Inapproximability After Uniqueness Phase Transition in Two-Spin Systems

Jin-Yi Cai* Xi Chen† Heng Guo‡ Pinyan Lu§

Abstract

A two-state spin system is specified by a matrix

$$\mathbf{A} = \begin{bmatrix} A_{0,0} & A_{0,1} \\ A_{1,0} & A_{1,1} \end{bmatrix} = \begin{bmatrix} \beta & 1 \\ 1 & \gamma \end{bmatrix}$$

where $\beta, \gamma \geq 0$. Given an input graph $G = (V, E)$, the partition function $Z_{\mathbf{A}}(G)$ of a system is defined as

$$Z_{\mathbf{A}}(G) = \sum_{\sigma: V \to \{0,1\}} \prod_{(u,v) \in E} A_{\sigma(u), \sigma(v)}. \tag{1}$$

We prove inapproximability results for the partition function in the region specified by the non-uniqueness condition from phase transition for the Gibbs measure. More specifically, assuming NP $\neq$ RP, for any fixed β, γ in the unit square, there is no randomized polynomial-time algorithm that approximates $Z_{\mathbf{A}}(G)$ for d-regular graphs G with relative error $\epsilon = 10^{-4}$, if $d = \Omega(\Delta(\beta, \gamma))$, where $\Delta(\beta, \gamma) > 1/(1 - \beta\gamma)$ is the uniqueness threshold. Up to a constant factor, this hardness result confirms the conjecture that the uniqueness phase transition coincides with the transition from computational tractability to intractability for $Z_{\mathbf{A}}(G)$. We also show a matching inapproximability result for a region of parameters β, γ outside the unit square, and all our results generalize to partition functions with an external field.

*University of Wisconsin, Madison.
†Columbia University.
‡University of Wisconsin, Madison.
§Microsoft Research Asia.

1 Introduction

Spin systems are well studied in statistical physics and applied probability. We focus on two-state spin systems. An instance of a spin system is a graph $G = (V, E)$. A configuration $\sigma : V \to \{0, 1\}$ assigns to every vertex one of two states. The contributions of local interactions between adjacent vertices are quantified by

$$\mathbf{A} = \begin{bmatrix} A_{0,0} & A_{0,1} \\ A_{1,0} & A_{1,1} \end{bmatrix} = \begin{bmatrix} \beta & 1 \\ 1 & \gamma \end{bmatrix},$$

a 2×2 matrix with $\beta, \gamma \geq 0$. The partition function $Z_{\mathbf{A}}(G)$ of a system is defined as

$$Z_{\mathbf{A}}(G) = \sum_{\sigma: V \to \{0,1\}} \prod_{(u,v) \in E} A_{\sigma(u),\sigma(v)}, \tag{2}$$

and we use $\omega(G, \sigma) = \prod_{(u,v) \in E} A_{\sigma(u),\sigma(v)}$ to denote the weight of σ.

For a fixed $\mathbf{A}$, we are interested in the complexity of computing $Z_{\mathbf{A}}(G)$, where G is given as an input. Many natural combinatorial counting problems can be formulated as two-state spin systems. For example, with $\beta = 0$ and $\gamma = 1$, $Z_{\mathbf{A}}(G)$ is exactly the number of independent sets (or vertex covers) of G. The definition of $Z_{\mathbf{A}}(G)$ in (2) can be generalized to larger $\mathbf{A}$, and the problem is also known as counting (weighted) graph homomorphisms [20, 16]. On the other hand, the Ising model is the special case where $\beta = \gamma$.

The *exact* complexity of computing $Z_{\mathbf{A}}(G)$ has been completely solved for any fixed symmetric matrix $\mathbf{A}$ [11, 3, 13, 6] and even for not necessarily symmetric $\mathbf{A}$ [8, 4, 2, 9, 7, 5] as part of the dichotomy theorems for the general counting CSP problem. When specialized to two-state spin systems, $Z_{\mathbf{A}}(G)$ is #P-hard to compute, except for the two restricted settings of $\beta\gamma = 1$ or $\beta = \gamma = 0$, in which cases it is polynomial-time computable. Consequently the study on two-state spin systems has focused on the approximation of $Z_{\mathbf{A}}(G)$, and this is the subject of the present paper.

Following standard definitions, a fully polynomial-time approximation scheme (FPTAS) for $Z_{\mathbf{A}}(G)$ is an algorithm that, given as input a graph G and a parameter $\epsilon > 0$, outputs a number Z that satisfies

$$(1 - \epsilon) \cdot Z_{\mathbf{A}}(G) \leq Z \leq (1 + \epsilon) \cdot Z_{\mathbf{A}}(G) \tag{3}$$

in time $\text{poly}(|G|, 1/\epsilon)$. A fully polynomial-time randomized approximation scheme (FPRAS) is then a randomized algorithm that, with probability $1 - \delta$, outputs a Z satisfying (3) in time $\text{poly}(|G|, 1/\epsilon, \log(1/\delta))$.

For the Ising model, in a seminal paper [17] Jerrum and Sinclair gave an FPRAS for $Z_{\mathbf{A}}(G)$ when $\beta = \gamma > 1$. It was further extended to the entire region of $\beta\gamma > 1$ by Goldberg, Jerrum and Paterson [14]. A two-state spin system is called *ferromagnetic* if $\beta\gamma > 1$ and *anti-ferromagnetic* if $\beta\gamma < 1$. The approximability of $Z_{\mathbf{A}}(G)$ for anti-ferromagnetic systems is less well understood. Starting with counting independent sets in sparse graphs [10], the approximability of $Z_{\mathbf{A}}(\cdot)$ in bounded degree graphs is also widely studied. Significant progress has been made recently on the algorithmic side, and approximation algorithms for anti-ferromagnetic two-state spin systems have been developed in [24, 22, 19, 18], based on the technique of correlation decay introduced by Bandyopadhyay and Gamarnik [1] and Weitz [24]. Finally, a unified FPTAS was found [18] to approximate $Z_{\mathbf{A}}(\cdot)$ for all anti-ferromagnetic two-state spin systems of either bounded degree graphs or general graphs, when the system satisfies a *uniqueness condition*.

This uniqueness condition is named for, and closely related to, phase transitions that occur for the Gibbs measure. It depends on not only β and γ but also the degree of the underlying graph as well. Such phase transitions from statistical physics are believed to frequently coincide with the transitions of computational complexity from tractability to intractability. However, there are only very few examples where the conjectured link is rigorously proved. One notable example is for the hardcore gas model (independent set), for which such a conjecture was rigorously proved (for almost all degree bounds) both for the algorithmic side [24] and for the hardness side [23, 12]. As discussed above [24, 22, 19, 18], for general anti-ferromagnetic two-state spin systems the algorithmic part of the conjecture has recently been established. In this paper, we make substantial progress on the hardness part of the conjecture.

Our Results. For $\beta, \gamma : 0 \le \beta, \gamma \le 1$ except at $(\beta, \gamma) = (0, 0)$ or $(1, 1)$, Goldberg, Jerrum, and Paterson proved that the problem does not admit an FPRAS on general graphs (when there is no degree bound) unless NP = RP [14]. In their reduction, the degrees of the hard instances are unbounded. This is consistent with the uniqueness threshold conjecture. However, for any fixed β, γ in the unit square, the uniqueness condition states that there exists a finite threshold degree $\Delta(\beta, \gamma)$ [22, 19, 18] where

$$\Delta(\beta, \gamma) > \frac{1 + \sqrt{\beta\gamma}}{1 - \sqrt{\beta\gamma}} = \frac{(1 + \sqrt{\beta\gamma})^2}{1 - \beta\gamma} \ge \frac{1}{1 - \beta\gamma}, \tag{4}$$

such that the system satisfies the uniqueness condition if the degree $d < \Delta(\beta, \gamma)$, and the non-uniqueness condition if $d \ge \Delta(\beta, \gamma)$. The paper [18] gives an FPTAS for graphs with degree bounded by $\Delta(\beta, \gamma)$. The conjectured coincidence of phase transition with hardness in complexity suggests that as soon as the degree of the input graph goes beyond $\Delta(\beta, \gamma)$, the problem becomes hard to approximate. Towards this direction we show that for any fixed β, γ in the unit square, the problem does not have an FPRAS if the degree of the input graph is $\Omega(\Delta(\beta, \gamma))$, unless NP = RP. Our hardness also holds when restricted to input graphs that are regular. Formally, we prove

Theorem 1. *There exists a positive constant h such that: Given any $\beta, \gamma : 0 \le \beta, \gamma \le 1$ with $(\beta, \gamma) \ne (0, 0)$, $(1, 1)$ and any integer $d \ge h/(1 - \beta\gamma)$, there is no randomized polynomial-time algorithm that approximates $Z_{\mathbf{A}}(G)$ in d-regular graphs G with relative error $\epsilon = 10^{-4}$, unless* NP = RP.

Note the relation between our degree bound $h/(1 - \beta\gamma)$ and $\Delta(\beta, \gamma)$ from (4).

We also make progress on (β, γ) outside the unit square. While the uniqueness condition is *monotone* inside the unit square, its behavior outside is significantly different. By symmetry we consider the region defined by $\beta\gamma < 1$ with $\beta, \gamma : 0 < \beta < 1 < \gamma$. There is a uniqueness curve, connecting the point $(1, 1)$ and the γ-axis, above which the system satisfies the uniqueness condition for any graph [19, 18]. Hence, hardness is only possible below this uniqueness curve. Furthermore, when (β, γ) is outside the unit square but below this uniqueness curve, there is only a finite range of degrees d for which the system does not satisfy the uniqueness condition. This makes it very challenging to prove hardness result for them. Previously, the hardness was only obtained in [14] for a very tiny square $0 \le \beta \le \eta$ and $1 \le \gamma \le 1 + \eta$ where η is roughly 10^{-7}, near the point $(0, 1)$ corresponding to independent set or the hardcore gas model. In this paper, we prove the following hardness result for (β, γ) outside the unit square:

Theorem 2. *Given β and γ such that $0 < \beta < 1$, $\gamma > 1$ and $\beta\gamma < 1$, let*

$$\Delta' = \left\lceil -1/(\ln \beta + \ln \gamma) \right\rceil \quad \text{and} \quad \Delta^* = \left\lceil 1/\ln \gamma \right\rceil. \tag{5}$$

When $\Delta^* \ge 8000\Delta'$, there is no randomized polynomial-time algorithm that approximates $Z_{\mathbf{A}}(G)$ in regular graphs of degree Δ^* with relative error $\epsilon = 10^{-4}$, unless NP = RP.

The new hardness region is pictured in Fig 1.[1] The two white squares are the hardness regions acquired before by Goldberg et.al. [14], and beyond the uniqueness threshold it is known that FPTAS exists. The region where we show is hard to approximate is from the vertical line segment $(0 < \beta < \gamma = 1)$ to the curve lower than the uniqueness threshold. Let us describe the new region we obtain a little further. Again we describe the region with $0 < \beta < 1 < \gamma$ within $\beta\gamma < 1$; there is a symmetric region where $0 < \gamma < 1 < \beta$. Near the point $(1, 1)$, the condition in Theorem 2 is almost linear, roughly a line with slope -8000. Approaching the line of $\beta = 0$, Δ' becomes 1 and the condition requires γ to be between 1 and roughly $1 + 1/8000$.

Moreover, using a standard translation described in Appendix C, we can generalize Theorem 1 and 2 to two-state spin systems with an external field. Formally, let $\mu \ge 0$, we have the following two corollaries for

[1] The reader should note that, for illustration purposes, here we are not drawing it according to the real scale nor in the precise shape.

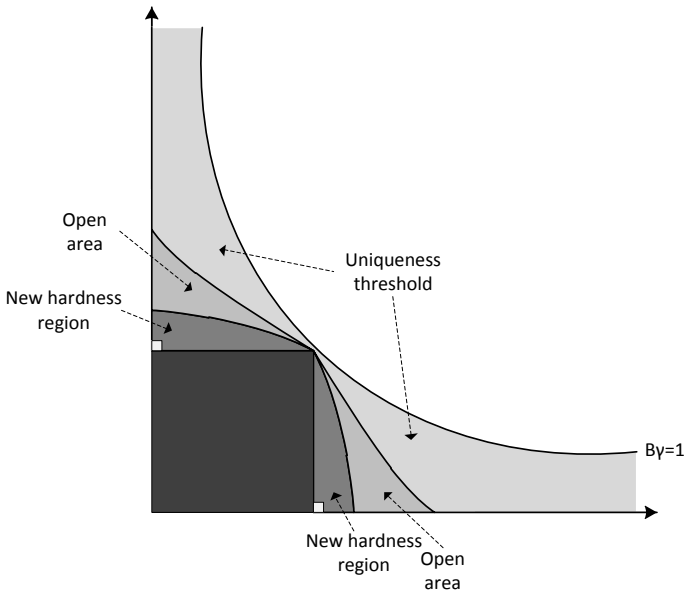

Figure 1: The new hardness region of Theorem 2.

$$Z_{\mathbf{A},\mu}(G) = \sum_{\sigma:V\to\{0,1\}} \mu^{\left|\{v\in V:\sigma(v)=0\}\right|} \prod_{(u,v)\in E} A_{\sigma(u),\sigma(v)}.$$

Corollary 1. *There exists a constant h such that, given any nonnegative β, γ, μ with $\beta\gamma < 1$, and an integer d satisfying that $\gamma \le \mu^{\frac{1}{d}} \le \frac{1}{\beta}$ and $d \ge \frac{h}{1-\beta\gamma}$, there is no randomized polynomial-time algorithm that approximates $Z_{\mathbf{A},\mu}(G)$ in d-regular graphs with relative error $\epsilon = 10^{-4}$ unless* NP = RP.

Corollary 2. *Given any nonnegative β, γ, μ, and an integer d that $e^{\frac{1}{d}} \le \gamma\mu^{-\frac{1}{d}} < e^{\frac{1}{d-1}}$ or $e^{\frac{1}{d}} \le \beta\mu^{\frac{1}{d}} < e^{\frac{1}{d-1}}$, if $d \ge 8000\lceil -1/(\ln\beta + \ln\gamma)\rceil$, then there is no randomized polynomial-time algorithm that approximates $Z_{\mathbf{A},\mu}(G)$ in d-regular graphs with relative error $\epsilon = 10^{-4}$ unless* NP = RP.

Proof Outline. The overall idea is to use the phase transition that occurs in the non-uniqueness region to encode a hard approximation problem. This approach has been used in previous hardness proofs for the hardcore gas model [10, 21, 23]. We reduce the approximation problem for E2LIN2 to the approximation of partition function in our two-state spin system. An instance of E2LIN2 is a set of equations of the form $x_i + x_j = 0$ or 1 over $\mathbb{Z}_2$ in n variables $x_1, x_2, \ldots, x_n$. It is known [15] that it is NP-hard to approximate the number of satisfiable equations for E2LIN2 within a constant factor 11/12. One uses a random bipartite regular graph to encode each variable x_i. Due to phase transition, we are in a non-uniqueness region, each bipartite regular graph would be in one of two types of configurations with high probability if it is sampled proportional to its weight in the partition function. This can be used to establish a correspondence between the configurations within the bipartite graph and the assignment of that variable. Furthermore, external connections of these bipartite graphs connected according to the given equations in the E2LIN2 instance contribute exponentially separated total weight in the partition function according to the proportion of the number of equations an assignment satisfies. Thus a sufficiently good approximation to the partition function can be used to decode an assignment that approximates the maximum satisfiability.

Our gadget is also randomly constructed. Then the probability should also be over the distribution of the gadgets. It is not hard to show that things work out beautifully if we simply substitute the expectation for the actual weight. But to make this proof rigorous, one must have a sufficiently good concentration result. Such a result is unknown and could be very difficult to prove (assuming it is true), as it is already a tour-de-force in the special case for the hardcore gas model [21, 23, 12].

Instead we use a detour: (1) We prove a lower bound for the weights of two types of configurations we expect, guided by the phase transition. (2) We prove that the total weight of other configurations is

exponentially small compared to the lower bound with a probability exponentially close to 1. The way we establish the lower bound in (1) is similar to the approach by Dyer, Frieze and Jerrum [10]. To prove (2), they used [10] the expectation and Markov's inequality. If we use the same approach, we could also get some hardness result for bounded degree graphs. However, the result is not in the same order of the uniqueness bound. Therefore, we use a new approach for (2).

In fact, we prove a high concentration result for an expander property of the gadgets we use, then we directly show that the total weight of other configurations is exponentially small, given that the gadgets satisfy that property. This circumvented our inability to prove a complete concentration result. But we do prove some limited concentration result regarding the gadget, which let us prove hardness results for degrees in the right order conjectured according to the uniqueness threshold. It remains open whether one can use a refined version of this reduction along the proof by Sly [23] to prove the exact right bound. As we discuss in Appendix A, this random regular graph quite closely follows the property of phase transition in infinite d-ary trees, when the parameter is below or beyond the uniqueness condition.

Although the high level idea is quite clear and similar for both Theorem 1 and Theorem 2, it is still a challenge to do the estimation right for all ranges of parameters and get the degree in the same order of the uniqueness bound. In order to do that, technically we need to use quite different approaches for Theorem 1 and Theorem 2, and even within the unit square (Theorem 1), we need to do the estimation differently for three different subcases.

2 Proof of the Main Theorems

From now on, we will simply use $Z(G)$ to denote $Z_{\mathbf{A}}(G)$ or $Z_{\mathbf{A},\mu}(G)$ whenever it is clear from the context.

Given positive integers N and Δ, let $\mathcal{H}(N,\Delta)$ denote the following probability distribution of Δ-regular bipartite graphs $H=(U\cup V,E)$ with bipartition U,V and $|U|=|V|=N$: here H is the union of Δ perfect matchings between U and V each selected independently and uniformly at random. (Because these perfect matchings are drawn independently, H may have parallel edges.)

In the proof of both Theorem 1 and Theorem 2 we give a polynomial-time reduction from E2LIN2 to the approximation of $Z(G)$. An instance of E2LIN2 consists of m equations over $\mathbb{Z}_2$ in n variables $x_1,x_2,\ldots,x_n$ where each equation has exactly two variables and is of the form $x_i+x_j=0$ or 1. Given an assignment S of the n variables, we use $\theta(S)$ to denote the number of equations that S satisfies, and let $\theta^*=\max_S\theta(S)$. In [15], Håstad showed that it is NP-hard to estimate θ^* within any constant factor smaller than $11/12$.

Given an instance of E2LIN2, we construct a random $(\Delta+\Delta')$-regular graph G as follows, with the two parameters Δ,Δ' to be specified later. This construction is used in the proof of both Theorem 1 and 2:

> **Construction of G from an instance of E2LIN2.** For each variable x_i, $i\in[n]$, we let U_i and V_i denote two sets of d_im vertices each, where $d_i\geq 1$ is the number of times that x_i appears in the equations (and thus, $\sum_i d_i=2m$). Enumerate all the m equations in the instance. Denote the kth equation by $x_i+x_j=b\in\{0,1\}$, and add the following $2\Delta' m$ edges:
>
> **(1)** Pick vertices $\{u_1,\ldots,u_m\}$ in U_i, $\{v_1,\ldots,v_m\}$ in V_i, $\{u'_1,\ldots,u'_m\}$ in U_j, and $\{v'_1,\ldots,v'_m\}$ in V_j, all of degree 0 at this moment. Denote these four sets of vertices by $U_{i,k}$, $V_{i,k}$, $U_{j,k}$ and $V_{j,k}$, respectively. If $b=0$, we add Δ' parallel edges between (u_i,v'_i) and (v_i,u'_i) for each $i\in[n]$; if $b=1$, we add Δ' parallel edges between (u_i,u'_i) and (v_i,v'_i) for each $i\in[n]$.
>
> By the end of this step, every vertex has degree Δ'. In the next step,
>
> **(2)** For each $i\in[n]$, we add a bipartite graph $H_i=(U_i\cup V_i,E_i)$ drawn from $\mathcal{H}(d_im,\Delta)$.

This finishes the construction, and we get a $(\Delta+\Delta')$-regular graph G with $4m^2$ vertices.

We need the following notation. Given an assignment S of the variables $x_1,\ldots,x_n$, we let $U_i(\sigma)$ denote the number of vertices $u\in U_i$ with $\sigma(u)=0$, and use $V_i(\sigma)$ denote the number of $v\in V_i$ with $\sigma(v)=0$.

Proof of Theorem 1. Without loss of generality, we assume $0 \le \beta \le \gamma \le 1$. We can also assume that $\beta > 0$ since the tight hardness to the exact uniqueness bound for $\beta = 0$ has been shown in [19] by generalizing the tight hardness result for the hardcore model [23, 12].

Given any assignment S of the n variables, we let $Z(G, S)$ denote the sum of $\omega(G, \sigma)$ over assignments $\sigma : V(G) \to \{0, 1\}$ that satisfy for each $i \in [n]$,

$$U_i(\sigma) \le V_i(\sigma) \text{ if } x_i = 0 \text{ in } S; \text{ or } U_i(\sigma) \ge V_i(\sigma) \text{ if } x_i = 1 \text{ in } S. \tag{6}$$

By definition, we have $Z(G, S) \le Z(G) \le \sum_S Z(G, S)$. We prove the following key lemma in Section 3:

Lemma 1. *There exists a positive constant h. For any $\beta, \gamma : 0 < \beta \le \gamma \le 1$ with $(\beta, \gamma) \ne (1, 1)$ and for any $\Delta^* \ge h/(1 - \beta\gamma)$, there are $D > 1$, $C > 0$ and positive integers Δ and Δ' with $\Delta + \Delta' = \Delta^*$, that satisfy the following property: Given any input instance of* E2LIN2 *with n variables $x_1, \ldots, x_n$ and m equations, except for probability $\exp(-\Omega(m))$, the $(\Delta + \Delta') = \Delta^*$-regular graph G constructed with parameters Δ, Δ' satisfies*

$$C^{m^2} \cdot D^{m\theta(S)} \le Z(G, S) \le C^{m^2} \cdot D^{m\left(\theta(S) + 0.03m\right)}, \quad \text{for any assignment } S \text{ of the } n \text{ variables.} \tag{7}$$

Given β, γ and Δ^*, we let C, w, Δ and Δ' denote the constants that satisfy the property in Lemma 1, then given an input instance of E2LIN2, (7) holds with probability $1 - \exp(-\Omega(m))$.

Now assume (7) holds. We use θ^* to denote the maximum number of consistent equation and use S^* to denote an assignment that satisfies θ^* equations. We also use Y to denote an estimate of $Z = Z(G)$, where $|Y/Z - 1| \le \epsilon = 10^{-4}$. From (7) and $Z(G) \le \sum_S Z(G, S) \le 2^n \cdot Z(G, S^*)$, we get

$$(1 + \epsilon) \cdot 2^n \cdot C^{m^2} \cdot D^{m(\theta^* + 0.03m)} \ge (1 + \epsilon) \cdot Z \ge Y \ge (1 - \epsilon) \cdot Z \ge (1 - \epsilon) \cdot C^{m^2} \cdot D^{m\theta^*} \tag{8}$$

Using Y, we set

$$Y' = \frac{\ln Y - \ln(1 + \epsilon) - n \ln 2 - nm \ln C - 0.03nm \ln D}{n \ln D}$$

and we get $Y' \le \theta^*$ since $\ln D > 0$. We finish the proof by showing that $Y' \ge (11/12) \cdot \theta^*$. By (8), we have

$$Y' \ge \theta^* - \frac{\ln(1 + \epsilon) - \ln(1 - \epsilon) + n \ln 2 + 0.03nm \ln D}{n \ln D}$$

As $\theta^* \ge m/2$, when m is large enough, it follows that $Y' \ge (11/12) \cdot \theta^*$ and the theorem is proven. □

Next, we prove Theorem 2:

Proof of Theorem 2. For β, γ with $0 < \beta < 1 < \gamma$ and $\beta\gamma < 1$, let Δ' and Δ^* be the two positive integers defined in (5) that satisfy $\Delta^* \ge 8000\Delta'$. We set $\Delta = \Delta^* - \Delta'$. Given any input instance of E2LIN2 with n variables and m equations, we use G to denote the Δ^*-regular graph constructed using Δ and Δ'.

First of all, we show that to get a good approximation of $Z(G)$, with high probability it suffices to sum $\omega(G, \sigma)$ only over assignments σ that satisfy the following condition:

$$\min\Big(\big|U_i(\sigma)\big|, \big|V_i(\sigma)\big|\Big) \le \lambda d_i m, \quad \text{for all } i \in [n], \text{ where } \lambda = 9 \times 10^{-5}. \tag{9}$$

We use Σ to denote the set of such assignments. Formally, we prove the following key lemma in Section 4:

Lemma 2. *Let G be the graph constructed from an instance of* E2LIN2 *with n variables $x_1, x_2, \ldots, x_n$ and m equations, with parameters Δ and Δ'. Then with probability $1 - \exp(-\Omega(m^{1/3}))$, it satisfies*

$$\sum_{\sigma \in \Sigma} \omega(G, \sigma) \le Z(G) \le (1 + o(1)) \cdot \sum_{\sigma \in \Sigma} \omega(G, \sigma). \tag{10}$$

Next, given any assignment S over the n variables, we use $Z_\Sigma(G,S)$ to denote the sum of $\omega(G,\sigma)$ over all assignments $\sigma \in \Sigma$ that satisfy (6) for all $i \in [n]$. We prove the following lemma in Appendix D:

Lemma 3. *There are $C > 0$ and $D > 1$ satisfying the following property: given an instance of* E2LIN2 *with n variables and m equations, the Δ^*-regular graph G constructed with parameters Δ and Δ' satisfies*

$$C^{m^2} \cdot D^{m\theta(S)} \le Z_\Sigma(G,S) \le C^{m^2} \cdot D^{m\left(\theta(S)+0.04m\right)}, \quad \text{for any assignment } S \text{ of the } n \text{ variables.} \tag{11}$$

Let $\theta^* \ge m/2$ be the maximum number of consistent equations in the instance of E2LIN2. Let S^* denote an assignment that satisfies θ^* equations. From these two lemmas, we have with high probability that

$$C^{m^2} \cdot D^{m\theta^*} \le Z_\Sigma(G,S^*) \le Z(G) \le \big(1+o(1)\big) \cdot \sum_S Z_\Sigma(G,S) \le \big(1+o(1)\big) \cdot 2^n \cdot C^{m^2} \cdot D^{m(\theta^*+0.04m)}$$

Theorem 2 then follows from the same argument used in the proof of Theorem 1. □

3 Proof of Lemma 1

Recall that β and γ satisfy $0 < \beta \le \gamma \le 1$ and $(\beta,\gamma) \ne (1,1)$.

Given any instance of E2LIN2 with n variables $x_1,\dots,x_n$ and m equations, the $(\Delta+\Delta')$-regular graph G we construct consists of Δ-regular bipartite graphs H_i, $i \in [n]$, and edges between them. For each H_i in G we use $U_i \cup V_i$ denote its vertex set with $|U_i| = |V_i| = d_i m$. Recall that $\epsilon = 10^{-4}$. Then we say $A \subseteq U_i$ is *big* if $|A| \ge \epsilon|U_i|$, and $B \subseteq V_i$ is *big* if $|B| \ge \epsilon|V_i|$. We also use $E(H_i,A,B)$ to denote the number of edges between A and B in H_i. Let $K = 48 \times 10^8$. We prove the following lemma:

Lemma 4. *When $\Delta \ge K$, with probability $1 - \exp(-\Omega(m))$, the graph G we construct satisfies*

$$E(H_i,A,B) \ge \frac{\Delta|A||B|}{4d_i m}, \quad \text{for all } i \in [n] \text{ and for all big } A \subseteq U_i \text{ and big } B \subseteq V_i. \tag{12}$$

Lemma 4 follows from an application of the Chernoff bound, and we give the proof in Appendix E.

To finish the proof we divide (β,γ) into three cases. For each case we show there exists a large enough constant h with the following property: For all (β,γ) of this case, and for all $\Delta^* \ge h/(1-\beta\gamma)$, there are $C > 0$, $D > 1$ and positive integers $\Delta \ge K$ and $\Delta' \ge 1$ with $\Delta + \Delta' = \Delta^*$, such that (7) holds whenever G satisfies (12). The lemma then follows by taking the maximum of the three h's.

Let $L = 12/\epsilon^2 = 12 \times 10^8$. In the rest of the proof, we assume that G satisfies (12), and let $M = m^2$.

3.1 Case 1: $0 < \beta < 1/2$ and $\beta \le \gamma^L$

We set h to be a large enough constant so that $h/(1-\beta\gamma) \ge 7(L+1)$. Given a $\Delta^* \ge h/(1-\beta\gamma)$, we then set $\Delta = \lfloor L\Delta^*/(L+1) \rfloor$ and $\Delta' = \lceil \Delta^*/(L+1) \rceil$ where $\Delta + \Delta' = \Delta^*$, $\Delta > K$ and $L\Delta' \ge \Delta \ge L(\Delta'-1) - 1$.

Let S be an assignment over the n variables $x_1,\dots,x_n$, then we start with a lower bound $Z^*(G,S)$ for $Z(G,S)$. To this end, we consider the sum of $\omega(G,\sigma)$ over all assignments σ that satisfy for each $i \in [n]$:

$$U_i(\sigma) = 0 \ \text{ if } x_i = 0 \text{ in } S; \text{ and } \ V_i(\sigma) = 0 \ \text{ otherwise.} \tag{13}$$

Denote this sum by $Z^*(G,S)$. It is clearly a lower bound for $Z(S)$, and is exactly equal to

$$\begin{aligned} Z^*(G,S) &= \left(1 + 2\gamma^{\Delta+\Delta'} + \gamma^{2\Delta+2\Delta'}\right)^{m\theta(S)} \cdot \left(\beta^{\Delta'}\gamma^{\Delta'} + 2\gamma^{\Delta+\Delta'} + \gamma^{2\Delta+2\Delta'}\right)^{m(m-\theta(S))} \\ &= \left(\beta^{\Delta'}\gamma^{\Delta'} + 2\gamma^{\Delta+\Delta'} + \gamma^{2\Delta+2\Delta'}\right)^{M} \cdot \left(\frac{1 + 2\gamma^{\Delta+\Delta'} + \gamma^{2\Delta+2\Delta'}}{\beta^{\Delta'}\gamma^{\Delta'} + 2\gamma^{\Delta+\Delta'} + \gamma^{2\Delta+2\Delta'}}\right)^{m\theta(S)}. \end{aligned} \tag{14}$$

Setting C and D appropriately, we have $Z^*(G,S) = C^M \cdot D^{m\theta(S)}$, where

$$C = \beta^{\Delta'}\gamma^{\Delta'} + 2\gamma^{\Delta+\Delta'} + \gamma^{2\Delta+2\Delta'} > 0 \quad \text{and} \quad D = \frac{1 + 2\gamma^{\Delta+\Delta'} + \gamma^{2\Delta+2\Delta'}}{\beta^{\Delta'}\gamma^{\Delta'} + 2\gamma^{\Delta+\Delta'} + \gamma^{2\Delta+2\Delta'}} > 1 \tag{15}$$

since $(\beta,\gamma) \neq (0,0),(1,1)$. It is also easy to give a lower bound of $8/7$ for D because the difference between the numerator and the denominator is $1 - \beta^{\Delta'}\gamma^{\Delta'} > 1/2$ as $\beta < 1/2$; and the denominator of D is $< 7/2$.

Next, to give an upper bound for $Z(G,S)$, we consider the sum of $\omega(G,\sigma)$ over σ that satisfies

$$U_i(\sigma) \leq \epsilon d_i n \ \text{ if } x_i = 0; \text{ and } \ V_i(\sigma) \leq \epsilon d_i n \ \text{ if } x_i = 1 \tag{16}$$

for every $i \in [n]$. We show that this sum is indeed a good approximation of $Z(G,S)$:

$$Z(G,S) \leq \big(1 + o(1)\big) \sum_{\sigma \text{ that satisfies (16)}} \omega(G,\sigma) \tag{17}$$

To prove (17) we randomly draw an assignment σ from those appear in the sum $Z(G,S)$ with probability proportional to $\omega(G,\sigma)$, and it suffices to show that the probability that σ satisfies (16) for all i is $1 - o(1)$. This then follows from the following lemma and the union bound:

Lemma 5. *For any $i \in [n]$, the probability that σ violates (16) is at most $\exp(-d_i m) \ll 1/m$.*

Proof. Without loss of generality, we assume $x_i = 0$ in S. Pick any partial assignment σ' over all vertices of G except those of H_i. To prove the lemma it suffices to show that the sum of $\omega(G,\sigma)$ over all assignments σ that are consistent with σ' but violate (16) is exponentially smaller than $\omega(G,\sigma^*)$, where σ^* denotes the unique assignment that is consistent with σ' and satisfies $U_i(\sigma^*) = 0$ and $V_i(\sigma^*) = d_i m$.

To this end, we let $\omega(\sigma')$ denote the product of the edge weights in σ' over all edges in G except those have at least one vertex in H_i. Then it is easy to give the following lower bound for $\omega(G,\sigma^*)$:

$$\omega(G,\sigma^*) \geq \omega(\sigma') \cdot \big(\beta\gamma\big)^{\Delta' d_i m} \tag{18}$$

On the other hand, for any σ that is consistent with σ' but violates (16), we have

$$\omega(G,\sigma) \leq \omega(\sigma') \cdot \beta^{\,\epsilon^2 \Delta d_i m/4}. \tag{19}$$

It follows from the assumption of (12) and $V_i(\sigma) \geq U_i(\sigma) \geq \epsilon d_i m$ as $x_i = 0$ and σ violates (16).

Plugging in Δ and Δ', we have

$$\frac{\omega(G,\sigma^*)}{\omega(G,\sigma)} \geq \frac{\omega(\sigma') \cdot \big(\beta\gamma\big)^{\Delta' d_i m}}{\omega(\sigma') \cdot \beta^{\,\epsilon^2 \Delta d_i m/4}} \geq \left(\frac{\beta^{2\Delta'}}{\beta^{3(\Delta'-1)-(\epsilon^2/4)}}\right)^{d_i m} = \left(\frac{1}{\beta^{\Delta'-3-(\epsilon^2/4)}}\right)^{d_i m} > 2^{3 d_i m}.$$

The lemma follows because the number of σ that is consistent with σ' but violates (16) is at most $2^{2 d_i m}$. □

We continue with (17). For each σ that satisfies (13), let T_σ denote the following set of assignments σ': (1) σ' satisfies (16) for all i; and (2) for each i, σ' is consistent with σ over V_i if $x_i = 0$; and over U_i if $x_i = 1$.

It is clear that $\{T_\sigma\}$ is a partition of the assignments that satisfy (16) for all $i \in [n]$. It is also easy to check that for any σ that satisfies (13), we know exactly the cardinality of $|T_\sigma|$:

$$|T_\sigma| = \prod_{i\in[n]} \left(\sum_{j=0}^{\epsilon d_i m} \binom{d_i m}{j}\right) \leq \prod_{i\in[n]} \left(\big(\epsilon d_i m + 1\big)\binom{d_i m}{\epsilon d_i m}\right) \leq m^{2n} \cdot \prod_{i\in[n]} e^{H(\epsilon)\, d_i m} = m^{2n} \cdot e^{2H(\epsilon) M} \tag{20}$$

where $H(\epsilon) \approx 0.00102$. For any $\sigma' \in T_\sigma$, we also have

$$\omega(G, \sigma') \leq \omega(G, \sigma) \cdot \left(1/\gamma^{\Delta+\Delta'}\right)^{2\epsilon M}, \tag{21}$$

because for any assignment, switching the value of a vertex from 0 to 1 can improve $\omega(G, \sigma')$ by at most a factor of $1/\gamma^{\Delta+\Delta'}$. Finally, by combining (17), (20) and (21) we get the following upper bound for $Z(G, S)$:

$$Z(G, S) \leq \big(1 + o(1)\big) \cdot m^{2n} \cdot e^{2H(\epsilon)M} \cdot Z^*(G, S) \cdot \left(1/\gamma^{\Delta+\Delta'}\right)^{2\epsilon M}. \tag{22}$$

To finish the proof and show (7): $Z(G, S) < Z^*(G, S) \cdot D^{0.03M}$, we also need to compare D with $1/\gamma^{\Delta+\Delta'}$. Since $\beta \leq \gamma^L$, we have $\beta^{\Delta'} \leq \gamma^{L\Delta'} \leq \gamma^{\Delta}$. It follows from the definition of D that $D \geq 1/(4\gamma^{\Delta+\Delta'})$. Then (7) follows directly from (22) by plugging in $D > 8/7$ and $\epsilon = 10^{-4}$.

3.2 Case 2: $\beta \geq 1/2$ and $\beta \leq \gamma^L$

The proof for Case 2 is similar, and can be found in Appendix B.

3.3 Case 3: $\beta > \gamma^L$

For this case, we need to use a different estimation for $Z(G, S)$.

We start by setting Δ' and Δ. Let h be a large enough constant such that for any $\Delta^* \geq h/(1 - \beta\gamma)$,

$$\Delta = \left\lceil \frac{L(L+1) \cdot \Delta^*}{L(L+1)+1} \right\rceil \quad \text{and} \quad \Delta' = \left\lfloor \frac{\Delta^*}{L(L+1)+1} \right\rfloor$$

satisfy $\Delta' \geq 1$ and $(\beta\gamma)^{\Delta'} < 1/4$. It follows from the definition that $\Delta^* = \Delta + \Delta'$ and $\Delta \geq L(L+1)\Delta'$.

Given an assignment S over the n variables $x_1, \ldots, x_n$, we use $\hat{\sigma}$ to denote the unique assignment with $U_i(\hat{\sigma}) = 0$ and $V_i(\hat{\sigma}) = d_i m$ when $x_i = 0$; $U_i(\hat{\sigma}) = d_i m$ and $V_i(\hat{\sigma}) = 0$ when $x_i = 1$, for all $i \in [n]$. Then

$$Z(G, S) > \omega(G, \hat{\sigma}) = \big(\beta\gamma\big)^{\Delta' m(m-\theta(S))} = \left(\big(\beta\gamma\big)^{\Delta'}\right)^M \cdot \left(\frac{1}{\big(\beta\gamma\big)^{\Delta'}}\right)^{m\theta(S)} \tag{23}$$

Setting $C = (\beta\gamma)^{\Delta'}$ and $D = 1/C$, we get $Z(G, S) > C^M \cdot D^{m\theta(S)}$, with $D > 4$ and $C > 0$.

Next, to give an upper bound for $Z(S)$, we consider the sum of $\omega(G, \sigma)$ over σ that satisfies

$$U_i(\sigma) \leq \epsilon \cdot d_i m \quad \text{and} \quad V_i(\sigma) \geq (1-\epsilon) \cdot d_i m, \quad \text{when } x_i = 0 \text{ in } S; \tag{24}$$

$$U_i(\sigma) \geq (1-\epsilon) \cdot d_i m \quad \text{and} \quad V_i(\sigma) \leq \epsilon \cdot d_i m, \quad \text{when } x_i = 1 \text{ in } S. \tag{25}$$

for every $i \in [n]$. We show that this sum is a good approximation of $Z(G, S)$:

$$Z(G, S) \leq \big(1 + o(1)\big) \sum_{\sigma \text{ that satisfies } (24,25)} \omega(G, \sigma) \tag{26}$$

To prove (26), we randomly draw a σ from those appear in the sum $Z(G, S)$ with probability proportional to $\omega(G, \sigma)$, and show that the probability that σ violates (24) or (25) is exponentially small.

For this purpose, we prove the same statement as in Lemma 5. Pick any $i \in [n]$, and assume $x_i = 0$ in S without loss of generality. Let σ' be any partial assignment over all vertices of G except those of H_i. We use σ^* to denote the unique assignment that is consistent with σ' and satisfies $U_i(\sigma^*) = 0$ and $V_i(\sigma^*) = d_i m$

and use σ to denote any assignment that is consistent with σ' but violates (24) in H_i. Then we get

$$\frac{\omega(G,\sigma^*)}{\omega(G,\sigma)} \geq \frac{\omega(\sigma')\cdot\big(\beta\gamma\big)^{\Delta' d_i m}}{\omega(\sigma')\cdot\gamma^{\epsilon^2\Delta d_i m/4}} > \left(\frac{\gamma^{(L+1)\Delta'}}{\gamma^{3(L+1)\Delta'}}\right)^{d_i m} = \left(\frac{1}{\gamma^{2(L+1)\Delta'}}\right)^{d_i m} > 2^{4d_i m}$$

Here the first inequality follows from (12) and the fact that, since σ violates (24), either

$$V_i(\sigma) \geq U_i(\sigma) > \epsilon\cdot d_i m \quad \text{or} \quad U_i(\sigma) \leq V_i(\sigma) < (1-\epsilon)\cdot d_i m \tag{27}$$

and the last inequality follows from $1/4 > (\beta\gamma)^{\Delta'} > \gamma^{(L+1)\Delta'}$.

This proves (26). Moreover, the number of σ that satisfies both (24) and (25) for all $i \in [n]$ can be easily bounded by $m^{4n}\cdot e^{4H(\epsilon)M}$. And for any σ that satisfies both (24) and (25), we also have

$$\omega(G,\sigma) \leq \omega(G,\hat{\sigma})\Big/\big(\beta\gamma\big)^{2\epsilon\Delta' M} \tag{28}$$

This is because, to obtain σ from $\hat{\sigma}$, each time we flip a vertex from 0 to 1, the weight increases by a factor of at most $1/\gamma^{\Delta'}$; each time we flip a vertex from 1 to 0, the weight increases by a factor of at most $1/\beta^{\Delta'}$.

Finally, combining (26), (27) and (28), we get

$$Z(G,S) \leq \big(1+o(1)\big)\cdot m^{4n}\cdot e^{4H(\epsilon)M}\cdot\omega(G,\hat{\sigma})\cdot D^{2\epsilon M}$$

Then (7) follows immediately by plugging in $D > 2$ and $\epsilon = 10^{-4}$. This finishes the proof of the lemma.

4 Proof of Lemma 2

Recall that β and γ satisfy $0 < \beta < 1 < \gamma$ and $\beta\gamma < 1$. Let Δ' and Δ^* be the two positive integers defined in Theorem 2 with $\Delta^* \geq 8000\Delta'$. From their definitions, we have $(\beta\gamma)^{\Delta'} \leq 1/e$ and $\gamma^{\Delta^*} \geq e$. Set $\Delta = \Delta^* - \Delta' \geq 7999\Delta' \geq 7999$. By the definition of Δ^*, we have $e > \gamma^{\Delta^*-1} \geq \gamma^{\Delta}$ and thus, $\gamma < 1.001$.

Given an instance of E2LIN2 with n variables $x_1,\ldots,x_n$ and m equations, we use G to denote the $(\Delta+\Delta') = \Delta^*$-regular graph constructed with parameters Δ and Δ'. We use H_i to denote the bipartite graph in G that corresponds to x_i and use $U_i \cup V_i$ to denote its vertices, with $|U_i| = |V_i| = d_i m$.

Before working on G and H_i, we start by proving a property that a bipartite graph sampled from the distribution $\mathcal{H}$ satisfies with high probability. Let H be a bipartite graph drawn from $\mathcal{H}(N,\Delta)$ for some $N \geq 1$ and Δ defined above, with $2N$ vertices $U \cup V$. We use $\rho : U\cup V \to \{0,1\}$ to denote an assignment and call it an (a,b)-assignment for some $a,b \in T_N$, where $T_N = \big\{0, 1/N, 2/N, \ldots, (N-1)/N, 1\big\}$ if $\big|u \in U : \rho(u) = 0\big| = aN$ and $\big|v \in V : \rho(v) = 0\big| = bN$. We also use $\mathcal{I}_N(a,b)$, where $a,b \in T_N$, to denote the set of all such (a,b)-assignments, and let

$$Z_{a,b}(H) = \sum_{\rho\in\mathcal{I}_N(a,b)} \omega(H,\rho)\cdot\gamma^{\Delta'(2-a-b)N} \tag{29}$$

with Δ' defined above. We are interested in the expectation of $Z_{a,b}(H)$ when $\min(a,b) \geq \lambda = 9\times 10^{-5}$:

Lemma 6. *For large enough N and $a,b \in T_N$ such that $\min(a,b) \geq \lambda$, we have*

$$\mathbf{E}_{H\leftarrow\mathcal{H}(N,\Delta)}\Big[Z_{a,b}(H)\Big] \leq \exp\big(1.21\cdot N\big).$$

We delay the proof of Lemma 6 to Appendix F. For now, using Lemma 6 we can impose the following condition on the graph G constructed from the input instance of E2LIN2:

$$\text{For all } i\in[n] \text{ and all } a,b\in T_{d_i m} \text{ with } \min(a,b)\geq\lambda,\ Z_{a,b}(H_i) \leq \exp\big(1.22\cdot d_i m\big). \tag{30}$$

Using Lemma 6, Markov's inequality and the union bound, it is easy to show that G satisfies this condition with probability $1-\exp(-\Omega(m))$. In the rest of the proof we show that G satisfies (10) whenever it satisfies (30). Lemma 2 then follows immediately.

We assume that G satisfies (30). To prove (10), we randomly sample an assignment σ with probability proportional to $\omega(G,\sigma)$. (10) follows if we can show that σ satisfies (9) with probability $1-o(1)$. For this purpose, we need the following lemmas which show properties that σ satisfies with high probability. Given any set L of vertices in G, we let $N_\sigma(L)=\{v\in L:\sigma(v)=0\}$. Also recall the definition of $U_{i,k}$ and $V_{i,k}$ in the construction of G. Let $k\in[m]$ and x_i be a variable that appears in the kth equation, then

Lemma 7. *Let σ be an assignment drawn according to its weight. Except for probability $\exp(-\Omega(m^{1/3}))$*

$$\big|N_\sigma(U_{i,k})\big| < \frac{1}{1+e}\cdot\Big(1+m^{-1/3}\Big)\cdot\big|U_{i,k}\big|.$$

Proof. Pick any partial assignment σ' over vertices of G except those in $U_{i,k}$. Conditioned on σ', it is easy to see that the values of vertices in $U_{i,k}$ are independent. Each vertex in $U_{i,k}$ has $\Delta+\Delta'$ neighbors, each of which contributes a vertex weight of either $(\beta,1)$ or $(1,\gamma)$. Since $\beta<1/\gamma$, the total weight for assignment 1 is at least $\gamma^{\Delta+\Delta'}\geq e$ times the weight for assignment 0. The lemma follows from the Chernoff bound. □

Given an assignment σ, we use σ_i to denote its partial assignment over vertices in H_i and σ_{-i} to denote its partial assignment over vertices in G except H_i. We let $M_{\sigma_{-i}}(U_i)$ denote the subset of U_i whose unique neighbor outside of H_i is assigned 1. Using Lemma 7 and the union bound, we have

Corollary 3. *Let σ be an assignment drawn according to its weight. Except for probability $\exp(-\Omega(m^{1/3}))$*

$$\big|M_{\sigma_{-i}}(U_i)\big| \geq \left(\frac{e}{1+e}-O\big(m^{-1/3}\big)\right)\cdot\big|U_i\big|. \tag{31}$$

It is also clear that Lemma 7 and Corollary 3 also hold for $V_{i,k}$ and V_i, respectively, by symmetry. Now we are ready to prove Lemma 2. Let $\sigma=(\sigma_i,\sigma_{-i})$ be an assignment drawn from this distribution. Recall the definition of Σ below (9). Then by Corollary 3 we have

$$\Pr\big[\,\sigma\notin\Sigma\,\big]\leq\exp(-\Omega(m^{1/3}))+\Pr\Big[\,\sigma\notin\Sigma\;\Big|\;\sigma_{-i}\text{ satisfies (31) for both }U_i\text{ and }V_i\,\Big] \tag{32}$$

To prove an upper bound for (32) we fix σ_{-i} to be any partial assignment over the vertices of G except those of H_i, which satisfies (31) for both U_i and V_i. Then it suffices to prove that the sum of $\omega(G,\sigma)$ over all $\sigma\in\Sigma$ that are consistent with σ_{-i}, denoted by Z_1, is exponentially larger than the sum of $\omega(G,\sigma)$ over all $\sigma\notin\Sigma$ that are consistent with σ_{-i}, denoted by Z_2.

Let $\omega(\sigma_{-i})$ denote the product of the edge weights in σ_{-i} over all edges that have no vertex in H_i. By the definition of $Z_{a,b}(H)$ in (29), we have

$$Z_2\leq\omega(\sigma_{-i})\sum_{a,b\in T_{d_im}:\,a,b\geq\lambda}Z_{a,b}(H_i)\leq\omega(\sigma_{-i})\cdot(d_im)^2\cdot\exp\big(1.22\cdot d_im\big), \tag{33}$$

where the second inequality follows from (30). To prove a lower bound for Z_1, we let $L=|M_{\sigma_{-i}}(U_i)|$ and $R=\big|M_{\sigma_{-i}}(V_i)\big|$. Consider all assignments σ that are consistent with σ_{-i} and $U_i(\sigma)=0$. This gives us

$$Z_1\geq\omega(\sigma_{-i})\cdot\gamma^{\Delta' L}\cdot(1+\gamma^{\Delta+\Delta'})^R\cdot(\beta^{\Delta'}+\gamma^{\Delta})^{d_im-R}.$$

By plugging in $\gamma^{\Delta+\Delta'}\geq e$, $\gamma^{\Delta}\geq e^{7999/8000}$, as well as the lower bound for R in (31), we get

$$Z_1\geq\omega(\sigma_{-i})\cdot\exp(1.22897\cdot d_im),$$

and the lemma follows from (33).

References

[1] A. Bandyopadhyay and D. Gamarnik. Counting without sampling: Asymptotics of the log-partition function for certain statistical physics models. *Random Structures and Algorithms*, 33:452–479, 2008.

[2] A.A. Bulatov. The complexity of the counting constraint satisfaction problem. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, pages 646–661, 2008.

[3] A.A. Bulatov and M. Grohe. The complexity of partition functions. *Theoretical Computer Science*, 348(2–3):148–186, 2005.

[4] J.-Y. Cai and X. Chen. A decidable dichotomy theorem on directed graph homomorphisms with non-negative weights. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 437–446, 2010.

[5] J.-Y. Cai and X. Chen. Complexity of counting CSP with complex weights. *arXiv:1111.2384*, 2011.

[6] J.-Y. Cai, X. Chen, and P. Lu. Graph homomorphisms with complex values: A dichotomy theorem. In *Proceedings of the 37th International Colloquium on Automata, Languages and Programming*, pages 275–286, 2010.

[7] J.-Y. Cai, X. Chen, and P. Lu. Non-negatively weighted #CSP: An effective complexity dichotomy. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity*, pages 45–54, 2011.

[8] M. Dyer, L.A. Goldberg, and M. Paterson. On counting homomorphisms to directed acyclic graphs. *Journal of the ACM*, 54(6), 2007.

[9] M. Dyer and D. Richerby. On the complexity of #CSP. *Proceedings of the 42nd ACM symposium on Theory of computing*, pages 725–734, 2010.

[10] M.E. Dyer, A.M. Frieze, and M. Jerrum. On counting independent sets in sparse graphs. *SIAM Journal on Computing*, 31:1527–1541, 2002.

[11] M.E. Dyer and C. Greenhill. The complexity of counting graph homomorphisms. In *Proceedings of the 9th International Conference on Random Structures and Algorithms*, pages 260–289, 2000.

[12] A. Galanis, Q. Ge, D. Štefankovič, E. Vigoda, and L. Yang. Improved inapproximability results for counting independent sets in the hard-core model. In *Proceedings of the 15th International Workshop on Randomization and Computation*, pages 567–578, 2011.

[13] L.A. Goldberg, M. Grohe, M. Jerrum, and M. Thurley. A complexity dichotomy for partition functions with mixed signs. *SIAM Journal on Computing*, 39(7):3336–3402, 2010.

[14] L.A. Goldberg, M. Jerrum, and M. Paterson. The computational complexity of two-state spin systems. *Random Structures and Algorithms*, 23(2):133–154, 2003.

[15] J. Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48:798–859, 2001.

[16] P. Hell and J. Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.

[17] M. Jerrum and A. Sinclair. Polynomial-time approximation algorithms for the ising model. *SIAM Journal on Computing*, 22(5):1087–1116, 1993.

[18] L. Li, P. Lu, and Y. Yin. Correlation decay up to uniqueness in spin systems. *arXiv:1111.7064*, 2011.

[19] L. Li, P. Lu, and Y. Yin. Approximate counting via correlation decay in spin systems. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, 2012.

[20] L. Lovász. Operations with structures. *Acta Mathematica Hungarica*, 18:321–328, 1967.

[21] E. Mossel, D. Weitz, and N. Wormald. On the hardness of sampling independent sets beyond the tree threshold. *Probability Theory and Related Fields*, 143:401–439, 2009.

[22] A. Sinclair, P. Srivastava, and M. Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, 2012.

[23] A. Sly. Computational transition at the uniqueness threshold. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 287–296, 2010.

[24] D. Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the 38th Annual ACM symposium on Theory of computing*, pages 140–149, 2006.

A Uniqueness

In this section, we formally define the uniqueness condition and discuss some properties of it.

Definition 1. *Let $\hat{x}$ be the positive fixed point of function*

$$f(x) = \mu \left(\frac{\beta x + 1}{x + \gamma} \right)^d.$$

We say that a tuple (β, γ, μ) exhibits uniqueness on d regular graphs if

$$\left| f'(\hat{x}) \right| = \frac{\mu d (1 - \beta\gamma)(\beta\hat{x} + 1)^{d-1}}{(\hat{x} + \gamma)^{d+1}} = \frac{d(1 - \beta\gamma)\hat{x}}{(\beta\hat{x} + 1)(\hat{x} + \gamma)} < 1.$$

This condition specifies whether the Gibbs measure of the system on the $d + 1$-infinite regular tree is unique or not.

For a fixed tuple of parameters (β, γ, μ), if (β, γ) lies inside of the unit square, that is, $0 < \beta, \gamma \leq 1$, as d increases, the uniqueness will eventually fail. As observed before, in this region the monotonicity property with respect to d holds. That is, there exists some threshold Δ depending on (β, γ, μ) such that, for any $d \geq \Delta$, the uniqueness condition does not hold.

When one of β, γ is larger than 1, to satisfy the uniqueness condition, the dependence between (β, γ, μ) and d gets tricky. There exists a boundary such that, beyond it, the uniqueness always hold. However, within the boundary, when d increases, the uniqueness will eventually fail, but if d gets even larger, the uniqueness condition will hold again.

A.1 Random Regular Bipartite Graphs

A classical gadget to use is random regular bipartite graphs by picking d perfect matchings, which is also the gadget used in this paper. One intuitive reason why this gadget is good is that its local structure looks like a tree. More formally, we can show that the expected behavior of this gadget undergoes a phase transition when the parameters of the system go across the uniqueness boundary. Let $Z_{a,b}$ be the expected weight summing over only subsets of size an and bn assigning 0 on each side respectively. If the system is of uniqueness, then the maximum of $Z_{a,b}$ is achieved at the single point (p^*, p^*); while if the system is of non-uniqueness, $Z_{a,b}$ achieved its maximum at two points (p^+, p^-), (p^-, p^+), where $p^+ > p^-$. So the idea is that we can use these two maximum points to encode two states (two assignments of a variable, two parts of a cut, and so on), and reduce other problems to this.

The main difficult is that we can only prove the above connection by expectation. To make the reduction go through, we need the fact that this is true with high probability when we randomly choose a fixed gadget.

To get such a high concentration result is not easy. For the special case of Hardcore model ($\beta = 0, \gamma = 1$), such a high concentration result was almost obtained after a sequence of work by a careful analysis of its second moment [21, 23, 12]. To get such a result for general two-state spin systems seems beyond the reach of current techniques.

Instead we observed that as d increases, p^+ goes to 1 and p^- goes to 0. Thus, as long as d is large enough, we still manage to get some exponential gap between the weights of points near $(0, 1)$ or $(1, 0)$ and everywhere else that holds with high probability. We further argue that to achieve this gap, the degree d differs with the desired uniqueness threshold by only a constant factor.

A.2 The order of the threshold degree

From the uniqueness condition, we cannot give a close form for the boundary d of uniqueness and non-uniqueness in terms of β, γ, μ in general. But some properties were known and we summarize them as follows.

The following lemma is from [22] and [18].

Lemma 8. *If $\sqrt{\beta\gamma} > \frac{d-1}{d+1}$ (or $d < \frac{1+\sqrt{\beta\gamma}}{1-\sqrt{\beta\gamma}}$), the system is always unique for any external field μ.*

Since

$$\frac{1+\sqrt{\beta\gamma}}{1-\sqrt{\beta\gamma}} = \frac{(1+\sqrt{\beta\gamma})^2}{1-\beta\gamma} \geq \frac{1}{1-\beta\gamma},$$

So the degree bound in our Theorem 1 is tight up to a constant factor.

For $d > \frac{1+\sqrt{\beta\gamma}}{1-\sqrt{\beta\gamma}}$, we define

$$x_1(d) = \frac{-1-\beta\gamma+d(1-\beta\gamma)-\sqrt{(-1-\beta\gamma+d(1-\beta\gamma))^2-4\beta\gamma}}{2\beta}.$$
$$x_2(d) = \frac{-1-\beta\gamma+d(1-\beta\gamma)+\sqrt{(-1-\beta\gamma+d(1-\beta\gamma))^2-4\beta\gamma}}{2\beta}.$$

which are the two positive roots of equation $\frac{d(1-\beta\gamma)x}{(\beta x+1)(x+\gamma)} = 1$. Let $\mu_i(d) = x_i(d)\left(\frac{x_i(d)+\gamma}{\beta x_i(d)+1}\right)^d, i = 1, 2$. The following lemma is from [18].

Lemma 9. *Given $\gamma > \beta > 0$, $\beta\gamma < 1$ and $\sqrt{\beta\gamma} \leq \frac{d-1}{d+1}$. The system described by (β, γ, μ) exhibits uniqueness on d regular graph iff $\mu < \mu_1(d)$ or $\mu > \mu_2(d)$.*

When $0 \leq \beta, \gamma < 1$ are treated as constants and μ goes to zero or infinite, we can get the dependent of d in terms of μ as follows.

- When d is large, $x_1(d)$ is very small, $\mu_1(d) = x_1(d)\left(\frac{x_1(d)+\gamma}{\beta x_1(d)+1}\right)^d$ is in the order of $d\gamma^d$. To get non-uniqueness for very small μ, we need that $\mu > \mu_1(d)$. This gives the bound $\frac{\log\mu}{\log\gamma}$.
- When d is large, $x_2(d)$ is very large, $\mu_2(d) = x_2(d)\left(\frac{x_2(d)+\gamma}{\beta x_2(d)+1}\right)^d$ is in the order of $\frac{d}{\beta^d}$. To get non-uniqueness for very large μ, we need that $\mu < \mu_2(d)$. This gives the bound $-\frac{\log\mu}{\log\beta}$.

Therefore, the dependence of d in terms of μ in Corollary 1 is also tight up to a constant factor given that $0 \leq \beta, \gamma < 1$ are treated as constants.

B Proof of Case 2 in Lemma 1

We set h, Δ and Δ' as follows. We pick h to be a large enough constant so that for any $\Delta^* \geq h/(1-\beta\gamma)$, $\Delta = \lfloor L\Delta^*/(L+1) \rfloor$ and $\Delta' = \lceil \Delta^*/(L+1) \rfloor$ satisfy $\Delta' \geq 7$ and $(\beta\gamma)^{\Delta'} < 2^{-12}$. Since $\gamma \geq \beta$, we also have $\beta^{\Delta'} < 1/64$. By the definition of Δ and Δ', we have $L\Delta' \geq \Delta \geq L(\Delta'-1) - 1$.

The proof follows the same flow as that for Case 1. First of all, we use the same lower bound in (14):

$$Z(G,S) \geq Z^*(G,S) = C^M \cdot D^{m\theta(S)}$$

where $Z^*(G,S)$ is the sum of $\omega(G,\sigma)$ over σ that satisfies (13). $C > 0$ and $D > 1$ are defined as in (15). It then follows from $(\beta\gamma)^{\Delta'} < 2^{-12}$ that $D > 4/(3+2^{-12}) \approx 4/3$.

We also use the same upper bound argument (17) for $Z(G,S)$:

$$Z(G,S) \leq \big(1+o(1)\big) \sum_{\sigma \text{ that satisfies (16)}} \omega(G,\sigma). \tag{34}$$

To prove the same statement as in Lemma 5, we pick an $i \in [n]$ and any partial assignment σ' over vertices of G except those of H_i. Without loss of generality, assume $x_i = 0$ in S. We let σ^* denote the assignment that is consistent with σ' and satisfies $U_i(\sigma^*) = 0$ and $V_i(\sigma^*) = d_i m$, and let σ denote any assignment that is consistent with σ' but violates (16). Then from (18) and (19), we have

$$\frac{\omega(G,\sigma^*)}{\omega(G,\sigma)} \geq \frac{\big(\beta\gamma\big)^{\Delta' d_i m}}{\beta^{\epsilon^2 \Delta d_i m/4}} \geq \left(\frac{\beta^{2\Delta'}}{\beta^{3(\Delta'-1)-(\epsilon^2/4)}}\right)^{d_i m} \geq \left(\frac{1}{\beta^{\Delta'/2}}\right)^{d_i m} > 2^{3 d_i m}$$

The upper bound (34) then follows directly. To finish the proof we define T_σ similarly for each assignment σ that satisfies (16). By combining (17), (20) and (21) we get the same upper bound (22) for $Z(G,S)$. It also follows from $\beta \leq \gamma^L$ that $D > 1/(4\gamma^{\Delta+\Delta'})$. Then (7) is proven by plugging in $D \approx 4/3$ and $\epsilon = 10^{-4}$.

C Spin Systems with External Field

It is is easy to verified that this $Z_{\mathbf{A},\mu}(G)$ can be written as

$$Z_{\mathbf{A},\mu}(G) = \sum_{\sigma: V \to \{0,1\}} \mu^{s(\sigma)} \cdot \beta^{t_0(\sigma)} \cdot \gamma^{t_1(\sigma)}$$

where we use $s(\sigma)$ to denote the number of $v \in V$ with $\sigma(v) = 0$; $t_0(\sigma)$ to denote the number of $(u,v) \in E$ with $\sigma(u) = \sigma(v) = 0$; and $t_1(\sigma)$ to denote the number of $(u,v) \in E$ with $\sigma(u) = \sigma(v) = 1$.

Let $t_2(\sigma)$ denote the number of edges whose two ends are assigned different spin states in σ. For a regular graph of degree d, we have $d \cdot s(\sigma) = 2t_0(\sigma) + t_2(\sigma)$, and $t_0(\sigma) + t_1(\sigma) + t_2(\sigma) = |E|$. Thus we can get

$$s(\sigma) = \frac{2t_0(\sigma) + t_2(\sigma)}{d} = \frac{|E| + t_0(\sigma) - t_1(\sigma)}{d}$$

and rewrite $Z_{\mathbf{A},\mu}(G)$ as

$$Z_{\mathbf{A},\mu}(G) = \mu^{\frac{|E|}{d}} \sum_{\sigma: V \to \{0,1\}} \mu^{\frac{t_0(\sigma) - t_1(\sigma)}{d}} \cdot \beta^{t_0(\sigma)} \cdot \gamma^{t_1(\sigma)} = \mu^{\frac{|E|}{d}} \sum_{\sigma: V \to \{0,1\}} (\beta \mu^{\frac{1}{d}})^{t_0(\sigma)} \cdot \left(\frac{\gamma}{\mu^{\frac{1}{d}}}\right)^{t_1(\sigma)}.$$

The global factor $\mu^{\frac{|E|}{d}}$ can be computed in polynomial time and the summation part can be read as the partition function on the same graph with new parameter $(\beta', \gamma', \mu') = (\beta\mu^{\frac{1}{d}}, \frac{\gamma}{\mu^{\frac{1}{d}}}, 1)$. This is a two-state spin system without external field and we can apply Theorem 1 directly to get Corollary 1. To apply Theorem 2

we need to further impose that the degree satisfy a certain relationship with the weight after transformation. This explains the conditions specified in Corollary 2.

D Proof of Lemma 3

Let S be an assignment over the n variables, and we use the same lower bound

$$Z_\Sigma(G,S) \geq Z^*(G,S) = C^{nm} \cdot w^{n\theta(S)}$$

with $Z^*(G,S)$ defined in (14) and C, D defined in (15). It is a lower bound for $Z_\Sigma(G,S)$ because every σ that satisfies (13) is in Σ by definition. Since $\beta\gamma < 1$, we have $C > 0$ and $D > 1$.

Now we give an upper bound for $Z_\Sigma(G,S)$. For each σ in the sum $Z^*(G,S)$, we use Q_σ to denote the following set of assignments σ' in the sum of $Z_\Sigma(G,S)$: for each $i \in [n]$, if $x_i = 0$ in S then σ' agrees with σ on V_i and $|U_i(\sigma')| \leq \lambda d_i m$ (while $|U_i(\sigma)| = 0$); or if $x_i = 1$ then σ' agrees with σ on U_i and $|V_i(\sigma')| \leq \lambda d_i m$ (while $|V_i(\sigma)| = 0$). It is easy to show that $\{Q_\sigma\}$ is a partition of the assignments in $Z^*(G,S)$. Moreover, as in (20) we can show that $|Q_\sigma| \leq m^{2n} \cdot e^{2H(\lambda)M}$. Also every σ' in Q_σ has weight $\omega(G,\sigma') \leq \omega(G,\sigma)$ because flipping a bit from 1 to 0 cannot increase the weight. Finally, we get the following bound for $Z_\Sigma(G,S)$:

$$Z_\Sigma(G,S) \leq Z^*(G,S) \cdot m^{2n} \cdot e^{2H(\lambda)M}$$

To finish the proof, we plug in $\lambda = 9 \times 10^{-5}$ and $H(\lambda) < 0.000929$ to compare $2H(\lambda)$ with $0.04 \ln D$. Recall the definition of D in (15). Use the assumption that $(\beta\gamma)^{\Delta'} \leq 1/e$ and $\gamma^{\Delta+\Delta'} < e\gamma < 1.001e$, we get $\ln D = 0.04673$ and $0.04 \ln D > 2H(\lambda)$. This finishes the proof of the lemma.

E Proof of Lemma 4

Let H be a random d-regular bipartite graph generated by picking d perfect matching between U_n and V_n uniformly at random. We prove the following useful lemma:

Lemma 10. *There exists a positive constant $c > 0$. Given any $A \subseteq U_n$ and $B \subseteq V_n$ with $|A| = an$ and $|B| = bn$, where $b \geq a \geq 10^{-4}$, we have*

$$\Pr\Big[\text{the number of edges between } A \text{ and } B \text{ in } H \leq abdn/4\Big] \leq 2^{-cdn}.$$

With this lemma, we can then choose a large enough d and apply union bound on A and B.

Proof. The intuition is that H is drawn from a distribution that is very similar to $G(n,d)$.

We use $u_1, \ldots, u_{an}$ to denote the vertices in A. For each $k \in [d]$ and $i \in [an]$, we use $Y_{k,i}$ to denote the random $\{0,1\}$-variable such that $Y_{k,i} = 1$ if u_i is matched with a vertex in B in the kth perfect matching; and is 0 otherwise. From this, we have

$$\Pr\Big[\text{the number of edges between } A \text{ and } B \text{ in } H \leq abdn/4\Big] = \Pr\Big[\sum_{k\in[d], i\in[an]} Y_{k,i} \leq abdn/4\Big].$$

However, the variables $Y_{k,1}, \ldots, Y_{k,an}$ are clearly not independent.

To deal with this issue, we introduce the following *independent* random $\{0,1\}$-variables $X_{k,i}$, for each $k \in [d]$ and $i \in [an]$. Here $X_{k,i} = 1$ with probability

$$\rho_i = \frac{bn - (i-1)}{n}$$

and $X_{k,i} = 0$ with probability $1 - \rho_i$. It is easy to see that $Y_{k,i}$ dominates $X_{k,i}$: For all k and i, we have

$$\Pr\Big[Y_{k,i} = 1 \Big| Y_{k,1} \cdots Y_{k,i-1}\Big] \geq \rho_i = \Pr\Big[X_{k,i} = 1\Big].$$

As a result, we can show that

$$\Pr\Big[\sum Y_{k,i} \geq abdn/4\Big] \geq \Pr\Big[\sum X_{k,i} \geq abdn/4\Big]. \tag{35}$$

To see this, fixing k, we generate $Y_{k,1}, \ldots, Y_{k,an}$ and $X_{k,1}, \ldots, X_{k,an}$ jointly as follows: for i from 1 to an, pick a $r_i \in [0,1]$ uniformly at random, then

$$Y_{k,i} = 1 \text{ if } r_i \leq \Pr\Big[Y_{k,i} = 1 \Big| Y_{k,1} \cdots Y_{k,i-1}\Big];$$

and

$$X_{k,i} = 1 \text{ if } r_i \leq \rho_i.$$

It is easy to see that $Y_{k,i} \geq X_{k,i}$ for all $i \in [an]$ and (35) follows.

By (35), it now suffices to prove an upper bound for

$$\Pr\Big[\sum X_{k,i} \leq abdn/4\Big].$$

We can now use the Chernoff bound. First, the expectation is

$$\mu = \sum_{k,i} \mathbf{E}\big[X_{k,i}\big] = d \cdot \sum_i \rho_i \geq abdn/2.$$

By the Chernoff bound (and setting $\delta = 1/2$), we have

$$\Pr\Big[\sum X_{k,i} \leq abdn/4\Big] \leq \Pr\Big[\sum X_{k,i} \leq (1-\delta)\mu\Big] \leq \exp(-\mu\delta^2/2).$$

The lemma then follows from

$$\exp(-\mu\delta^2/2) \leq \exp\left(-\frac{abdn}{2}\left(\frac{1}{2}\right)^2 \cdot \frac{1}{2}\right) \leq \exp\left(-\frac{10^{-8}dn}{2}\left(\frac{1}{2}\right)^2 \cdot \frac{1}{2}\right)$$

by setting c to be a small enough positive constant. □

F Proof of Lemma 6

We recall the definition

$$Z_{a,b}(H) = \sum_{\rho \in \mathcal{I}_n(a,b)} \omega(H,\rho) \cdot \gamma^{\Delta'(2-a-b)N}.$$

We want to compute the expectation $\mathbf{E}_{H \leftarrow \mathcal{H}(n,\Delta)}\Big[Z_{a,b}(H)\Big]$. Since the distribution $\mathcal{H}(N,\Delta)$ is totally symmetric, each term in the summation of $Z_{a,b}(H)$ has the exact the same expectation. Thus, we have

$$\mathbf{E}_{H \leftarrow \mathcal{H}(N,\Delta)}\Big[Z_{a,b}(H)\Big] = \gamma^{\Delta'(2-a-b)N}\binom{N}{aN}\binom{N}{bN}\left(\sum_{k \in T_N, a+b-1 \leq k \leq \min(a,b)} \frac{\beta^{kN}\gamma^{(1-a-b+k)N}\binom{bN}{kN}\binom{(1-b)N}{(a-k)N}}{\binom{N}{aN}}\right)^{\Delta}.$$

Since we will only care about the exponent of $\mathbf{E}_{H \leftarrow \mathcal{H}(n,\Delta)}\Big[Z_{a,b}(H)\Big]$ and the summation of k is only

over linear number of terms, we can replace this summation by maximum without change the leading term of the exponent. Form above, we conclude that the coefficient of the exponent $\Psi(a,b)$ of the expectation $\mathbf{E}_{H\leftarrow\mathcal{H}(n,\Delta)}\Big[Z_{a,b}(H)\Big] = \exp(\Psi(a,b)N)$ is of the following form, in which the function $H(x) = -x\ln x - (1-x)\ln(1-x)$:

$$\begin{aligned}
& \Psi(a,b) \\
= \; & \max_k (2-a-b)\Delta'\ln\gamma + H(a) + H(b) + \Delta(k\ln\beta + (1-a-b+k)\ln\gamma + bH(\frac{k}{b}) + (1-b)H(\frac{a-k}{1-b}) - H(a)) \\
= \; & \max_k \Delta'\ln\gamma + (1-a-b)(\Delta+\Delta')\ln\gamma + H(a) + H(b) + \Delta(k\ln(\beta\gamma) + bH(\frac{k}{b}) + (1-b)H(\frac{a-k}{1-b}) - H(a)) \\
\leq \; & \max_k \frac{1}{c} + (1-a-b)\frac{c}{c-1} + H(a) + H(b) + (c-1)(k\Delta'\ln(\beta\gamma) + bH(\frac{k}{b}) + (1-b)H(\frac{a-k}{1-b}) - H(a)) \\
\leq \; & \max_k \frac{1}{c} + (1-a-b)\frac{c}{c-1} + H(a) + H(b) + (c-1)(-k + bH(\frac{k}{b}) + (1-b)H(\frac{a-k}{1-b}) - H(a))
\end{aligned}$$

For the last formula, we can use computer to get that the value is less than 1.21 when we require that $\min(a,b) \geq \lambda = 9\times 10^{-5}$.