$.

2. If there exists $\lambda^* \in \mathbb{R}_+$ such that $\lambda_{v,k} \in (0, \lambda^*) \subseteq \Gamma_{v,k}$ for every $(v,k) \in \mathcal{P}_1$, then one has

$$\begin{aligned} \alpha_{v,k} &= \inf\left(\tilde{\Gamma}_{v,k} \cap \mathbb{R}_+\right) = 0, & \beta_{v,k} &= \sup\left(\tilde{\Gamma}_{v,k} \cap \mathbb{R}_+\right) \geq \frac{\lambda^*}{\lambda_{v,k}}, \\ \alpha_v &= \inf\left(\tilde{\Gamma}_v \cap \mathbb{R}_+\right) = 0, & \beta_v &= \sup\left(\tilde{\Gamma}_v \cap \mathbb{R}_+\right) \geq \frac{\lambda^*}{\lambda_{v,\max}}, \end{aligned}$$

where $\lambda_{v,\max} = \max_{k \in Q_1} \lambda_{v,k} \leq \lambda_{\max}$. Thus, we deduce from the second part of Lemmas 3.5 and 3.6 that for all $\tau \in \mathcal{T}$ and $(v,k) \in \mathcal{P}^\tau$,

$$\begin{aligned} \text{dist}\left(1, C_{v,k}^\tau\right) &\leq \min\left\{\frac{1 - p_{v,k}^\tau}{p_{v,k}^\tau}, \frac{\lambda_{v,k}}{p_{v,k}^\tau(\lambda^* - \lambda_{v,k})} + 1\right\} \\ &\leq \min\left\{\frac{1 - b}{b}, \frac{\lambda_{\max}}{b(\lambda^* - \lambda_{\max})} + 1\right\}, \quad \text{for } k \neq 0; \\ \text{dist}\left(1, C_{v,0}^\tau\right) &\leq \min\left\{1, \frac{\lambda_{v,\max}}{p_{v,0}^\tau(\lambda^* - \lambda_{v,\max})} + \frac{1 - p_{v,0}^\tau}{p_{v,0}^\tau}\right\} \leq 1. \end{aligned}$$

The second bound in Lemma 3.4 then yields the desired bound on spectral independence. Note that we may assume μ is supported on at least two configurations so that $(1-b)/b \geq 1$, namely $b \leq 1/2$; otherwise μ is concentrated on a single configuration and spectral independence holds with constant 0.

3. If there exists $\lambda^* \in \mathbb{R}_+$ such that $\lambda_{v,k} \in (\lambda^*, \infty) \subseteq \Gamma_{v,k}$ for every $(v, k) \in \mathcal{P}_1$, then one has

$$\alpha_{v,k} = \inf \left(\tilde{\Gamma}_{v,k} \cap \mathbb{R}_+ \right) \leq \frac{\lambda^*}{\lambda_{v,k}}, \quad \beta_{v,k} = \sup \left(\tilde{\Gamma}_{v,k} \cap \mathbb{R}_+ \right) = \infty,$$

$$\alpha_v = \inf \left(\tilde{\Gamma}_v \cap \mathbb{R}_+ \right) \leq \frac{\lambda^*}{\lambda_{v,\min}}, \quad \beta_v = \sup \left(\tilde{\Gamma}_v \cap \mathbb{R}_+ \right) = \infty,$$

where $\lambda_{v,\min} = \min_{k \in Q_1} \lambda_{v,k}$. Thus, we deduce from the second part of Lemmas 3.5 and 3.6 that for all $\tau \in \mathcal{T}$ and $(v, k) \in \mathcal{P}^\tau$,

$$\text{dist} \left(1, C_{v,k}^\tau \right) \leq \min \left\{ \frac{\lambda^*}{p_{v,k}^\tau (\lambda_{v,k} - \lambda^*)} + \frac{1 - p_{v,k}^\tau}{p_{v,k}^\tau}, 1 \right\} \leq 1, \quad \text{for } k \neq 0;$$

$$\text{dist} \left(1, C_{v,0}^\tau \right) \leq \min \left\{ \frac{\lambda^*}{p_{v,0}^\tau (\lambda_{v,\min} - \lambda^*)} + 1, \frac{1 - p_{v,0}^\tau}{p_{v,0}^\tau} \right\}$$

$$\leq \min \left\{ \frac{\lambda^*}{b(\lambda_{\min} - \lambda^*)} + 1, \frac{1 - b}{b} \right\}.$$

The second bound in Lemma 3.4 then yields the bound on spectral independence as wanted. ■

REMARK 3.7. For Item 2 of Theorem 3.2 our proof actually yields a more complicated but stronger constant for spectral independence:

$$\eta = \frac{8}{\delta} \max \left\{ 1, \max_{\substack{\tau \in \mathcal{T} \\ (v,k) \in \mathcal{P}_1^\tau}} \left\{ \min \left\{ \frac{1 - p_{v,k}^\tau}{p_{v,k}^\tau}, \frac{\lambda_{v,k}}{p_{v,k}^\tau (\lambda^* - \lambda_{v,k})} + 1 \right\} \right\} \right\}.$$

For Item 3, the constant is:

$$\eta = \frac{8}{\delta} \max \left\{ 1, \max_{\substack{\tau \in \mathcal{T} \\ (v,0) \in \mathcal{P}^\tau}} \left\{ \min \left\{ \frac{1 - p_{v,0}^\tau}{p_{v,0}^\tau}, \frac{\lambda^*}{p_{v,0}^\tau (\lambda_{v,\min} - \lambda^*)} + 1 \right\} \right\} \right\}.$$

These two bounds are more robust in the sense of Remark 1.9, namely, when some external fields are close to λ^* while others are close to 0 (respectively, ∞).

3.3 Bounding the Absolute Sum of Influences: Proof of Lemma 3.4

Let $\tau \in \mathcal{T}$ be an arbitrary pinning and fix τ . We will give an upper bound on the absolute row sum of the associated influence matrix Ψ_μ^τ under τ , which then provides an upper bound on

the maximum eigenvalue of Ψ_μ^τ . In particular, for $(u, j) \in \mathcal{P}^\tau$ we define

$$\tilde{S}_\mu^\tau(u, j) := \sum_{(v,k) \in \mathcal{P}^\tau} |\Psi_\mu^\tau(u, j; v, k)| = \sum_{(v,k) \in \mathcal{P}^\tau: v \neq u} |\mu^\tau(\sigma_v = k \mid \sigma_u = j) - \mu^\tau(\sigma_v = k)| \quad (11)$$

to be the absolute sum of influences in the row (u, j) , and define

$$S_\mu^\tau(u, j) := \sum_{(v,k) \in \mathcal{P}_1^\tau} |\Psi_\mu^\tau(u, j; v, k)| = \sum_{(v,k) \in \mathcal{P}_1^\tau: v \neq u} |\mu^\tau(\sigma_v = k \mid \sigma_u = j) - \mu^\tau(\sigma_v = k)| \quad (12)$$

to be the partial absolute sum of influences for pairs (v, k) with $k \neq 0$ in the row (u, j) . Notice that one has $\tilde{S}_\mu^\tau(u, j) \leq 2S_\mu^\tau(u, j)$, because for each $(v, 0) \in \mathcal{P}^\tau$ it follows from the triangle inequality that

$$|\mu^\tau(\sigma_v = 0 \mid \sigma_u = j) - \mu^\tau(\sigma_v = 0)| \leq \sum_{k \in \Omega_v^\tau \setminus \{0\}} |\mu^\tau(\sigma_v = k \mid \sigma_u = j) - \mu^\tau(\sigma_v = k)|.$$

Hence,

$$\text{EigMax}(\Psi_\mu^\tau) \leq \|\Psi_\mu^\tau\|_\infty = \max_{(u,j) \in \mathcal{P}^\tau} \tilde{S}_\mu^\tau(u, j) \leq 2 \max_{(u,j) \in \mathcal{P}^\tau} S_\mu^\tau(u, j). \quad (13)$$

The rest of the proof aims to bound $S_\mu^\tau(u, j)$ for a fixed $(u, j) \in \mathcal{P}^\tau$. We consider two cases $j \neq 0$ and $j = 0$ separately, and prove the following.

LEMMA 3.8. *For $j \neq 0$ we have*

$$S_\mu^\tau(u, j) \leq \min \left\{ \frac{1}{b\delta^2}, \frac{4}{\delta} \text{dist} \left(1, C_{u,j}^\tau \right) \right\}.$$

LEMMA 3.9. *For $j = 0$ we have*

$$S_\mu^\tau(u, 0) \leq \min \left\{ \frac{1}{b\delta^2}, \frac{4}{\delta} \text{dist} \left(1, C_{u,0}^\tau \right) \right\}.$$

Lemma 3.4 follows immediately from these two lemmas.

PROOF OF LEMMA 3.4. Combining Equation (13) and Lemmas 3.8 and 3.9. ■

3.3.1 Proof of Lemma 3.8

Fix $(u, j) \in \mathcal{P}_1^\tau$. We follow the proof approach of [2] and view $S_\mu^\tau(u, j)$ as the derivative of a certain function related to the partition function; the lemma then follows from an application of the Schwarz-Pick Theorem (Theorem 2.5) for bounding the derivative. For ease of notation we write

$$\mathcal{P}' := \{(v, k) \in \mathcal{P}_1^\tau : v \neq u\} \quad \text{and} \quad \mathcal{K} := \prod_{(v,k) \in \mathcal{P}'} \Gamma_{v,k}.$$

Recall that $p_{u,j}^\tau = \mu^\tau(\sigma_u = j)$. Define the multivariate complex function $f : \mathcal{K} \rightarrow \mathbb{C}$ as

$$f(\lambda) = \frac{1}{p_{u,j}^\tau} \frac{Z_w^{\tau \cup (u,j)}(\lambda)}{Z_w^\tau(\lambda \cup \mathbf{1}_u)}, \quad \text{for } \lambda \in \mathcal{K} \quad (14)$$

where $\tau \cup (u, j) \in \mathcal{T}$ is the pinning that combines τ and $\sigma_u = j$, and $\lambda \cup \mathbf{1}_u$ is the vector of external fields that combines λ and the all-one fields $\mathbf{1}_u$ at u (i.e., $\lambda_{u,j'} = 1$ for all $0 \neq j' \in \Omega_u^\tau$); notice that

$$Z_w^\tau(\lambda \cup \mathbf{1}_u) = \sum_{j' \in \Omega_u^\tau} Z_w^{\tau \cup (u,j')}(\lambda).$$

Note that f is well-defined since by our assumption $Z_w^\tau(\lambda \cup \mathbf{1}_u) \neq 0$ whenever each $\lambda_{v,k} \in \Gamma_{v,k}$. The following claim summarizes several important properties of the function f .

CLAIM 3.10. *Let $f : \mathcal{K} \rightarrow \mathbb{C}$ be the multivariate complex function defined by Equation (14).*

1. *The function f is well-defined and holomorphic on $\mathcal{K}$, and $f(\mathbf{1}) = 1$.*
2. *For every $(v, k) \in \mathcal{P}'$,*

$$\left. \frac{\partial f}{\partial \lambda_{v,k}} \right|_{\lambda=\mathbf{1}} = \Psi_\mu^\tau(u, j; v, k).$$

3. *Suppose that $f \neq 1$. Let $\mathcal{A} \subseteq \mathbb{C}$ be an open region defined as*

$$\mathcal{A} := -\frac{1}{p_{u,j}^\tau} (\Gamma_{u,j} - 1)^{-1}. \quad (15)$$

Then $1 \notin \overline{\mathcal{A}}$. Let $\mathcal{A}_1$ be the connected component of $\overline{\mathcal{A}}^c$ which contains 1. Then $\mathcal{A}_1$ is open and simply connected, and

$$\text{image}(f) \subseteq \mathcal{A}_1.$$

If $f \equiv 1$, then Item 2 of Claim 3.10 implies that $\Psi_\mu^\tau(u, j; v, k) = 0$ for all $(v, k) \in \mathcal{P}'$, and hence $S_\mu^\tau(u, j) = 0$. In the rest of the proof we assume that $f \neq 1$.

Given Claim 3.10, in order to bound $S_\mu^\tau(u, j)$ it suffices to bound $\|\nabla f(\mathbf{1})\|_1$. We do this by taking holomorphic functions $\varphi : \mathbb{D}(0, 1) \rightarrow \mathcal{K}$, $\psi : \mathcal{A}_1 \rightarrow \mathbb{D}(0, 1)$ and considering their composition with the function f . The bound on $\|\nabla f(\mathbf{1})\|_1$ would then follow from the Schwarz-Pick Theorem which bounds the derivative of a holomorphic function from the open unit disk into itself.

We now formalize this idea. Let $\varphi : \mathbb{D}(0, 1) \rightarrow \mathcal{K}$ be a holomorphic vector-valued function such that for every $(v, k) \in \mathcal{P}'$, the (v, k) -coordinate function $\varphi_{v,k} : \mathbb{D}(0, 1) \rightarrow \Gamma_{v,k}$ is holomorphic and satisfies $\varphi_{v,k}(0) = 1$ and $\varphi'_{v,k}(0) \in \mathbb{R}_+$ if $\Psi_\mu^\tau(u, j; v, k) \geq 0$ while $\varphi'_{v,k}(0) \in \mathbb{R}_-$ if $\Psi_\mu^\tau(u, j; v, k) \leq 0$. Hence, $\varphi(0) = \mathbf{1}$ and $\varphi'_{v,k}(0) \Psi_\mu^\tau(u, j; v, k) \geq 0$ for all (v, k) . Meanwhile, for the region $\mathcal{A}_1$ given in Item 3 of Claim 3.10, let $\psi : \mathcal{A}_1 \rightarrow \mathbb{D}(0, 1)$ be a holomorphic function such that $\psi'(1) \in \mathbb{R}_+$. We will specify our choice of φ and ψ soon. Also, we point out here that our assumptions $\varphi'_{v,k}(0) \in \mathbb{R}_+/\mathbb{R}_-$ and $\psi'(1) \in \mathbb{R}_+$ would not cause strong restrictions; they can be easily satisfied by considering rotations $\varphi(e^{i\theta}z)$ and $e^{i\theta}\psi(z)$.

Given such φ and ψ , we define the holomorphic function $F : \mathbb{D}(0, 1) \rightarrow \mathbb{D}(0, 1)$ given by $F = \psi \circ f \circ \varphi$. Notice that $F(0) = \psi(1)$. The derivative $F'(0)$ at 0 is real and can be bounded by

$$\begin{aligned} F'(0) &= \psi'(1) \sum_{(v,k) \in \mathcal{P}'} \varphi'_{v,k}(0) \left. \frac{\partial f}{\partial \lambda_{v,k}} \right|_{\lambda=1} \\ &= \psi'(1) \sum_{(v,k) \in \mathcal{P}'} \varphi'_{v,k}(0) \Psi_\mu^\tau(u, j; v, k) \\ &\geq \psi'(1) \min_{(v,k) \in \mathcal{P}'} \left\{ \left| \varphi'_{v,k}(0) \right| \right\} S_\mu^\tau(u, j), \end{aligned} \quad (16)$$

where the second equality follows from Item 2 of Claim 3.10 and the inequality is due to our assumption that $\varphi'_{v,k}(0) \Psi_\mu^\tau(u, j; v, k) \geq 0$ for each (v, k) . The Schwarz-Pick Theorem (Theorem 2.5) implies that $F'(0) \leq 1$, and hence we obtain

$$S_\mu^\tau(u, j) \leq \frac{1}{\psi'(1)} \left(\min_{(v,k) \in \mathcal{P}'} \left| \varphi'_{v,k}(0) \right| \right)^{-1}. \quad (17)$$

It remains to choose φ and ψ . Consider first the function φ . For each $(v, k) \in \mathcal{P}'$ we let

$$\delta_{v,k} := \text{dist}(1, \partial\Gamma_{v,k}) \quad \text{and} \quad \chi_{v,k} := \text{sgn}(\Psi_\mu^\tau(u, j; v, k)) = \begin{cases} +1, & \Psi_\mu^\tau(u, j; v, k) \geq 0; \\ -1, & \Psi_\mu^\tau(u, j; v, k) < 0. \end{cases} \quad (18)$$

We then define $\varphi_{v,k} : \mathbb{D}(0, 1) \rightarrow \Gamma_{v,k}$ by

$$\varphi_{v,k}(z) = 1 + \chi_{v,k} \delta_{v,k} z. \quad (19)$$

Observe that φ is holomorphic, $\varphi(0) = 1$, and $\varphi'_{v,k}(z) = \chi_{v,k} \delta_{v,k}$ has the same sign as $\Psi_\mu^\tau(u, j; v, k)$ for each (v, k) . Recall that $\delta = \min_{(v,k) \in \mathcal{P}_1} \text{dist}(1, \partial\Gamma_{v,k})$, and thus

$$\min_{(v,k) \in \mathcal{P}'} \left| \varphi'_{v,k}(0) \right| = \min_{(v,k) \in \mathcal{P}'} \delta_{v,k} \geq \delta. \quad (20)$$

Next, we decide ψ . We will actually give two choices of ψ , denoted by ψ_1 and ψ_2 respectively, which correspond to the two bounds in Lemma 3.8.

We first consider the simpler choice ψ_1 . Let $\delta_{u,j} = \text{dist}(1, \partial\Gamma_{u,j}) \geq \delta$, and so $\mathbb{D}(1, \delta_{u,j}) \subseteq \Gamma_{u,j}$. Then, the region $\mathcal{A}$ from Item 3 of Claim 3.10 satisfies

$$\mathcal{A} = -\frac{1}{p_{u,j}^\tau} (\Gamma_{u,j} - 1)^{-1} \supseteq -\frac{1}{p_{u,j}^\tau} (\mathbb{D}(1, \delta_{u,j}) - 1)^{-1} = \frac{1}{p_{u,j}^\tau \delta_{u,j}} \overline{\mathbb{D}(0, 1)}^c.$$

It follows that

$$\mathcal{A}_1 \subseteq \overline{\mathcal{A}}^c \subseteq \frac{1}{p_{u,j}^\tau \delta_{u,j}} \mathbb{D}(0, 1).$$

We can define $\psi_1 : \mathcal{A}_1 \rightarrow \mathbb{D}(0, 1)$ as

$$\psi_1(z) = p_{u,j}^\tau \delta_{u,j} z.$$

Then, ψ_1 is holomorphic, $\psi'_1(z) = p_{u,j}^\tau \delta_{u,j} \in \mathbb{R}_+$, and

$$\frac{1}{\psi'_1(1)} = \frac{1}{p_{u,j}^\tau \delta_{u,j}} \leq \frac{1}{b\delta}. \quad (21)$$

Combining Equations (17), (20) and (21), we obtain

$$S_\mu^\tau(u, j) \leq \frac{1}{b\delta^2}.$$

This shows the first bound in Lemma 3.8.

Next, we define ψ_2 . Since $\emptyset \neq \mathcal{A}_1 \subsetneq \mathbb{C}$ is open and simply connected by Item 3 of Claim 3.10, the Riemann Mapping Theorem (Theorem 2.6) implies that there exists a (unique) biholomorphic mapping $\psi_2 : \mathcal{A}_1 \rightarrow \mathbb{D}(0, 1)$ such that $\psi_2(1) = 0$ and $\psi'_2(1) \in \mathbb{R}_+$. Write $h = \psi_2^{-1}$, which is a bijective holomorphic function from $\mathbb{D}(0, 1)$ to $\mathcal{A}_1$ satisfying $h(0) = 1$. Then, Koebe's One-Quarter Theorem (Theorem 2.7) shows that

$$\frac{1}{4}|h'(0)| \leq \text{dist}(1, \partial\mathcal{A}_1) \leq \text{dist}(1, \mathcal{A}) = \text{dist}\left(1, C_{u,j}^\tau\right).$$

It follows that

$$\frac{1}{\psi'_2(1)} = h'(0) \leq 4\text{dist}\left(1, C_{u,j}^\tau\right). \quad (22)$$

Combining Equations (17), (20) and (22), we get

$$S_\mu^\tau(u, j) \leq \frac{4}{\delta}\text{dist}\left(1, C_{u,j}^\tau\right),$$

which is the second bound in Lemma 3.8.

REMARK 3.11. The proof of Lemma 3.8 (and also Lemma 3.9 in Section 3.3.2) leaves the possibility of further improvements on the spectral independence bounds for specific problems. Here in the proof we are given regions $\mathcal{K}$ and $\mathcal{A}_1$ in abstract forms and the choices of φ and ψ may not be optimal for specific instances; in particular, the Riemann Mapping Theorem only shows the existence of a biholomorphic mapping and there is no guarantee that such a choice is the best possible. Hence, for specific problems and specific zero-free regions, one may be able to pick φ and ψ in a smarter way to achieve a better bound on spectral independence.

It remains to prove Claim 3.10. The following lemma is helpful to us.

LEMMA 3.12. *Let $S \subseteq \mathbb{C}$ be a non-empty open connected region such that S is unbounded and $\overline{S} \neq \mathbb{C}$. If S_1 is a connected component of $\overline{S}^c$, then S_1 is open and simply connected.*

PROOF. Clearly S_1 is open and connected. If S_1 is not simply connected, then there exists a Jordan curve (simple closed curve) γ in S_1 whose interior region contains a point $z_0 \notin S_1$. Note that we can actually find a point z from the interior of γ such that $z \in \overline{S}$; if not, then the whole interior of γ is contained in $\overline{S}^c$ and thus $z_0 \notin S_1$ is connected to S_1 in $\overline{S}^c$, contradicting

to the assumption that S_1 is a connected component of $\overline{S}^c$. Since the interior of γ is open, this further implies that the interior of γ contains a point $z \in S$. Meanwhile, since S is unbounded the exterior of γ contains a point $w \in S$. Now, as S is connected there exists a path p in S connecting z and w . Note that p must intersect with γ , because the interior and exterior of γ are disconnected. This yields a contradiction since $\gamma \subseteq S_1 \subseteq S^c$ while $p \subseteq S$. ■

Notice that the assumption of S being unbounded is necessary; for example, if $S = \mathbb{D}(0, 1)$ then $\overline{S}^c = \mathbb{C} \setminus \overline{\mathbb{D}(0, 1)} = \{z \in \mathbb{C} : |z| > 1\}$ is open and connected, but not simply connected.

We complete the proof of Lemma 3.8 with the proof of Claim 3.10.

PROOF OF CLAIM 3.10. 1. Since $Z_w^\tau(\lambda \cup \mathbf{1}_u) \neq 0$ whenever $\lambda \in \mathcal{K}$ by our stability assumption, the function f is well-defined and holomorphic on $\mathcal{K}$. Also, by definition we have $f(\mathbf{1}) = 1$.

2. Let $(v, k) \in \mathcal{P}'$. Then one has

$$\frac{\partial f}{\partial \lambda_{v,k}} = \frac{1}{p_{u,j}^\tau} \left(\frac{1}{Z_w^\tau(\lambda \cup \mathbf{1}_u)} \left(\frac{\partial}{\partial \lambda_{v,k}} Z_w^{\tau \cup (u,j)}(\lambda) \right) - \frac{Z_w^{\tau \cup (u,j)}(\lambda)}{Z_w^\tau(\lambda \cup \mathbf{1}_u)^2} \left(\frac{\partial}{\partial \lambda_{v,k}} Z_w^\tau(\lambda \cup \mathbf{1}_u) \right) \right).$$

Suppose τ is a pinning on $\Lambda \subseteq V$ and let $U = V^{\tau \cup (u,j)} = V \setminus \Lambda \setminus \{u\}$ be the set of unpinned vertices under the pinning $\tau \cup (u, j)$. We deduce that,

$$\begin{aligned} \frac{\partial}{\partial \lambda_{v,k}} Z_w^{\tau \cup (u,j)}(\lambda) &= \sum_{\sigma \in \Omega: \sigma_\Lambda = \tau, \sigma_u = j} w(\sigma) \cdot \frac{\partial}{\partial \lambda_{v,k}} \lambda^{\sigma_U} \\ &= \sum_{\sigma \in \Omega: \sigma_\Lambda = \tau, \sigma_u = j, \sigma_v = k} w(\sigma) \lambda^{\sigma_{U \setminus \{v\}}} \\ &= Z_w^{\tau \cup (u,j) \cup (v,k)}(\lambda). \end{aligned}$$

Similarly,

$$\frac{\partial}{\partial \lambda_{v,k}} Z_w^\tau(\lambda \cup \mathbf{1}_u) = Z_w^{\tau \cup (v,k)}(\lambda \cup \mathbf{1}_u).$$

We then get

$$\begin{aligned} \left. \frac{\partial f}{\partial \lambda_{v,k}} \right|_{\lambda=\mathbf{1}} &= \frac{1}{p_{u,j}^\tau} \left(\frac{Z_w^{\tau \cup (u,j) \cup (v,k)}(\mathbf{1})}{Z_w^\tau(\mathbf{1})} - \frac{Z_w^{\tau \cup (u,j)}(\mathbf{1}) \cdot Z_w^{\tau \cup (v,k)}(\mathbf{1})}{Z_w^\tau(\mathbf{1})^2} \right) \\ &= \frac{1}{\mu^\tau(\sigma_u = j)} (\mu^\tau(\sigma_u = j, \sigma_v = k) - \mu^\tau(\sigma_u = j) \mu^\tau(\sigma_v = k)) \\ &= \Psi_\mu^\tau(u, j; v, k), \end{aligned}$$

as claimed.

3. We first show that $\text{image}(f) \subseteq \mathcal{A}^c$. Suppose for sake of contradiction that $f(\lambda) \in \mathcal{A}$ for some $\lambda \in \mathcal{K}$. Then there exists $1 \neq y \in \Gamma_{u,j}$ such that

$$-\frac{1}{p_{u,j}^\tau(y-1)} = f(\lambda) = \frac{1}{p_{u,j}^\tau} \frac{Z_w^{\tau \cup (u,j)}(\lambda)}{Z_w^\tau(\lambda \cup \mathbf{1}_u)}.$$

It follows that

$$Z_w^\tau(\lambda \cup \lambda_u) = y Z_w^{\tau \cup (u,j)}(\lambda) + \sum_{j \neq j' \in \Omega_u^\tau} Z_w^{\tau \cup (u,j')}(\lambda) = 0,$$

where λ_u is the vector of external fields at u defined by $\lambda_{u,j} = y$ and $\lambda_{u,j'} = 1$ for $j' \in \Omega_u^\tau \setminus \{0, j\}$. This contradicts our stability assumption that $Z_w^\tau(\lambda \cup \lambda_u) \neq 0$. Therefore, we have shown that $\text{image}(f) \subseteq \mathcal{A}^c$.

Now, since $\mathcal{K}$ is open and connected and f is a non-constant holomorphic function, the Open Mapping Theorem (Theorem 2.8) implies that $\text{image}(f)$ is open and connected. Thus, we have $\text{image}(f) \subseteq (\mathcal{A}^c)^o = \overline{\mathcal{A}}^c$; note that in particular $1 \in \overline{\mathcal{A}}^c$. Furthermore, since $\text{image}(f)$ is connected one has $\text{image}(f) \subseteq \mathcal{A}_1$, the connected component of $\overline{\mathcal{A}}^c$ containing 1. The region $\mathcal{A}_1$ is open and connected by definition. It remains to show that $\mathcal{A}_1$ is simply connected, which follows immediately from Lemma 3.12 and the fact that $\mathcal{A}$ is connected and unbounded. ■

3.3.2 Proof of Lemma 3.9

The proof of Lemma 3.9 is similar to that of Lemma 3.8. We will use the same notations and only emphasize a few key steps that differ.

Recall that

$$\mathcal{P}' = \{(v, k) \in \mathcal{P}_1^\tau : v \neq u\} \quad \text{and} \quad \mathcal{K} = \prod_{(v,k) \in \mathcal{P}'} \Gamma_{v,k}.$$

Define the multivariate complex function $g : \mathcal{K} \rightarrow \mathbb{C}$ as

$$g(\lambda) = \frac{1}{p_{u,0}^\tau} \frac{Z_w^{\tau \cup (u,0)}(\lambda)}{Z_w^\tau(\lambda \cup \mathbf{1}_u)}, \quad \text{for } \lambda \in \mathcal{K} \quad (23)$$

where $\tau \cup (u, 0) \in \mathcal{T}$ is the pinning combining τ and $\sigma_u = 0$, and $\lambda \cup \mathbf{1}_u$ is the vector of external fields that combines λ and $\mathbf{1}_u$. The following claim is analogous to Claim 3.10 and summarizes key properties of the function g .

CLAIM 3.13. *Let $g : \mathcal{K} \rightarrow \mathbb{C}$ be the multivariate complex function defined by Equation (23).*

1. *The function g is well-defined and holomorphic on $\mathcal{K}$, and $g(\mathbf{1}) = 1$.*
2. *For every $(v, k) \in \mathcal{P}'$,*

$$\left. \frac{\partial g}{\partial \lambda_{v,k}} \right|_{\lambda=\mathbf{1}} = \Psi_\mu^\tau(u, 0; v, k).$$

3. *Suppose that $g \not\equiv 1$. Let $\mathcal{B} \subseteq \mathbb{C}$ be an open region defined as*

$$\mathcal{B} := \frac{1}{p_{u,0}^\tau} \left((\Gamma_u - 1)^{-1} + 1 \right). \quad (24)$$

Then $1 \notin \overline{\mathcal{B}}$. Let $\mathcal{B}_1$ be the connected component of $\overline{\mathcal{B}}^c$ which contains 1. Then $\mathcal{B}_1$ is open and simply connected, and

$$\text{image}(g) \subseteq \mathcal{B}_1.$$

We may assume that $g \neq 1$ since otherwise $S_\mu^\tau(u, 0) = 0$ and the lemma is trivial. Again we choose holomorphic functions $\varphi : \mathbb{D}(0, 1) \rightarrow \mathcal{K}$, $\psi : \mathcal{B}_1 \rightarrow \mathbb{D}(0, 1)$ and consider the holomorphic function $G : \mathbb{D}(0, 1) \rightarrow \mathbb{D}(0, 1)$ defined as $G = \psi \circ g \circ \varphi$. Just as in the proof of Lemma 3.8, we let $\varphi : \mathbb{D}(0, 1) \rightarrow \mathcal{K}$ be a holomorphic vector-valued function such that for every $(v, k) \in \mathcal{P}'$, the (v, k) -coordinate function $\varphi_{v,k} : \mathbb{D}(0, 1) \rightarrow \Gamma_{v,k}$ is holomorphic and satisfies $\varphi_{v,k}(0) = 1$ and $\varphi'_{v,k}(0) \in \mathbb{R}_+$ if $\Psi_\mu^\tau(u, j; v, k) \geq 0$ while $\varphi'_{v,k}(0) \in \mathbb{R}_-$ if $\Psi_\mu^\tau(u, j; v, k) \leq 0$. Meanwhile, let $\psi : \mathcal{B}_1 \rightarrow \mathbb{D}(0, 1)$ be a holomorphic function such that $\psi'(1) \in \mathbb{R}_+$. Hence, we have $G(0) = \psi(1)$, and by Claim 3.13 $G'(0) \in \mathbb{R}$ can be bounded by

$$G'(0) \geq \psi'(1) \min_{(v,k) \in \mathcal{P}'} \left\{ \left| \varphi'_{v,k}(0) \right| \right\} S_\mu^\tau(u, 0),$$

which is analogous to Equation (16). We then deduce the analog of Equation (17) from the Schwarz-Pick Theorem (Theorem 2.5):

$$S_\mu^\tau(u, 0) \leq \frac{1}{\psi'(1)} \left(\min_{(v,k) \in \mathcal{P}'} \left| \varphi'_{v,k}(0) \right| \right)^{-1}. \quad (25)$$

We specify next our choice of φ and ψ . The function φ is the same one as in the proof of Lemma 3.8, and is given by Equation (19). In particular, Equation (20) still holds. We also give two choices of the function ψ , denoted by ψ_3 and ψ_4 respectively, corresponding to the two bounds in Lemma 3.9.

Consider first ψ_3 . Recall that $\Gamma_u \subseteq \mathbb{C}$ is defined to be the connected component of the intersection $\bigcap_{0 \neq j \in \Omega_u^\tau} \Gamma_{u,j}$ that contains 1. Let $\delta_u = \text{dist}(1, \partial \Gamma_u) \geq \delta$ and thus $\mathbb{D}(1, \delta_u) \subseteq \Gamma_u$. Then we have

$$\mathcal{B} = \frac{1}{p_{u,0}^\tau} \left((\Gamma_u - 1)^{-1} + 1 \right) \supseteq \frac{1}{p_{u,0}^\tau} \left(\mathbb{D}(0, \delta_u)^{-1} + 1 \right) = \frac{1}{p_{u,0}^\tau \delta_u} \overline{\mathbb{D}}(\delta_u, 1)^c,$$

and hence

$$\mathcal{B}_1 \subseteq \overline{\mathcal{B}}^c \subseteq \frac{1}{p_{u,0}^\tau \delta_u} \mathbb{D}(\delta_u, 1).$$

We define $\psi_3 : \mathcal{B}_1 \rightarrow \mathbb{D}(0, 1)$ as

$$\psi_3(z) = p_{u,0}^\tau \delta_u z - \delta_u.$$

Observe that ψ_3 is holomorphic, $\psi'_3(z) = p_{u,0}^\tau \delta_u \in \mathbb{R}_+$, and

$$\frac{1}{\psi'_3(1)} = \frac{1}{p_{u,0}^\tau \delta_u} \leq \frac{1}{b\delta}. \quad (26)$$

Combining Equations (20), (25) and (26), we obtain

$$S_\mu^\tau(u, 0) \leq \frac{1}{b\delta^2}.$$

This shows the first bound in Lemma 3.9.

Finally, we define ψ_4 . Since $\emptyset \neq \mathcal{B}_1 \subsetneq \mathbb{C}$ is open and simply connected by Claim 3.13, there exists a (unique) biholomorphic mapping $\psi_4 : \mathcal{B}_1 \rightarrow \mathbb{D}(0, 1)$ such that $\psi_4(1) = 0$ and $\psi'_4(1) \in \mathbb{R}_+$ by the Riemann Mapping Theorem (Theorem 2.6). Let $h = \psi_4^{-1}$ be the holomorphic mapping from $\mathbb{D}(0, 1)$ to $\mathcal{B}_1$ with $h(0) = 1$. We deduce from the Koebe's One-Quarter Theorem (Theorem 2.7) that

$$\frac{1}{4}|h'(0)| \leq \text{dist}(1, \partial\mathcal{B}_1) \leq \text{dist}(1, \mathcal{B}) = \text{dist}\left(1, C_{u,0}^\tau\right),$$

and hence

$$\frac{1}{\psi'_4(1)} = h'(0) \leq 4\text{dist}\left(1, C_{u,0}^\tau\right). \quad (27)$$

Combining Equations (20), (25) and (27), we get

$$S_\mu^\tau(u, 0) \leq \frac{4}{\delta}\text{dist}\left(1, C_{u,0}^\tau\right),$$

which is the second bound in Lemma 3.9.

We end this section with the proof of Claim 3.13.

PROOF OF CLAIM 3.13. Item 1 follows from the stability of the partition function and Item 2 can be deduced by direct calculations. We omit the details here and refer to the proof of Claim 3.10.

For Item 3, again we first show that $\text{image}(g) \subseteq \mathcal{B}^c$. Suppose for sake of contradiction that $g(\lambda) \in \mathcal{B}$ for some $\lambda \in \mathcal{K}$. Then there exists $1 \neq y \in \Gamma_u \subseteq \bigcap_{0 \neq j \in \Omega_u^\tau} \Gamma_{u,j}$ such that

$$\frac{y}{p_{u,0}^\tau(y-1)} = g(\lambda) = \frac{1}{p_{u,0}^\tau} \frac{Z_w^{\tau \cup (u,0)}(\lambda)}{Z_w^\tau(\lambda \cup \mathbf{1}_u)}.$$

It follows that

$$Z_w^\tau(\lambda \cup y\mathbf{1}_u) = Z_w^{\tau \cup (u,0)}(\lambda) + \sum_{0 \neq j \in \Omega_u^\tau} y Z_w^{\tau \cup (u,j)}(\lambda) = 0,$$

where $y\mathbf{1}_u$ represents the vector of external fields at u defined by $\lambda_{u,j} = y$ for all $0 \neq j \in \Omega_u^\tau$. This contradicts the stability assumption of the partition function. Therefore, $\text{image}(g) \subseteq \mathcal{B}^c$. The Open Mapping Theorem (Theorem 2.8) then implies that $\text{image}(g) \subseteq \mathcal{B}_1$ which is the connected component of $\overline{\mathcal{B}}^c$ containing 1. Meanwhile, notice that the region Γ_u is open and connected since it is a connected component of the open set $\bigcap_{0 \neq j \in \Omega_u^\tau} \Gamma_{u,j}$, and so $\mathcal{B}$ is open, connected, and unbounded. Hence, Lemma 3.12 shows that $\mathcal{B}_1$ is open and simply connected. This completes the proof of the claim. ■

4. Spectral Independence for Binary Symmetric Holant Problems

Let $G = (V, E)$ be a graph of maximum degree Δ . We consider the Holant problem in the binary symmetric case, which we now describe. Let $\{f_v\}_{v \in V} : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ be a family of functions, one

for each vertex $v \in V$ in the input graph. One should think of each f_v as representing a local constraint on the assignments to edges incident to v . Since we are restricting ourselves to the binary case, our configurations σ will map edges to $\{0, 1\}$. Furthermore, since we are restricting ourselves to the symmetric case, our local functions f_v will only depend on the number of edges incident to v which are mapped to 1. With these $\{f_v\}_{v \in V}$ in hand, we may write the multivariate partition function as

$$Z_G(\lambda) = \sum_{\sigma: E \rightarrow \{0,1\}} \prod_{v \in V} f_v(|\sigma_{E(v)}|) \prod_{e \in E} \lambda_e^{\mathbb{1}_{\{\sigma_e=1\}}}, \quad (28)$$

where $E(v)$ is the set of all edges adjacent to v , $\sigma_{E(v)}$ is the configuration restricted on $E(v)$, and $|\sigma_{E(v)}|$ is the number of edges in $E(v)$ with assignment 1.

This class of problems is already incredibly rich, and encompasses many classical objects studied in combinatorics and statistical physics including the following:

- *Matchings/Monomer-Dimer Model*: Assume all f_v are the same and given by the “at-most-one” function:

$$f_v(k) = \begin{cases} 1, & \text{if } k = 0, 1; \\ 0, & \text{if } k \geq 2. \end{cases}$$

Then $Z_G(\mathbf{1})$ yields the number of matchings (of any size) in G .

- *Weighted Edge Covers*: Assume all f_v are the same and given by the weighted “at-least-one” function:

$$f_v(k) = \begin{cases} \rho, & \text{if } k = 0; \\ 1, & \text{if } k \geq 1. \end{cases}$$

In the case $\rho = 0$, then $Z_G(\mathbf{1})$ yields the number of edge covers of G , that is, subsets of edges such that every vertex is incident to at least one selected edge.

- *Weighted Even Subgraphs*: In this case, all f_v are the same and given by the weighted “parity” function. More specifically, for a fixed positive parameter $\rho > 0$, we have

$$f_v(k) = \begin{cases} 1, & \text{if } k \text{ is even;} \\ \rho, & \text{if } k \text{ is odd.} \end{cases}$$

In the case $\rho = 0$, then $Z_G(\mathbf{1})$ counts the number of even subgraphs, that is, subsets of edges such that all vertices have even degrees in the resulting subgraph. (Note that when $\rho = 0$, the Glauber dynamics is not ergodic.)

- *Ising Model on Line Graphs*: In this case, each f_v depends on the degree of v . If $\beta > 0$ is some fixed parameter (independent of v), and $d = \deg(v)$, then we have

$$f_v(k) = \begin{cases} \beta^{\binom{k}{2}} \beta^{\binom{d-k}{2}}, & \text{if } 0 \leq k \leq d; \\ 0, & \text{o/w.} \end{cases}$$

In all of the above examples, prior works managed to show that the Glauber dynamics admits an inverse polynomial spectral gap ([36] for matchings, [35] for edge covers, [37] for weighted even subgraphs, and [27] for the Ising model in the antiferromagnetic $\beta < 1$ regime). Furthermore, all of these results were obtained via the canonical paths method [36], and its winding extension [45]. However, one down-side behind these results is that the spectral gap bounds are suboptimal, and do not yield optimal mixing times nor sub-Gaussian concentration estimates. In contrast, by combining our framework with known zero-free regions for these models and the local-to-global mixing result of [24], we obtain optimal mixing times and sub-Gaussian concentration results for these problems in the bounded-degree regime.

One of the convenient aspects of our approach is that establishing the required zero-free region for the complicated multivariate partition function can be boiled down to establishing stability for a bounded-degree univariate polynomial with coefficients coming from the local functions f_v . This was one of the main insights of [53, 34, 15]. More specifically, if Δ is the maximum degree of the input graph $G = (V, E)$, and $f_v : [d] \rightarrow \mathbb{R}_{\geq 0}$ is the local function for some vertex $v \in V$, where $d = \deg(v) \leq \Delta$, then define the corresponding local polynomial at v by

$$P_v(z) = \sum_{k=0}^d \binom{d}{k} f_v(k) z^k. \quad (29)$$

A circular region on the complex plane is the interior or exterior of a disk, or an open half-plane. [34] showed using Asano–Ruelle contractions [5, 52] that in the case all f_v are the same, and all P_v are Φ -stable for an open half-plane $\Phi \subseteq \mathbb{C}$, the multivariate partition function is Γ -stable where $\Gamma = [-(\Phi^c)^2]^c$. This result actually holds for any circular region $\Phi \subseteq \mathbb{C}$ assuming that either Φ is convex or every local polynomial P_v has degree $\deg(v)$; under these assumptions one can apply the famous Grace–Walsh–Szegő Coincidence Theorem to the local polynomials, see [34, 18]. A straightforward generalization of their techniques yields the following.

THEOREM 4.1 ([34]). *Let $G = (V, E)$ be a graph. Let $\{f_v\}_{v \in V} : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ be a family of local functions, and let $\{\Phi_v\}_{v \in V}$ be a family of circular regions containing 0 such that for every $v \in V$, either Φ_v is convex or $f_v(\deg(v)) > 0$. If for every $v \in V$, the local polynomial P_v is Φ_v -stable, then the multivariate partition function $Z_G(\lambda)$ is $\prod_{e \in E} \Gamma_e$ stable, where for each edge $e = \{u, v\}$, $\Gamma_e = (-\Phi_u^c \cdot \Phi_v^c)^c \subseteq \mathbb{C}$.*

Using Theorem 4.1, [34] established zero-free regions for a large class of Holant problems satisfying generalized second-order recurrences, including matchings, weighted edge covers, and weighted even subgraphs. Our main theorems Theorems 1.1, 1.2 and 1.4 build upon these zero-free results as well as Theorems 1.7 and 2.2 (note that we can obtain spectral independence for matchings from Theorems 1.7 and 4.1, which was already known in [24] with a better bound by correlation decay proofs). Zero-free regions were also established for weighted edge covers and the antiferromagnetic Ising model on line graphs in [15], using techniques from [53].

Before proving the main theorems, we will need the following simple lemma concerning the case where the regions Φ_u are half-planes. Recall that $\mathbb{H}_\varepsilon = \{x + iy : x < -\varepsilon\}$ and $\overline{\mathbb{H}}_\varepsilon = \{x + iy : x \leq -\varepsilon\}$ for $\varepsilon \in \mathbb{R}_+$.

LEMMA 4.2 (Lemma 5 in [34]). *For $\varepsilon > 0$, let $\Gamma = (-\overline{\mathbb{H}}_\varepsilon^2)^c$ be a region. Then Γ contains $\mathbb{R}_+$, and for every $\lambda \in \mathbb{R}_+$ we have $\text{dist}(\lambda, \partial\Gamma) = \lambda + \varepsilon^2$ if $\lambda \in (0, \varepsilon^2)$, and $\text{dist}(\lambda, \partial\Gamma) = 2\varepsilon\sqrt{\lambda}$ if $\lambda \in [\varepsilon^2, \infty)$.*

For completeness, we provide a proof in Appendix B. With these tools in hand, we deduce strong zero-free regions for the above examples. We use these to prove our main mixing results Theorems 1.1, 1.2 and 1.4. Note that by Lemma 3.1 and Theorem 3.2, one can in fact establish rapid mixing results for these models with non-uniform external fields, though we only state the uniform case for simplicity.

PROOF OF THEOREM 1.1. By Theorem 2.2, it suffices to prove η -spectral independence for $\eta = O_{\Delta, \lambda, \rho}(1)$. By Theorem 1.7, it suffices to prove that the multivariate partition function Equation (28) is Γ -stable, where $\Gamma \subseteq \mathbb{C}$ is an open connected region containing $\mathbb{R}_+$ and $\delta = \frac{1}{\lambda} \text{dist}(\lambda, \partial\Gamma) = \Omega_{\Delta, \lambda, \rho}(1)$.

It is more convenient for us to work with the model on complements of weighted edge covers, whose partition function is the inversion of that for weighted edge covers. For this, the local polynomial is given by

$$P_v(z) = (1 + z)^{\deg(v)} - (1 - \rho)z^{\deg(v)},$$

which is $\overline{\mathbb{H}}_{1/2}^c$ -stable. Then by Theorem 4.1, the inversion of the weighted edge cover partition function $Z_G(\lambda)$ is $(-\overline{\mathbb{H}}_{1/2}^2)^c$ -stable, and therefore $Z_G(\lambda)$ is Γ -stable for

$$\Gamma = \left[\left(-\overline{\mathbb{H}}_{1/2}^2 \right)^c \right]^{-1} = [-\overline{\mathbb{D}}(-1, 1)^2]^c.$$

This region Γ is also derived in [15]. We remark that the region $-\overline{\mathbb{D}}(-1, 1)^2$ is cardioid-shaped, and its complement Γ is an open connected region containing $\mathbb{R}_+$; see Lemma 3.9 and Figure 1 in [15]. Hence, we have $\mathbb{R}_+ \subseteq \Gamma$ and $\delta = \Omega_{\Delta, \lambda, \rho}(1)$ as wanted. ■

PROOF OF THEOREM 1.2. We may assume $\rho \in (0, 1)$ since if $\rho = 1$ then we get a trivial product distribution. Once again, by Theorem 2.2, it suffices to prove η -spectral independence for $\eta = O_{\Delta, \lambda, \rho}(1)$, and by Theorem 1.7, it suffices to prove that the multivariate partition function Equation (28) is Γ -stable, where $\Gamma \subseteq \mathbb{C}$ is an open connected region containing $\mathbb{R}_+$ and $\delta = \frac{1}{\lambda} \text{dist}(\lambda, \partial\Gamma) = \Omega_{\Delta, \lambda, \rho}(1)$.

For this, observe that the local polynomial is given by

$$\begin{aligned} P_v(z) &= \sum_{k=0}^{\deg(v)} \binom{\deg(v)}{k} \left(\frac{1+\rho}{2} + \frac{1-\rho}{2}(-1)^k \right) z^k \\ &= \frac{1+\rho}{2} (1+z)^{\deg(v)} + \frac{1-\rho}{2} (1-z)^{\deg(v)}. \end{aligned}$$

Since $0 < \rho < 1$, the roots of P_v are given by $\frac{\omega - t_v}{\omega + t_v}$ where $\omega \in \mathbb{C}$ satisfies $\omega^{\deg(v)} = -1$, and $t_v \in \mathbb{R}_+$ is given by

$$t_v = \left(\frac{1 + \rho}{1 - \rho} \right)^{1/\deg(v)} > 1.$$

It follows that P_v is $\left[\overline{\mathbb{D}} \left(-\frac{t_v^2 + 1}{t_v^2 - 1}, \frac{2t_v}{t_v^2 - 1} \right) \right]^c$ -stable. Then by Theorem 4.1, $Z_G(\lambda)$ is $\prod_{e \in E} \Gamma_e$ -stable, where for each edge $e = uv \in E$,

$$\Gamma_e = \left[-\overline{\mathbb{D}} \left(-\frac{t_u^2 + 1}{t_u^2 - 1}, \frac{2t_u}{t_u^2 - 1} \right) \cdot \overline{\mathbb{D}} \left(-\frac{t_v^2 + 1}{t_v^2 - 1}, \frac{2t_v}{t_v^2 - 1} \right) \right]^c.$$

In particular, $Z_G(\lambda)$ is Γ -stable for

$$\Gamma = \left[-\overline{\mathbb{D}} \left(-\frac{t^2 + 1}{t^2 - 1}, \frac{2t}{t^2 - 1} \right) \right]^c \subseteq \Gamma_e, \quad \forall e \in E, \quad \text{where } t = \left(\frac{1 + \rho}{1 - \rho} \right)^{1/\Delta} > 1.$$

The region Γ is open and connected. Observe that we have $\Gamma \supseteq \left(-\overline{\mathbb{H}}_{\frac{t-1}{t+1}}^2 \right)^c$. Hence, by Lemma 4.2 we have $\mathbb{R}_+ \subseteq \Gamma$ and $\delta = \Omega_{\Delta, \lambda, \rho}(1)$ as wanted. ■

PROOF OF THEOREM 1.4. By Theorem 2.2 it suffices to prove η -spectral independence for $\eta = O_{\Delta, \beta, \gamma, \lambda}(1)$. By Theorem 1.7 it suffices to prove that the multivariate partition function Equation (28) is Γ -stable, where $\Gamma \subseteq \mathbb{C}$ is an open connected region containing $\mathbb{R}_+$ and $\delta = \frac{1}{\lambda} \text{dist}(\lambda, \partial\Gamma) = \Omega_{\Delta, \beta, \gamma, \lambda}(1)$.

For this, observe that the local polynomial is given by

$$P_v(z) = \sum_{k=0}^{\deg(v)} \binom{\deg(v)}{k} \beta^{(k)} \gamma^{\binom{\deg(v)-k}{2}} z^k.$$

By Proposition 4.3 below (see Section 4.1 for the proof), all roots of these polynomials are strictly negative reals, i.e., they are contained in $(-\infty, -\varepsilon_{\deg(v)}]$ for some constant $\varepsilon_{\deg(v)} = \varepsilon_{\deg(v)}(\beta, \gamma) > 0$ depending only on $\deg(v)$, β , γ . Then by Theorem 4.1, $Z_G(\lambda)$ is $\prod_{e \in E} \Gamma_e$ -stable, where for each edge $e = uv \in E$,

$$\Gamma_e = \left(-\overline{\mathbb{H}}_{\varepsilon_{\deg(u)}} \cdot \overline{\mathbb{H}}_{\varepsilon_{\deg(v)}} \right)^c.$$

In particular, $Z_G(\lambda)$ is Γ -stable for $\Gamma = \left(-\overline{\mathbb{H}}_{\varepsilon}^2 \right)^c$ where $\varepsilon = \min_{1 \leq d \leq \Delta} \varepsilon_d$ depends only on Δ , β , γ . The region Γ is open and connected, and by Lemma 4.2 it contains $\mathbb{R}_+$ and we have $\delta = \Omega_{\Delta, \beta, \gamma, \lambda}(1)$ as wanted. ■

4.1 Stability for Antiferromagnetic Two-Spin Edge Models

In this subsection, we analyze the roots of the local polynomial for antiferromagnetic two-spin edge models, which is needed in the proof of Theorem 1.4 above. We generalize a result due to [15] which proves that the local polynomial for the antiferromagnetic edge Ising model has

strictly negative real roots. We achieve this by generalizing their arguments to all antiferromagnetic two-spin edge models.

PROPOSITION 4.3 (Generalization of Lemma 4.3 in [15]). *For every $\beta \geq 0$, $\gamma > 0$ with $\beta\gamma < 1$ and every positive integer $d \geq 1$, the univariate polynomial*

$$P_d(z) = \sum_{k=0}^d \binom{d}{k} \beta^{\binom{k}{2}} \gamma^{\binom{d-k}{2}} z^k$$

has strictly negative real roots.

We prove this via an inductive approach, relying on the following decomposition of P_d .

LEMMA 4.4. *For every $\beta \geq 0$, $\gamma > 0$ and every positive integer $d \geq 1$, we have that*

$$P_{d+1}(z) = \gamma^d P_d(z/\gamma) + z P_d(\beta z).$$

PROOF. We have

$$\begin{aligned} P_{d+1}(z) &= \sum_{k=0}^{d+1} \underbrace{\binom{d+1}{k}}_{=\binom{d}{k} + \binom{d}{k-1}} \beta^{\binom{k}{2}} \gamma^{\binom{d+1-k}{2}} z^k \\ &= \sum_{k=0}^d \binom{d}{k} \beta^{\binom{k}{2}} \gamma^{\binom{d+1-k}{2}} z^k + \sum_{k=0}^d \binom{d}{k} \beta^{\binom{k+1}{2}} \gamma^{\binom{d-k}{2}} z^{k+1} \\ &= \sum_{k=0}^d \binom{d}{k} \beta^{\binom{k}{2}} \gamma^{\binom{d-k}{2}} \gamma^{d-k} z^k + z \sum_{k=0}^d \binom{d}{k} \beta^{\binom{k}{2}} \gamma^{\binom{d-k}{2}} \beta^k z^k \\ &= \gamma^d P_d(z/\gamma) + z P_d(\beta z). \end{aligned}$$

PROOF OF PROPOSITION 4.3. If $\beta = 0$ then P_d is linear and the proposition is immediate. We may assume $\beta > 0$. We prove via induction the following stronger claim: The roots $r_1 > \dots > r_d$ of P_d are distinct, real, and strictly negative, and further satisfy $r_i/r_{i+1} < \beta\gamma$. The cases $d = 0, 1$ are vacuous. When $d = 2$, the polynomial $P_2(z) = \beta z^2 + 2z + \gamma$ has roots $(-1 \pm \sqrt{1 - \beta\gamma})/\beta$, which are distinct, real, and strictly negative since $\beta\gamma < 1$. One can also check that $r_1/r_2 < \beta\gamma$ via a straightforward calculation. This establishes the base case.

Assume the stronger conclusion holds for some $d \geq 2$. By Lemma 4.4, we may write $P_{d+1}(z) = \gamma^d P_d(z/\gamma) + z P_d(\beta z)$. If $r_1 > \dots > r_d$ are the roots of P_d , then $\gamma r_1 > \dots > \gamma r_d$ are the roots of $\gamma^d P_d(z/\gamma)$, and $0 = r_0/\beta > r_1/\beta > \dots > r_d/\beta$ are the roots of $z P_d(\beta z)$, where for convenience we define $r_0 = 0$. First, we claim that the roots of $\gamma^d P_d(z/\gamma)$ interlace the roots of $z P_d(\beta z)$, i.e.,

$$0 = r_0/\beta > \gamma r_1 > r_1/\beta > \gamma r_2 > \dots > r_{d-1}/\beta > \gamma r_d > r_d/\beta.$$

To see this, observe that $\gamma r_i > r_i/\beta$ since $\beta\gamma < 1$, and $r_{i-1}/\beta > \gamma r_i$ since $r_{i-1}/r_i < \beta\gamma$ by the induction hypothesis for P_d .

Now, we claim that for each $i = 2, \dots, d$, the evaluations

$$P_{d+1}(\gamma r_i) = \gamma r_i P_d(\beta \gamma r_i) \quad \text{and} \quad P_{d+1}(r_{i-1}/\beta) = \gamma^d P_d(r_{i-1}/\beta \gamma)$$

are nonzero and have different signs. Observe that $\beta \gamma r_i, r_{i-1}/\beta \gamma \in (r_i, r_{i-1})$; hence, the evaluations $P_d(\beta \gamma r_i)$ and $P_d(r_{i-1}/\beta \gamma)$ are nonzero and have the same sign, and we deduce the claim by $r_i < 0$. It then follows from the Intermediate Value Theorem that P_{d+1} has a root $s_i \in (\gamma r_i, r_{i-1}/\beta)$ for each $i = 2, \dots, d$.

Moreover, P_{d+1} also has a root $s_1 \in (\gamma r_1, 0)$ and a root $s_{d+1} \in (-\infty, r_d/\beta)$. Observe that the evaluations $P_{d+1}(\gamma r_1) = \gamma r_1 P_d(\beta \gamma r_1)$ and $P_{d+1}(0) = \gamma^d P_d(0)$ are nonzero and have different signs since $0 > \beta \gamma r_1 > r_1$, and the Intermediate Value Theorem implies a root $s_1 \in (\gamma r_1, 0)$. Meanwhile, $P_{d+1}(r_d/\beta) = \gamma^d P_d(r_d/\beta \gamma)$ and $P_d(-\infty)$ are nonzero and have the same sign since $-\infty < r_d/\beta \gamma < r_d$. Also, $P_d(-\infty)$ and $P_{d+1}(-\infty)$ have different signs since the two polynomials differ in degree by 1. This shows that $P_{d+1}(r_d/\beta)$ and $P_{d+1}(-\infty)$ are nonzero and have different signs, and the Intermediate Value Theorem shows the existence of a root $s_{d+1} \in (-\infty, r_d/\beta)$.

To summarize, we prove that P_{d+1} has roots $s_1 > \dots > s_{d+1}$ which are distinct strictly negative real numbers and (taking $r_0 = 0$ and $r_{d+1} = -\infty$ for convenience) satisfy $s_i \in (\gamma r_i, r_{i-1}/\beta)$ for any $i = 1, \dots, d+1$. To finish the induction, we need to show that $s_i/s_{i+1} < \beta \gamma$ for all $i = 1, \dots, d$, which follows by $s_i/s_{i+1} < (\gamma r_i)/(r_i/\beta) = \beta \gamma$. ■

5. Spectral Independence for Weighted Graph Homomorphisms and Tensor Network Contractions

In this section, we study spectral independence for general tensor network contractions and weighted graph homomorphisms. Unlike binary symmetric Holant problems, where rapid mixing of the Glauber dynamics was already known for our main examples such as matchings [36], Ising model on line graphs [27], edge covers [35], and weighted even subgraphs [37], in the setting we consider here, rapid mixing for any local Markov chain was not known beyond simple and standard path coupling arguments. Prior works [14, 13, 50, 48] had studied these problems but only from the perspective of deterministic approximation algorithms using Barvinok's polynomial interpolation method [7]. While these algorithms run in polynomial time for bounded-degree graphs, the exponent typically depends on the maximum degree, and are more difficult to implement.

Here, we show that the Glauber dynamics mixes in $O(n \log n)$ steps for these problems on bounded-degree graphs, yielding significantly faster and simpler algorithms for computing the partition function. We again reduce rapid mixing to spectral independence via Theorem 2.2, and then reduce spectral independence to the existence of a sufficiently large zero-free region for the multivariate partition function via Theorem 3.2. Fortunately, such zero-free regions

were already obtained in prior works, as they are the entire basis for Barvinok's polynomial interpolation method. We leverage them here in a completely black-box manner.

5.1 Weighted Graph Homomorphisms

Here, we study weighted graph homomorphisms, which may also be viewed as spin systems on vertices. In the bounded-degree setting, we show that the Glauber dynamics on vertex configurations for these models mixes in $O(n \log n)$ steps, provided the weights are sufficiently close to 1. This is analogous to classical mixing results stating the Glauber dynamics mixes rapidly in the “high-temperature” regime.

THEOREM 5.1 (Spectral Independence for Weighted Graph Homomorphisms). *Fix a positive integer $q \geq 2$, let $G = (V, E)$ be a graph with maximum degree $\leq \Delta$, and for each edge $uv \in E$, let $A^{uv} \in \mathbb{R}_{\geq 0}^{q \times q}$ be a (not necessarily symmetric) nonnegative matrix. There exists a universal constant $\gamma \approx 0.56$ independent of $q, G, \{A^{uv}\}_{uv \in E}$ such that if $|A^{uv}(j, k) - 1| \leq \frac{\gamma}{\Delta + \gamma} - \varepsilon$ for some $\varepsilon > 0$, all $uv \in E$ and all $j, k \in [q]$, then the associated graph homomorphism distribution μ on vertex configurations $\sigma : V \rightarrow [q]$ given by*

$$\mu(\sigma) \propto \prod_{uv \in E} A^{uv}(\sigma(u), \sigma(v))$$

is η -spectrally independent for some constant $\eta = \eta(\Delta, \varepsilon)$. In particular, if $\Delta, \varepsilon = \Theta(1)$, then the Glauber dynamics for sampling from μ mixes in $O(n \log n)$ steps.

REMARK 5.2. A straightforward application of the classical Dobrushin uniqueness condition yields rapid mixing when $|A^{uv}(j, k) - 1| \lesssim \frac{1}{2\Delta}$ for all $uv \in E$ and all $j, k \in [q]$.

The zero-free region for the graph homomorphism partition function was studied in [13]. We state here a slightly more general theorem, the proof of which is included in Appendix A.1 for completeness.

THEOREM 5.3 (Zeros for Weighted Graph Homomorphisms; [13]). *Fix a positive integer $q \geq 2$, let $G = (V, E)$ be a graph with maximum degree $\leq \Delta$, and for each edge $e = uv \in E$, let $A^{uv} \in \mathbb{C}^{q \times q}$ be a (not necessarily symmetric or Hermitian) complex matrix. There exists a universal constant $\gamma \approx 0.56$ independent of $q, G, \{A^{uv}\}_{uv \in E}$ such that if $|A^{uv}(j, k) - 1| < \frac{\gamma}{\Delta + \gamma}$ for all $uv \in E$ and all $j, k \in [q]$, then for every $S \subseteq V$ and every $\phi : S \rightarrow [q]$, the graph homomorphism partition function*

$$\sum_{\substack{\sigma : V \rightarrow [q] \\ \sigma|_S = \phi}} \prod_{uv \in E} A^{uv}(\sigma(u), \sigma(v))$$

with pinning ϕ is nonzero.

We give below the proof of Theorem 5.1.

PROOF OF THEOREM 5.1. By Theorem 3.2, it suffices to prove that the multivariate partition function

$$\sum_{\substack{\sigma: V \rightarrow [q] \\ \sigma|_S = \phi}} \prod_{uv \in E} A^{uv}(\sigma(u), \sigma(v)) \prod_{v \in V} \lambda_{v, \sigma(v)} \quad (30)$$

is nonzero in the polydisk $\mathcal{D} = \{\lambda \in \mathbb{C}^{V \times [q]} : |\lambda_{v,k} - 1| < c, \forall v \in V, \forall k \in [q]\}$ for all pinnings ϕ , where $c = c(\Delta, \varepsilon) > 0$ is some constant depending only on Δ, ε but not G . Define a new set of matrices $\{\tilde{A}^{uv}\}_{uv \in E}$ by

$$\tilde{A}^{uv}(j, k) = A^{uv}(j, k) \cdot \lambda_{u,j}^{1/\deg(u)} \cdot \lambda_{v,k}^{1/\deg(v)}, \quad \forall uv \in E, \forall j, k \in [q].$$

Note that the partition function for $G, \{\tilde{A}^{uv}\}_{uv \in E}$ is precisely given in Equation (30).

Since $|A^{uv}(j, k) - 1| \leq \frac{\gamma}{\Delta + \gamma} - \varepsilon$, there exists our desired $c(\Delta, \varepsilon) > 0$ such that $|\lambda_{u,j} - 1|, |\lambda_{v,k} - 1| < c(\Delta, \varepsilon)$ implies $|\tilde{A}^{uv}(j, k) - 1| < \frac{\gamma}{\Delta + \gamma}$, for all $uv \in E$ and all $j, k \in [q]$. It follows from Theorem 5.3 that the multivariate partition function Equation (30) is nonzero. As this holds for all $\lambda \in \mathcal{D}$, we are done. ■

5.2 Tensor Network Contractions

Here, we study general tensor network contractions, which is a partition function of a distribution over configurations on edges of a graph. Tensor networks are heavily studied in quantum computing [44, 4, 47] and are also used to model Holant problems [19, 21, 20]. In the bounded-degree setting, we also show that the Glauber dynamics on edge configurations for these models mixes in $O(n \log n)$ steps, provided the weights are sufficiently close to 1. Again, this is analogous to classical mixing results stating the Glauber dynamics mixes rapidly in the “high-temperature” regime.

To state our main result, let us first define tensor network contraction. Given a graph $G = (V, E)$ and a collection of local functions $\{f_v : [q]^{E(v)} \rightarrow \mathbb{R}_{\geq 0}\}_{v \in V}$ on configurations on edges, we define the associated tensor network distribution μ over edge configurations $\sigma : E \rightarrow [q]$ to be given by

$$\mu(\sigma) \propto \prod_{v \in V} f_v(\sigma|_{E(v)}). \quad (31)$$

The associated partition function, known as a tensor network contraction, is given by

$$\sum_{\sigma: E \rightarrow [q]} \prod_{v \in V} f_v(\sigma|_{E(v)}).$$

The name “tensor network” comes from the fact that each f_v may be viewed as a tensor with axes corresponding to edges in $E(v)$ and indexed by $[q]$. This is a vast generalization of the Holant problems considered in Section 4 (see, for instance, Equation (28)), where $q = 2$ and

each local function f_v is symmetric. Zeros for tensor network contractions were analyzed in [50] in the symmetric case.

THEOREM 5.4 (Spectral Independence for Tensor Network Distribution). *Fix a positive integer $q \geq 2$, let $G = (V, E)$ be a graph with maximum degree $\leq \Delta$, and for each vertex $v \in V$, let $f_v : [q]^{E(v)} \rightarrow \mathbb{R}_{\geq 0}$ be a nonnegative function on configurations of edges incident to v . There exists a universal constant $\gamma \approx 0.56$ independent of $q, G, \{f_v\}_{v \in V}$ such that if $|f_v(\alpha) - 1| \leq \frac{\gamma}{\Delta+1+\gamma} - \varepsilon$ for some $\varepsilon > 0$, all $v \in V$ and all $\alpha : E(v) \rightarrow [q]$, then the tensor network distribution μ on edge configurations $\sigma : E \rightarrow [q]$ given by Equation (31) is η -spectrally independent for some constant $\eta = \eta(\Delta, \varepsilon)$. In particular, if $\Delta, \varepsilon = \Theta(1)$, then the Glauber dynamics for sampling from μ mixes in $O(n \log n)$ steps.*

REMARK 5.5. A straightforward application of the classical Dobrushin uniqueness condition yields rapid mixing when $|f_v(\alpha) - 1| \lesssim \frac{1}{4(\Delta-1)}$ for all $v \in V$ and all $\alpha : E(v) \rightarrow [q]$.

To establish this spectral independence, we need a sufficiently large zero-free region. This was proved by [50] in the symmetric case, where each local function f_v depends only on the number of incident edges that are mapped to each color in $[q]$. It turns out using nearly identical arguments, one can obtain the following more general theorem. We provide a proof in Appendix A.2 for completeness.

THEOREM 5.6 (Zeros of Tensor Network Contractions; [50]). *Fix a positive integer $q \geq 2$, let $G = (V, E)$ be a graph with maximum degree $\leq \Delta$, and for each vertex $v \in V$, let $f_v : [q]^{E(v)} \rightarrow \mathbb{C}$ be a complex function on configurations of edges incident to v . There exists a universal constant $\gamma \approx 0.56$ independent of $q, G, \{f_v\}_{v \in V}$ such that if $|f_v(\alpha) - 1| < \frac{\gamma}{\Delta+1+\gamma}$ for all $v \in V$ and all $\alpha : E(v) \rightarrow [q]$, then for every $F \subseteq E$ and every $\phi : F \rightarrow [q]$, the tensor network contraction*

$$\sum_{\substack{\sigma : E \rightarrow [q] \\ \sigma|_F = \phi}} \prod_{v \in V} f_v(\sigma|_{E(v)})$$

with pinning ϕ is nonzero.

We give below the proof of Theorem 5.4.

PROOF OF THEOREM 5.4. By Theorem 3.2, it suffices to prove that the multivariate partition function

$$\sum_{\substack{\sigma : E \rightarrow [q] \\ \sigma|_F = \phi}} \prod_{v \in V} f_v(\sigma|_{E(v)}) \prod_{e \in E} \lambda_{e, \sigma(e)} \quad (32)$$

is nonzero whenever λ lies in the polydisk $\mathcal{D} = \{\lambda \in \mathbb{C}^{E \times [q]} : |\lambda_{e,k} - 1| < c, \forall e \in E, \forall k \in [q]\}$ for all pinnings ϕ , where $c = c(\Delta, \varepsilon) > 0$ is some constant depending only on Δ, ε but not G . Define

a new set of local constraint functions $\{\tilde{f}_v\}_{v \in V}$ by

$$\tilde{f}_v(\alpha) = f_v(\alpha) \cdot \prod_{e \in E(v)} \lambda_{e, \alpha(e)}^{1/2}, \quad \forall v \in V, \forall \alpha : E(v) \rightarrow [q].$$

Note that the partition function for $G, \{\tilde{f}_v\}_{v \in V}$ is precisely given in Equation (32).

Since $|f_v(\alpha) - 1| \leq \frac{\gamma}{\Delta+1+\gamma} - \varepsilon$, there exists our desired $c(\Delta, \varepsilon) > 0$ such that $|\lambda_{e,k} - 1| < c(\Delta, \varepsilon)$ for all $e \in E(v)$ implies $|\tilde{f}_v(\alpha) - 1| < \frac{\gamma}{\Delta+1+\gamma}$, for all $v \in V$ and all $\alpha : E(v) \rightarrow [q]$. It follows from Theorem 5.6 that the multivariate partition function Equation (32) is nonzero. As this holds for all $\lambda \in \mathcal{D}$, we are done. $\blacksquare$

6. Arbitrary Measures on the Discrete Cube

In this section, we state a general result for mixing of an arbitrary measure on the discrete cube $\{-1, 1\}^n$. For this, we fix an arbitrary potential $f : \{-1, 1\}^n \rightarrow \mathbb{R}$. A standard result from analysis of Boolean functions says that f admits a unique representation as a multilinear polynomial $f(x) = \sum_{S \subseteq [n]} \hat{f}(S) \prod_{i \in S} x_i$. This representation is known as the Fourier-Walsh transform of f (see [46] and references therein), and the coefficients $\hat{f}(S)$ are known as the Fourier coefficients. [9] showed that when the Fourier coefficients, as well as the degree $\deg f$ of f as a multilinear polynomial, are sufficiently small, then one has a zero-free disk for the corresponding partition function $\sum_{x \in \{-1, 1\}^n} e^{f(x)}$. We convert this via Theorem 3.2 into a corresponding statement for the spectral independence of the distribution. Since we do not assume that f arises from a spin system (or, more generally, tensor network) on a bounded-degree graph, we only obtain a spectral gap bound with a relatively large exponent using Theorem 2.1 proved in [1, 3].

THEOREM 6.1. *Let $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ and $\varepsilon > 0$ be given, and assume that*

$$\sqrt{\deg f} \cdot L(f) \leq C - \varepsilon,$$

where $C \approx 0.55$ is an absolute constant, and $L(f) := \max_{i \in [n]} \sum_{S \subseteq [n] : S \ni i} |\hat{f}(S)|$. Further assume that the associated Gibbs distribution μ on $\{-1, 1\}^n$ given by

$$\mu(x) \propto \exp(f(x))$$

is b -marginally bounded for some $b > 0$. Let $\xi := \frac{2C}{\sqrt{\deg f}} - 2L(f)$. Then μ is η -spectrally independent where η is a constant depending only on b and ξ . In particular, μ is $O(1)$ -spectrally independent if $\varepsilon, \deg(f), b = \Theta(1)$.

REMARK 6.2. One may also view $L(f)$ as bounding the Lipschitz constant of f .

REMARK 6.3. A standard calculation using Dobrushin uniqueness condition yields that the Glauber dynamics is rapidly mixing when

$$\max_{i \in [n]} \sum_{S \subseteq [n]: S \ni i} (|S| - 1) \cdot |\hat{f}(S)| < 1,$$

which can be weakened to $(\deg(f) - 1) \cdot L(f) < 1$. These bounds are in general not comparable with the above due to the square root. While this bound is stronger when $\deg(f)$ is small, the above is stronger when most of the Fourier mass of f is on high-degree monomials.

REMARK 6.4. A standard notion in analysis of Boolean functions is also that of “influence”, which to avoid confusion with the notion of pairwise influence used to define spectral independence, we refer to as “voter influence”. This terminology is consistent with the traditional applications of analysis of Boolean functions to social choice theory and voting systems; see [46] and references therein. A standard result in analysis of Boolean functions says that the “voter influence” of coordinate i is precisely

$$\sum_{S \subseteq [n]: S \ni i} |\hat{f}(S)|^2.$$

Hence, while we do not establish a formal connection between small “voter influence” and strong spectral independence guarantees, our result Theorem 6.1 says this is true at least morally.

We need the following zero-free result from [9].

THEOREM 6.5 ([9]). *Let $f : \{-1, 1\}^n \rightarrow \mathbb{C}$ be given, and assume that*

$$\sqrt{\deg(f)} \cdot L(f) < C,$$

where $C \approx 0.55$ is an absolute constant, and $L(f) := \max_{i \in [n]} \sum_{S \subseteq [n]: S \ni i} |\hat{f}(S)|$. Then for every $S \subseteq [n]$ and every pinning $\phi : S \rightarrow \{-1, 1\}$, we have that the partition function of the associated Gibbs measure on $\{-1, 1\}^n$ with pinning ϕ is nonzero:

$$\sum_{x \in \{-1, 1\}^n: x|_S = \phi} \exp(f(x)) \neq 0.$$

We now prove Theorem 6.1.

PROOF OF THEOREM 6.1. By Theorem 3.2, it suffices to prove that the multivariate partition function

$$\sum_{x \in \{-1, 1\}^n: x|_S = \phi} \exp(f(x)) \prod_{i \in [n]: x_i = 1} \lambda_i \tag{33}$$

is nonzero whenever λ lies in the set $\mathcal{D} = \{\lambda \in \mathbb{C}^n : |\lambda_i - 1| < c, \forall i \in [n]\}$ for all pinnings ϕ , where $c = c(\xi) > 0$ is a constant depending only on ξ but not n . Define a new function

$g : \{-1, 1\}^n \rightarrow \mathbb{C}$ by

$$g(x) = f(x) + \sum_{i=1}^n \frac{1+x_i}{2} \log \lambda_i = f(x) + \sum_{i \in [n]: x_i=1} \log \lambda_i.$$

Then $\exp(g(x)) = \exp(f(x)) \prod_{i \in [n]: x_i=1} \lambda_i$ and the partition function $\sum_{x \in \{-1,1\}^n: x|_S=\phi} \exp(g(x))$ associated with g is precisely our desired multivariate partition function Equation (33). Our goal is to apply Theorem 6.5 to g and deduce our desired stability statement.

First, it is clear from the definition of g that the Fourier coefficients of g are given by

$$\hat{g}(S) = \begin{cases} \hat{f}(S), & \text{if } |S| > 1; \\ \hat{f}(i) + \frac{1}{2} \log \lambda_i, & \text{if } S = \{i\} \text{ for some } i \in [n]; \\ \hat{f}(\emptyset) + \frac{1}{2} \sum_{i=1}^n \log \lambda_i, & \text{if } S = \emptyset. \end{cases}$$

It follows that

$$L(g) \leq L(f) + \frac{1}{2} \max_{i \in [n]} |\log \lambda_i|.$$

Note that $\deg(g) = \deg(f)$ (unless $\deg(f) \leq 1$, in which case spectral independence and rapid mixing is trivial). Hence, if $\lambda \in \mathbb{C}^n$ satisfies $|\log \lambda_i| < \xi$ for all $i \in [n]$, then rearranging yields precisely that $\sqrt{\deg(g)} \cdot L(g) < C$ and the zero-freeness follows from Theorem 6.5. Furthermore, it is clear that the set $\{\lambda \in \mathbb{C}^n : |\log \lambda_i| < \xi, \forall i \in [n]\}$ contains $\mathcal{D}$ for a value of $c(\xi) > 0$ which depends only on ξ , just by continuity of the logarithm. ■

7. Future Directions

One open problem of our work is to obtain spectral independence for sampling b -matchings and b -edge-covers for general b on all graphs. It was shown in [35] that the Glauber dynamics for sampling b -matchings (respectively, b -edge-covers) mixes rapidly when $b \leq 7$ (respectively, $b \leq 2$); their proof approach relies on the canonical paths technique and in particular, the winding method [45]. In a recent paper [23], spectral independence was established for b -matchings and b -edge-covers for all b on bounded-degree graphs.

References

- [1] Vedat Levi Alev and Lap Chi Lau. Improved analysis of higher order random walks and applications. *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 1198–1211, 2020. DOI (11, 37)
- [2] Yeganeh Alimohammadi, Nima Anari, Kirankumar Shiragur, and Thuy-Duong Vuong. Fractionally log-concave and sector-stable polynomials: counting planar matchings and more. *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing (STOC)*, 2021. DOI (8, 14, 16, 17, 20)

- [3] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1319–1330, 2020. DOI (2, 6, 10, 11, 37)
- [4] Itai Arad and Zeph Landau. Quantum computation and the evaluation of tensor networks. *SIAM Journal on Computing*, 39(7):3089–3121, 2010. DOI (35)
- [5] Taro Asano. Theorems on the partition functions of the Heisenberg ferromagnets. *Journal of the Physical Society of Japan*, 29(2):350–359, 1970. DOI (9, 29)
- [6] Alexander Barvinok. Approximating permanents and hafnians. *Discrete Analysis*, 2, 2017. DOI (9, 42)
- [7] Alexander Barvinok. *Combinatorics and Complexity of Partition Functions*, volume 30. Springer Algorithms and Combinatorics, 2016. DOI (2, 9, 33, 42)
- [8] Alexander Barvinok. Computing the partition function for cliques in a graph. *Theory of Computing*, 11(13):339–355, 2015. DOI (42)
- [9] Alexander Barvinok. Computing the partition function of a polynomial on the Boolean cube. *A Journey Through Discrete Mathematics*:135–164, 2017. DOI (9, 37, 38)
- [10] Alexander Barvinok. Computing the permanent of (some) complex matrices. *Foundations of Computational Mathematics*, 16:329–342, 2016. DOI (9, 42)
- [11] Alexander Barvinok and Nicholas Barvinok. More on zeros and approximation of the Ising partition function. *Forum of Mathematics, Sigma*, 9:e46, 2021. DOI (42)
- [12] Alexander Barvinok and Anthony Della Pella. Testing for dense subsets in a graph via the partition function. *SIAM Journal on Discrete Mathematics*, 34:308–327, 1, 2020. DOI (42)
- [13] Alexander Barvinok and Pablo Soberón. Computing the partition function for graph homomorphisms. *Combinatorica*, 37:633–650, 2017. DOI (9, 33, 34, 42, 43)
- [14] Alexander Barvinok and Pablo Soberón. Computing the partition function for graph homomorphisms with multiplicities. *Journal of Combinatorial Theory, Series A*, 137:1–26, 2016. DOI (9, 33)
- [15] Ferenc Bencs, Péter Csikvári, and Guus Regts. Some applications of Wagner’s weighted subgraph counting polynomial. *The Electronic Journal of Combinatorics*, 28(4):P4.14, 2021. DOI (3, 5, 9, 29–32)
- [16] Ivona Bezáková and William A. Rummier. Sampling edge covers in 3-regular graphs. *Proceedings of the 34th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pages 137–148, 2009. DOI (3)
- [17] Antonio Blanca, Pietro Caputo, Zongchen Chen, Daniel Parisi, Daniel Štefankovič, and Eric Vigoda. On mixing of Markov chains: coupling, spectral independence, and entropy factorization. *Proceedings of the 33rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3670–3692, 2022. DOI (2, 3, 10, 11)
- [18] Julius Borcea and Petter Brändén. The Lee-Yang and Pólya-Schur programs. I. linear operators preserving stability. *Inventiones mathematicae*, 177(3):541, 2009. DOI (29)
- [19] Jin-Yi Cai, Sangxia Huang, and Pinyan Lu. From Holant to #CSP and back: dichotomy for Holant^c problems. *Algorithmica*, 64(3):511–533, 2012. DOI (35)
- [20] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Computational complexity of Holant problems. *SIAM Journal on Computing*, 40:1101–1132, 4, 2011. DOI (35)
- [21] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holant problems and counting CSP. *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 715–724, 2009. DOI (35)
- [22] Zongchen Chen, Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Rapid mixing for colorings via spectral independence. *Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1548–1557, 2021. DOI (2, 6)
- [23] Zongchen Chen and Yuzhou Gu. Fast sampling of b -matchings and b -edge covers. *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4972–4987, 2024. DOI (39)
- [24] Zongchen Chen, Kuikui Liu, and Eric Vigoda. Optimal mixing of Glauber dynamics: Entropy factorization via high-dimensional expansion. *SIAM Journal on Computing*:STOC21-104–153. DOI (2, 3, 10, 11, 29)
- [25] Zongchen Chen, Kuikui Liu, and Eric Vigoda. Rapid mixing of Glauber dynamics up to uniqueness via contraction. *SIAM Journal on Computing*, 52(1):196–237, 2023. DOI (2)
- [26] Zongchen Chen, Kuikui Liu, and Eric Vigoda. Spectral independence via stability and applications to holant-type problems. *62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 149–160. IEEE, 2021. DOI (1)
- [27] 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*:1–17, 2021. DOI (5, 29, 33)
- [28] Robert G Edwards and Alan D Sokal. Generalization of the Fortuin-Kasteleyn-Swendsen-Wang representation and Monte Carlo algorithm. *Physical review D*, 38(6):2009, 1988. DOI (4)
- [29] Weiming Feng, Heng Guo, and Jiaheng Wang. Swendsen-Wang dynamics for the ferromagnetic Ising model with external fields. *Information and Computation*, 294:105066, 2023. DOI (4)

- [30] Weiming Feng, Heng Guo, Yitong Yin, and Chihao Zhang. Rapid mixing from spectral independence beyond the Boolean domain. *ACM Transactions on Algorithms (TALG)*, 18(3):1–32, 2022. DOI (2, 6)
- [31] Cornelius Marius Fortuin and Piet W Kasteleyn. On the random-cluster model: I. Introduction and relation to other models. *Physica*, 57(4):536–564, 1972. DOI (4)
- [32] Antoine Gerschenfeld and Andrea Montanari. Reconstruction for models on random graphs. *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 194–204, 2007. DOI (8)
- [33] Geoffrey Grimmett and Svante Janson. Random even graphs. *The Electronic Journal of Combinatorics*, 16(1):R46, 2009. DOI (4)
- [34] 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. DOI (3, 9, 29, 30, 49)
- [35] Lingxiao Huang, Pinyan Lu, and Chihao Zhang. Canonical paths for MCMC: from art to science. *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 514–527, 2016. DOI (3, 29, 33, 39)
- [36] Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM Journal on Computing*, 18(6):1149–1178, 1989. DOI (5, 29, 33)
- [37] Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on Computing*, 22(5):1087–1116, 1993. DOI (4, 5, 29, 33)
- [38] Jacob Korevaar and Jan Wiegerinck. Several complex variables. Korteweg-de Vries Institute for Mathematics, 2017. URL (13)
- [39] Tsung-Dao Lee and Chen-Ning Yang. Statistical theory of equations of state and phase transitions. II. Lattice gas and Ising model. *Physical Review*, 87(3):410–419, 1952. DOI (3)
- [40] Chengyu Lin, Jingcheng Liu, and Pinyan Lu. A simple FPTAS for counting edge covers. *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 341–348, 2014. DOI (3)
- [41] Jingcheng Liu, Pinyan Lu, and Chihao Zhang. FPTAS for counting weighted edge covers. *Proceedings of the 22nd Annual European Symposium on Algorithms (ESA)*, pages 654–665, 2014. DOI (3)
- [42] Jingcheng Liu, Alistair Sinclair, and Piyush Srivastava. The Ising partition function: zeros and deterministic approximation. *Journal of Statistical Physics*, 174(2):287–315, 2019. DOI (4)
- [43] Kuikui Liu. From coupling to spectral independence and blackbox comparison with the down-up walk. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2021)*, volume 207, 32:1–32:21, 2021. DOI (2)
- [44] Igor L. Markov and Yaoyun Shi. Simulating quantum computation by contracting tensor networks. *SIAM Journal on Computing*, 38:963–981, 3, 2008. DOI (35)
- [45] Colin McQuillan. Approximating Holant problems by winding. *arXiv preprint arXiv:1301.2880*, 2013. DOI (29, 39)
- [46] Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, USA, 2014. DOI (37, 38)
- [47] Román Orús. A practical introduction to tensor networks: matrix product states and projected entangled pair states. *Annals of Physics*, 349:117–158, 2014. DOI (35)
- [48] 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. DOI (2, 9, 33)
- [49] Dana Randall and David Wilson. Sampling spin configurations of an Ising system. *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 959–960, 1999. URL (4)
- [50] Guus Regts. Zero-free regions of partition functions with applications to algorithms and graph limits. *Combinatorica*, 38:987–1015, 2018. DOI (9, 33, 36, 44–46)
- [51] Walter Rudin. *Real and Complex Analysis*. Mathematics series. McGraw-Hill, 1987. (12)
- [52] David Ruelle. Extension of the Lee–Yang circle theorem. *Physical Review Letters*, 26(6):303–304, 1971. DOI (9, 29)
- [53] David G. Wagner. Weighted enumeration of spanning subgraphs with degree constraints. *Journal of Combinatorial Theory, Series B*, 99(2):347–357, 2009. DOI (9, 29)
- [54] Dror Weitz. Counting independent sets up to the tree threshold. *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 140–149, 2006. DOI (2)

A. Proofs of Zero-Free Results

In this section, we supply proofs of the main zero-free statements used in Section 5. As noted earlier, for technical reasons, we need straightforward generalizations of prior results which do not make symmetry assumptions. We manage to adapt previous arguments without much additional effort, which we provide here for completeness.

The main idea in these zero-free proofs is to do induction by conditioning on the assignment of fewer and fewer vertices (respectively, edges) for weighted homomorphisms (respectively, tensor networks). However, one needs to strengthen the inductive hypothesis beyond simple zero-freeness. To the best of our knowledge, this type of argument was first pioneered by Barvinok, and has had a wide range of applications; see [8, 6, 12, 11] for applications besides those discussed in this paper.

The crucial tool is the following geometric lemma, which provides a kind of “reverse triangle inequality”. The version below is due to Boris Bukh; a weaker version, with $\cos(\theta/2)$ replaced by $\sqrt{\cos \theta}$, was known due to [10]. See [7] for a proof.

LEMMA A.1 (Angle Lemma). *Let $x_1, \dots, x_n \in \mathbb{C}$ be nonzero complex numbers viewed as vectors in $\mathbb{R}^2$. Suppose there is an angle $0 \leq \theta < 2\pi/3$ such that for all i, j , the angle between x_i, x_j is at most θ . Then we have the lower bound $|\sum_{i=1}^n x_i| \geq \cos(\theta/2) \sum_{i=1}^n |x_i|$.*

A.1 Proofs for Weighted Graph Homomorphisms

Our goal in this subsection is to prove Theorem 5.3, i.e., that the weighted graph homomorphism partition function

$$Z_\phi^S(A) = \sum_{\substack{\sigma: V \rightarrow [q] \\ \sigma|_S = \phi}} \prod_{uv \in E} A^{uv}(\sigma(u), \sigma(v))$$

is nonzero in a large polydisk around 1, where $S \subseteq V$, $\phi : S \rightarrow [q]$, and we view $Z_\phi^S(A)$ as a polynomial with variables $\{A^{uv}(j, k)\}_{uv \in E, j, k \in [q]}$. For convenience, for a $\delta > 0$, define

$$\mathcal{U}(\delta) = \{A = \{A^{uv}\}_{uv \in E} : |A^{uv}(j, k) - 1| < \delta, \forall uv \in E, \forall j, k \in [q]\}.$$

Additionally, for a partial configuration $\phi : S \rightarrow [q]$, a vertex $u \in V \setminus S$ and a spin $j \in [q]$, we write $\phi_{u,j} : S \cup \{u\} \rightarrow [q]$ for the unique extension of ϕ with $\phi_{u,j}(u) = j$.

We will need the following lemmas to implement an inductive approach.

LEMMA A.2 (Lemma 3.3 from [13]). *Let $\tau, \delta > 0$, and suppose $A \in \mathcal{U}(\delta)$. Let $S \subseteq V$, $\phi : S \rightarrow [q]$, $u \in V \setminus S$ be arbitrary. Assume the following hold:*

- (1) $Z_{\phi_{u,j}}^{S \cup \{u\}}(A) \neq 0$ for every $u \in V \setminus S$ and every $j \in [q]$;

(2) For every $u \in V \setminus S$ and every $j \in [q]$, we have

$$|Z_{\phi_{u,j}}^{S \cup \{u\}}(A)| \geq \frac{\tau}{\Delta} \sum_{v \sim u} \sum_{k \in [q]} |A^{uv}(j, k)| \cdot \left| \frac{\partial}{\partial A^{uv}(j, k)} Z_{\phi_{u,j}}^{S \cup \{u\}}(A) \right|.$$

Then for every $u \in V \setminus S$ and every $j, k \in [q]$, the angle between $Z_{\phi_{u,j}}^{S \cup \{u\}}(A)$ and $Z_{\phi_{u,k}}^{S \cup \{u\}}(A)$ in $\mathbb{C}$ is at most $\frac{2\delta\Delta}{\tau(1-\delta)}$.

PROOF. By assumption (1), the relevant partition functions are nonzero, and so the logarithm is well-defined when applied to these partition functions and we may bound the angle between $Z_{\phi_{u,j}}^{S \cup \{u\}}(A)$ and $Z_{\phi_{u,k}}^{S \cup \{u\}}(A)$ by

$$\left| \log Z_{\phi_{u,j}}^{S \cup \{u\}}(A) - \log Z_{\phi_{u,k}}^{S \cup \{u\}}(A) \right|. \quad (34)$$

The strategy is to write $Z_{\phi_{u,k}}^{S \cup \{u\}}(A)$ as $Z_{\phi_{u,j}}^{S \cup \{u\}}(\tilde{A})$ for some $\tilde{A} \in \mathcal{U}(\delta)$ which differs from A by a small number of coordinates, and then apply the Fundamental Theorem of Calculus and assumption (2). For every $v \sim u$, we set $\tilde{A}^{uv}(j, c) = A^{uv}(k, c)$ for every $c \in [q]$, and $\tilde{A}^{uv}(\ell, c) = A^{uv}(\ell, c)$ for all $\ell \neq j$. For all other edges $vw \in E$, we set $\tilde{A}^{vw} = A^{vw}$.

It is clear that $Z_{\phi_{u,k}}^{S \cup \{u\}}(A) = Z_{\phi_{u,j}}^{S \cup \{u\}}(\tilde{A})$. By the Fundamental Theorem of Calculus, we may upper bound Equation (34) by

$$\begin{aligned} & \max_{B \in \mathcal{U}(\delta)} \sum_{v \sim u} \sum_{c \in [q]} \left| \frac{\partial}{\partial A^{uv}(j, c)} \log Z_{\phi_{u,j}}^{S \cup \{u\}}(B) \right| \cdot \underbrace{|A^{uv}(j, c) - \tilde{A}^{uv}(j, c)|}_{\leq 2\delta \text{ since } A, \tilde{A} \in \mathcal{U}(\delta)} \\ & \leq \frac{2\delta}{1-\delta} \max_{B \in \mathcal{U}(\delta)} \underbrace{\sum_{v \sim u} \sum_{c \in [q]} |A^{uv}(j, c)| \cdot \frac{1}{|Z_{\phi_{u,j}}^{S \cup \{u\}}(B)|}}_{\leq \Delta/\tau \text{ by assumption (2)}} \cdot \left| \frac{\partial}{\partial A^{uv}(j, c)} Z_{\phi_{u,j}}^{S \cup \{u\}}(B) \right| \\ & \leq \frac{2\delta\Delta}{\tau(1-\delta)}. \end{aligned} \quad \blacksquare$$

LEMMA A.3 (Lemma 3.4 from [13]). Let $0 \leq \theta < 2\pi/3$, $\delta > 0$, and suppose $A \in \mathcal{U}(\delta)$. Let $S \subseteq V$, $\phi : S \rightarrow [q]$ be arbitrary. Assume the following hold:

- (1) $Z_{\phi_{u,j}}^{S \cup \{u\}}(A) \neq 0$ for every $u \in V \setminus S$ and every $j \in [q]$;
- (2) The angle between $Z_{\phi_{u,j}}^{S \cup \{u\}}(A)$ and $Z_{\phi_{u,k}}^{S \cup \{u\}}(A)$ in $\mathbb{C}$ is at most θ , for every $u \in V \setminus S$ and every $j, k \in [q]$.

Then for every $u \in S$, we have the lower bound

$$|Z_{\phi}^S(A)| \geq \frac{\cos(\theta/2)}{\Delta} \sum_{v \sim u} \sum_{k \in [q]} |A^{uv}(\phi(u), k)| \cdot \left| \frac{\partial}{\partial A^{uv}(\phi(u), k)} Z_{\phi}^S(A) \right|.$$

PROOF. If $v \in S$ as well, then there is a unique $k \in [q]$ for which $\frac{\partial}{\partial A^{uv}(\phi(u), k)} Z_{\phi}^S(A) \neq 0$, namely $k = \phi(v)$. In this case, $A^{uv}(\phi(u), k) \cdot \frac{\partial}{\partial A^{uv}(\phi(u), k)} Z_{\phi}^S(A) = Z_{\phi}^S(A)$. Otherwise, $v \notin S$ and $\frac{\partial}{\partial A^{uv}(\phi(u), k)} Z_{\phi}^S(A) = \frac{1}{A^{uv}(\phi(u), k)} \cdot Z_{\phi_{v,k}}^{S \cup \{v\}}(A)$, where $\phi_{v,k}$ is the unique extension of ϕ mapping v to k .

Combining these two observations, we obtain

$$\begin{aligned}
& \sum_{v \sim u} \sum_{k \in [q]} |A^{uv}(\phi(u), k)| \cdot \left| \frac{\partial}{\partial A^{uv}(\phi(u), k)} Z_{\phi}^S(A) \right| \\
&= |N(u) \cap S| \cdot |Z_{\phi}^S(A)| + \sum_{v \sim u: v \notin S} \sum_{k \in [q]} |Z_{\phi_{v,k}}^{S \cup \{v\}}(A)| \\
&\leq |N(u) \cap S| \cdot |Z_{\phi}^S(A)| + \frac{1}{\cos(\theta/2)} \underbrace{\left| \sum_{v \sim u: v \notin S} \sum_{k \in [q]} Z_{\phi_{v,k}}^{S \cup \{v\}}(A) \right|}_{=|N(u) \setminus S| \cdot |Z_{\phi}^S(A)|} \quad (\text{Lemma A.1}) \\
&\leq \frac{\Delta}{\cos(\theta/2)} \cdot |Z_{\phi}^S(A)|.
\end{aligned}$$

Rearranging yields the desired result. ■

With these lemmas in hand, we can now prove the main zero-free result.

PROOF OF THEOREM 5.3. Let $0 < \theta < 2\pi/3$ be a parameter to be determined later, set $\tau = \cos(\theta/2)$, and let $\delta > 0$ satisfy $\theta = \frac{2\delta\Delta}{\tau(1-\delta)}$; in particular, $\delta = \frac{\frac{1}{2\Delta}\theta \cos(\theta/2)}{1 + \frac{1}{2\Delta}\theta \cos(\theta/2)}$. We show by descending induction on $|S|$ that the following three statements are all true:

- (i) For every $S \subseteq V$, $\phi : S \rightarrow [q]$ and $A \in \mathcal{U}(\delta)$, we have $Z_{\phi}^S(A) \neq 0$.
- (ii) For every $S \subseteq V$, $u \in V \setminus S$, $\phi : S \rightarrow [q]$, $A \in \mathcal{U}(\delta)$ and $j, k \in [q]$, the angle between $Z_{\phi_{u,j}}^{S \cup \{u\}}(A)$ and $Z_{\phi_{u,k}}^{S \cup \{u\}}(A)$ in $\mathbb{C}$ is at most θ .
- (iii) For every $S \subseteq V$, $u \in S$, $A \in \mathcal{U}(\delta)$, we have the inequality

$$|Z_{\phi}^S(A)| \geq \frac{\cos(\theta/2)}{\Delta} \sum_{v \sim u} \sum_{k \in [q]} |A^{uv}(\phi(u), k)| \cdot \left| \frac{\partial}{\partial A^{uv}(\phi(u), k)} Z_{\phi}^S(A) \right|.$$

The base case $S = V$ is easily verified since $Z_{\phi}^S(A) = \prod_{uv \in E} A^{uv}(\phi(u), \phi(v))$, a product of nonzero complex numbers.

Now, let $S \subseteq V$ with $|S| < |V|$.

Proof of (i) Let $u \in V \setminus S$, which exists since $|S| < |V|$. It follows that (i) holds for $S \cup \{u\}$ by the inductive hypothesis. Since $Z_{\phi}^S(A) = \sum_{k \in [q]} Z_{\phi_{u,k}}^{S \cup \{u\}}(A)$, Lemma A.1 applied to $Z_{\phi}^S(A)$ yields (i) assuming that (ii) holds. We prove (ii) below.

Proof of (ii) Let $u \in V \setminus S$, which exists since $|S| < |V|$. Then (i) and (iii) hold for $S \cup \{u\}$ by the inductive hypothesis. (ii) then follows by Lemma A.2.

Proof of (iii) Let $u \in S$. Then (i) holds for $S \cup \{u\}$ by the inductive hypothesis. Since (ii) holds for S (as proved earlier), we may then apply Lemma A.3, yielding (iii) for S .

Now, we choose $0 < \theta < 2\pi/3$. As we wish to maximize the size of our zero-free region, namely δ , we need to maximize $\theta \cos(\theta/2)$. As shown in [50], the maximum is attained when $2/\theta =$

$\tan(\theta/2)$, which has solution $\theta^* \approx 1.72067$ and has objective value $x^* = \theta^* \cos(\theta^*/2) \approx 1.12219$. This yields $\delta = \frac{\frac{x^*}{2}}{\Delta + \frac{x^*}{2}}$ as claimed. $\blacksquare$

A.2 Proofs for Tensor Network Contractions

Our goal in this subsection is to prove Theorem 5.6, i.e., that the tensor network partition function

$$Z_\phi^F(h) = \sum_{\substack{\sigma: E \rightarrow [q] \\ \sigma|_F = \phi}} \prod_{v \in V} h_v(\sigma|_{E(v)})$$

is nonzero in a large polydisk around 1, where $F \subseteq E$, $\phi : F \rightarrow [q]$, and we view $Z_\phi^F(\cdot)$ as a polynomial with variables $\{h_v(\alpha)\}_{v, \alpha}$. We prove the following stronger result.

THEOREM A.4 (Generalization of Theorem 6 from [50]). *Let $G = (V, E)$ be a graph of maximum degree $\leq \Delta$. Then for every $F \subseteq E$, $\phi : F \rightarrow [q]$, $\eta > 0$, and $0 \leq \theta < 2\pi/3$, the function $Z_\phi^F(h)$ is nonzero whenever $h \in \prod_{v \in V} S_v(\delta, \eta)$, where*

$$S_v(\delta, \eta) = \left\{ h_v : [q]^{E(v)} \rightarrow \mathbb{C} : \begin{array}{l} |h_v(\alpha) - h_v(\beta)| < \delta, \forall \alpha, \beta : E(v) \rightarrow [q], \\ |h_v(\alpha)| \geq \eta, \forall \alpha : E(v) \rightarrow [q] \end{array} \right\}$$

$$\text{and } \delta = \eta \cdot \min \left\{ 1, \frac{\theta \cos(\theta/2)}{\Delta+1} \right\}.$$

Before we prove this result, let us see how this gives Theorem 5.6.

PROOF OF THEOREM 5.6. Observe that $S_v(\delta, \eta)$ contains a disk around 1 of radius $\min\{\delta/2, 1 - \eta\}$. Using Theorem A.4 and given that $\delta = \eta \cdot \min \left\{ 1, \frac{\theta \cos(\theta/2)}{\Delta+1} \right\}$, where $0 < \theta < 2\pi/3$, our goal is to maximize $\theta \cos(\theta/2)$ over $0 < \theta < 2\pi/3$ to obtain the largest zero-free disk. As shown in [50], this maximum is attained when $2/\theta = \tan(\theta/2)$, which has solution $\theta^* \approx 1.72067$ and has objective value $x^* = \theta^* \cos(\theta^*/2) \approx 1.12219$. Given this, to obtain the largest possible radius disk, we equalize $1 - \eta$ and $\delta/2 = \eta \cdot \frac{x^*}{2(\Delta+1)}$. Solving, we obtain $\eta = \frac{1}{1 + \frac{x^*}{2(\Delta+1)}}$, yielding radius $\frac{\frac{x^*}{2(\Delta+1)}}{1 + \frac{x^*}{2(\Delta+1)}}$ as desired. $\blacksquare$

It remains to prove Theorem A.4. We will need the following lemmas to implement an inductive approach.

LEMMA A.5 (Lemma 8 from [50]). *Let $\tau > 0$, $F \subseteq E$, $\phi : F \rightarrow [q]$ and $u \in V$ be arbitrary. Suppose for all $h \in \prod_{v \in V} S_v(\delta, \eta)$ and all $\psi : F \cup E(u) \rightarrow [q]$ extending ϕ , the following hold:*

- (1) $Z_\psi^{F \cup E(u)}(h) \neq 0$;
- (2) For all $v \in N(u) \cup \{u\}$, we have

$$|Z_\psi^{F \cup E(u)}(h)| \geq \tau \sum_{\substack{\alpha: E(v) \rightarrow [q] \\ \text{compatible with } \psi}} |h_v(\alpha)| \cdot \left| \frac{\partial}{\partial h_v(\alpha)} Z_\psi^{F \cup E(u)}(h) \right|.$$

Then for all extensions $\psi, \tilde{\psi} : F \cup E(u) \rightarrow \mathbb{C}$ of ϕ , the angle between $Z_{\psi}^{F \cup E(u)}$ and $Z_{\tilde{\psi}}^{F \cup E(u)}(h)$ in $\mathbb{C}$ is at most $\frac{\delta(\Delta+1)}{\tau\eta}$.

PROOF. By assumption (1), the relevant partition functions are nonzero, and so the logarithm is well-defined when applied to these partition functions and we may bound the angle between $Z_{\psi}^{F \cup E(u)}(h)$ and $Z_{\tilde{\psi}}^{F \cup E(u)}(h)$ by

$$\left| \log Z_{\psi}^{F \cup E(u)}(h) - \log Z_{\tilde{\psi}}^{F \cup E(u)}(h) \right|. \quad (35)$$

The strategy is to write $Z_{\tilde{\psi}}^{F \cup E(u)}(h)$ as $Z_{\psi}^{F \cup E(u)}(\tilde{h})$ for some $\tilde{h} \in \prod_{v \in V} S_v(\delta, \eta)$ which differs from h by a small number of coordinates, and then apply the Fundamental Theorem of Calculus and assumption (2). Let $v \in V$. We consider three cases.

- $v \notin N(u) \cup \{u\}$: In this case, $\psi, \tilde{\psi}$ agree on $E(v)$ and so we may simply take $h_v = \tilde{h}_v$.
- $v \in N(u)$: In this case, $\psi, \tilde{\psi}$ differ only on the single edge uv . If $\alpha : E(v) \rightarrow [q]$ agrees with ψ on uv , then let $\alpha' : E(v) \rightarrow [q]$ be given by replacing $\alpha(uv) = \psi(uv)$ with $\tilde{\psi}(uv)$, and take $\tilde{h}_v(\alpha) = h_v(\alpha')$. Otherwise, just set $\tilde{h}_v(\alpha) = h_v(\alpha)$. (Note that it does not really matter what we set $\tilde{h}_v(\alpha)$ to since $Z_{\psi}^{F \cup E(u)}(h)$ only has the term $h_v(\alpha)$ when α agrees with ψ on uv . However, we wish to minimize the number of coordinates in which $h, \tilde{h}$ differ.)
- $v = u$: In this case, just set $\tilde{h}_v(\psi|_{E(v)}) = h_v(\tilde{\psi}|_{E(v)})$ and $\tilde{h}_v(\alpha) = h_v(\alpha)$ for all $\alpha \neq \psi|_{E(v)}$.

It is clear that $Z_{\tilde{\psi}}^{F \cup E(u)}(h) = Z_{\psi}^{F \cup E(u)}(\tilde{h})$. By the Fundamental Theorem of Calculus, we may upper bound Equation (35) by

$$\begin{aligned} & \max_{x \in \prod_{v \in V} S_v(\delta, \eta)} \sum_{v \in N(u) \cup \{u\}} \sum_{\substack{\alpha : E(v) \rightarrow [q] \\ \text{compatible with } \psi}} \left| \frac{\partial}{\partial h_v(\alpha)} \log Z_{\psi}^{F \cup E(u)}(x) \right| \cdot |h_v(\alpha) - \tilde{h}_v(\alpha)| \\ & \leq \frac{\delta}{\eta} \max_{x \in \prod_{v \in V} S_v(\delta, \eta)} \underbrace{\sum_{v \in N(u) \cup \{u\}} \sum_{\substack{\alpha : E(v) \rightarrow [q] \\ \text{compatible with } \psi}}}_{\leq \Delta+1} |h_v(\alpha)| \cdot \underbrace{\frac{1}{|Z_{\psi}^{F \cup E(u)}(x)|} \cdot \left| \frac{\partial}{\partial h_v(\alpha)} Z_{\psi}^{F \cup E(u)}(x) \right|}_{\leq 1/\tau \text{ by assumption (2)}} \\ & \hspace{20em} (\text{Definition of } S_v(\delta, \eta)) \\ & \leq \frac{\delta(\Delta+1)}{\tau\eta}. \end{aligned} \quad \blacksquare$$

LEMMA A.6 (Lemma 9 from [50]). Let $0 \leq \theta < 2\pi/3$, $u \in V$, $F \subseteq E$ satisfying $F \supseteq E(u)$, and $\phi : F \rightarrow [q]$. Suppose for all $v \in N(u) \cup \{u\}$, all $h \in \prod_{v \in V} S_v(\delta, \eta)$, and all extensions $\psi, \tilde{\psi} : F \cup E(v) \rightarrow [q]$ of ϕ , the following hold:

- (1) $Z_{\psi}^{F \cup E(v)}(h) \neq 0$;
- (2) The angle between $Z_{\psi}^{F \cup E(v)}(h)$ and $Z_{\tilde{\psi}}^{F \cup E(v)}(h)$ in $\mathbb{C}$ is at most θ .

Then for all $v \in N(u) \cup \{u\}$ and all $h \in \prod_{v \in V} S_v(\delta, \eta)$, we have

$$|Z_\phi^F(h)| \geq \cos(\theta/2) \sum_{\substack{\alpha: E(v) \rightarrow [q] \\ \text{compatible with } \phi}} |h_v(\alpha)| \cdot \left| \frac{\partial}{\partial h_v(\alpha)} Z_\phi^F(h) \right|.$$

PROOF. The conclusion is trivially true if $v = u$, since by the assumption $E(u) \subseteq F$, there is only one $\alpha : E(v) \rightarrow [q]$ compatible with ϕ , namely $\phi|_{E(u)}$ itself. In this case, $h_v(\alpha)$ divides $Z_\phi^F(h)$ and we can replace $\cos(\theta/2)$ by 1.

Suppose $v \in N(u)$. Since $Z_\phi^F(h) = \sum_{\substack{\psi: F \cup E(v) \rightarrow [q] \\ \psi|_F = \phi}} Z_\psi^{F \cup E(v)}(h)$, assumptions (1) and (2) make Lemma A.1 applicable, yielding

$$\begin{aligned} |Z_\phi^F(h)| &\geq \cos(\theta/2) \sum_{\substack{\psi: F \cup E(v) \rightarrow [q] \\ \psi|_F = \phi}} |Z_\psi^{F \cup E(v)}(h)| \\ &= \cos(\theta/2) \sum_{\substack{\alpha: E(v) \rightarrow [q] \\ \text{compatible with } \psi}} |h_v(\alpha)| \cdot \left| \frac{\partial}{\partial h_v(\alpha)} Z_\psi^{F \cup E(v)}(h) \right| \end{aligned}$$

as desired. ■

With these lemmas in hand, we may now proceed with the proof of Theorem A.4.

PROOF OF THEOREM A.4. Let $\eta > 0$ and $0 \leq \theta < 2\pi/3$ be arbitrary, and take $\tau = \cos(\theta/2)$, $\delta = \eta \cdot \min \{1, \frac{\theta\tau}{\Delta+1}\}$. We show by descending induction on $|F|$ that the following three statements are all true:

- (i) For every $F \subseteq E$, $\phi : F \rightarrow [q]$ and $h \in \prod_{v \in V} S_v(\delta, \eta)$, we have $Z_\phi^F(h) \neq 0$.
- (ii) For every $F \subseteq E$, $u \in V$, $\phi : F \rightarrow [q]$, $h \in \prod_{v \in V} S_v(\delta, \eta)$ and $\psi, \tilde{\psi} : F \cup E(u) \rightarrow [q]$ extending ϕ , the angle between $Z_\psi^{F \cup E(u)}(h)$ and $Z_{\tilde{\psi}}^{F \cup E(u)}(h)$ in $\mathbb{C}$ is at most θ .
- (iii) For every $F \subseteq E$, $u \in V$ satisfying $E(u) \subseteq F$, $\phi : F \rightarrow [q]$, $h \in \prod_{v \in V} S_v(\delta, \eta)$ and $v \in N(u) \cup \{u\}$, we have the inequality

$$|Z_\phi^F(h)| \geq \cos(\theta/2) \sum_{\substack{\alpha: E(v) \rightarrow [q] \\ \text{compatible with } \phi}} |h_v(\alpha)| \cdot \left| \frac{\partial}{\partial h_v(\alpha)} Z_\phi^F(h) \right|.$$

The base case $F = E$ is easily verified since $Z_\phi^F(h) = \prod_{v \in V} h_v(\phi|_{E(v)})$, a product of nonzero complex numbers.

Now, let $F \subseteq E$ with $|F| < |E|$.

Proof of (i) Let $v \in V$ with $E(v) \not\subseteq F$. Since $|F \cup E(v)| > |F|$, (i) holds for $F \cup E(v)$ by the inductive hypothesis. Since $Z_\phi^F(h) = \sum_{\substack{\psi: F \cup E(v) \rightarrow [q] \\ \psi|_F = \phi}} Z_\psi^{F \cup E(v)}(h)$, Lemma A.1 applied to $Z_\psi^{F \cup E(v)}(h)$ yields (i) assuming that (ii) holds. We prove (ii) below.

Proof of (ii) Let $u \in V$ and $\phi : F \rightarrow [q]$. If $E(u) \subseteq F$, then the claim is trivially true since $\psi = \tilde{\psi} = \phi$. Otherwise, assume $E(u) \not\subseteq F$ and let $\psi, \tilde{\psi} : F \cup E(u) \rightarrow [q]$ extend ϕ . Since

$|F \cup E(u)| > |F|$, (i) and (iii) hold for $F \cup E(u)$ by the inductive hypothesis. Applying Lemma A.5 to $F \cup E(u)$ then yields (ii).

Proof of (iii) Let $u \in V$ with $E(u) \subseteq F$. Without loss of generality, we may assume such an u exists since otherwise, there is nothing to prove. Let $v \in N(u) \cup \{u\}$. If $E(v) \subseteq F$, then (iii) trivially holds with $\cos(\theta/2)$ replaced by 1, since there is only one term in the summation, namely $\alpha = \phi|_{E(v)}$. Hence, assume $E(v) \not\subseteq F$. In this case, $|F \cup E(v)| > |F|$ and so (i) holds for $F \cup E(v)$ by the inductive hypothesis. Since (ii) for F holds (as proved earlier), we may then apply Lemma A.6, yielding (iii) for F . ■

B. Proofs of Technical Lemmas

PROOF OF LEMMA 3.5. By definition we have

$$\text{dist}\left(1, C_{v,k}^\tau\right) = \inf_{1 \neq z \in \Gamma_{v,k}} \left| -\frac{1}{p_{v,k}^\tau(z-1)} - 1 \right|.$$

If $\Gamma_{v,k}$ is unbounded, then there exists a sequence $\{z_n\}$ such that $1 \neq z_n \in \Gamma_{v,k}$ and $\lim_{n \rightarrow \infty} |z_n| = \infty$. Therefore,

$$\text{dist}\left(1, C_{v,k}^\tau\right) \leq \liminf_{n \rightarrow \infty} \left| -\frac{1}{p_{v,k}^\tau(z_n-1)} - 1 \right| \leq 1 + \liminf_{n \rightarrow \infty} \frac{1}{p_{v,k}^\tau|z_n-1|} = 1.$$

This shows the first part.

For the second part, observe that $\alpha_{v,k} < 1 < \beta_{v,k}$ since $\Gamma_{v,k}$ is open and $1 \in \Gamma_{v,k}$. Hence, we obtain

$$\text{dist}\left(1, C_{v,k}^\tau\right) \leq \inf_{x \in \Gamma_{v,k} \cap (0,1)} \left| -\frac{1}{p_{v,k}^\tau(x-1)} - 1 \right| = \frac{1}{p_{v,k}^\tau(1-\alpha_{v,k})} - 1 = \frac{\alpha_{v,k}}{p_{v,k}^\tau(1-\alpha_{v,k})} + \frac{1-p_{v,k}^\tau}{p_{v,k}^\tau},$$

and also

$$\text{dist}\left(1, C_{v,k}^\tau\right) \leq \inf_{x \in \Gamma_{v,k} \cap (1,\infty)} \left| -\frac{1}{p_{v,k}^\tau(x-1)} - 1 \right| = \frac{1}{p_{v,k}^\tau(\beta_{v,k}-1)} + 1.$$

The second part follows. ■

PROOF OF LEMMA 3.6. By definition we have

$$\text{dist}\left(1, C_{v,0}^\tau\right) = \inf_{1 \neq z \in \Gamma_v} \left| \frac{z}{p_{v,0}^\tau(z-1)} - 1 \right|.$$

If $0 \in \overline{\Gamma_v}$, then there exists a sequence $\{z_n\}$ such that $1 \neq z_n \in \Gamma_{v,k}$ and $\lim_{n \rightarrow \infty} z_n = 0$. Therefore,

$$\text{dist}\left(1, C_{v,0}^\tau\right) \leq \liminf_{n \rightarrow \infty} \left| \frac{z_n}{p_{v,0}^\tau(z_n-1)} - 1 \right| \leq 1 + \liminf_{n \rightarrow \infty} \frac{|z_n|}{p_{v,0}^\tau|z_n-1|} = 1.$$

This shows the first part.

For the second part, observe that $\alpha_v < 1 < \beta_v$ since Γ_v is open and $1 \in \Gamma_{v,k}$. Hence, we obtain

$$\text{dist}\left(1, C_{v,0}^\tau\right) \leq \inf_{x \in \Gamma_v \cap (0,1)} \left| \frac{x}{p_{v,0}^\tau(x-1)} - 1 \right| = \frac{\alpha_v}{p_{v,0}^\tau(1-\alpha_v)} + 1,$$

and also

$$\text{dist}\left(1, C_{v,0}^\tau\right) \leq \inf_{x \in \Gamma_v \cap (1,\infty)} \left| \frac{x}{p_{v,0}^\tau(x-1)} - 1 \right| = \frac{\beta_v}{p_{v,0}^\tau(\beta_v-1)} - 1 = \frac{1}{p_{v,0}^\tau(\beta_v-1)} + \frac{1-p_{v,0}^\tau}{p_{v,0}^\tau}.$$

The second part follows. ■

PROOF OF LEMMA 4.2. It was shown in [34] that

$$\Gamma = \left(-\overline{\mathbb{H}}_\varepsilon^2\right)^c = \{\rho e^{i\theta} : \rho(1 - \cos \theta) < 2\varepsilon^2, 0 \leq \theta < 2\pi\}.$$

To make this more interpretable, we rewrite the set in Cartesian coordinates. If $z = \rho e^{i\theta}$, then by Euler's formula we may write $z = x + iy$ where $x = \rho \cos \theta$ and $y = \rho \sin \theta$. We then obtain

$$\begin{aligned} & \rho(1 - \cos \theta) < 2\varepsilon^2 \\ \iff & \rho < x + 2\varepsilon^2 \\ \iff & x^2 + y^2 < (x + 2\varepsilon^2)^2 \\ \iff & y^2 < 4\varepsilon^2(x + \varepsilon^2). \end{aligned}$$

Therefore, we see that

$$\Gamma = \{x + iy : y^2 < 4\varepsilon^2(x + \varepsilon^2)\},$$

which clearly contains $\mathbb{R}_+$.

Furthermore, for $\lambda \in \mathbb{R}_+$ we have

$$\begin{aligned} \text{dist}(\lambda, \partial\Gamma) &= \inf_{z \in \partial\Gamma} |z - \lambda| \\ &= \inf_{(x,y) \in \mathbb{R}^2: y^2 = 4\varepsilon^2(x+\varepsilon^2)} \sqrt{(x-\lambda)^2 + y^2} \\ &= \inf_{x \in [-\varepsilon^2, \infty)} \sqrt{(x-\lambda)^2 + 4\varepsilon^2(x+\varepsilon^2)} \\ &= \begin{cases} \lambda + \varepsilon^2, & \lambda \in (0, \varepsilon^2); \\ 2\varepsilon\sqrt{\lambda}, & \lambda \in [\varepsilon^2, \infty). \end{cases} \end{aligned}$$

This establishes the lemma. ■