

#BIS-Hardness for 2-Spin Systems on Bipartite Bounded Degree Graphs in the Tree Non-uniqueness Region*

Jin-Yi Cai¹, Andreas Galanis², Leslie Ann Goldberg², Heng Guo¹, Mark Jerrum³, Daniel Štefankovič⁴, and Eric Vigoda⁵

- 1 University of Wisconsin-Madison, Madison, WI, 53706, USA
jyc@cs.wisc.edu, hguo@cs.wisc.edu
- 2 University of Oxford, Wolfson Building, Parks Road, Oxford, OX1 3QD, UK
andreas.galanis@cs.ox.ac.uk, leslie.goldberg@cs.ox.ac.uk
- 3 Queen Mary, University of London, Mile End Road, London E1 4NS, UK
m.jerrum@qmul.ac.uk
- 4 University of Rochester, Rochester, NY, 14627, USA
stefanko@cs.rochester.edu
- 5 Georgia Institute of Technology, Atlanta, GA, 30332, USA
vigoda@cc.gatech.edu

Abstract

Counting independent sets on bipartite graphs (#BIS) is considered a canonical counting problem of intermediate approximation complexity. It is conjectured that #BIS neither has an FPRAS nor is as hard as #SAT to approximate. We study #BIS in the general framework of two-state spin systems in bipartite graphs. Such a system is parameterized by three numbers (β, γ, λ) , where β (respectively γ) represents the weight of an edge (or “interaction strength”) whose end points are of the same 0 (respectively 1) spin, and λ is the weight of a 1 vertex, also known as an “external field”. By convention, the edge weight with unequal 0/1 end points and the vertex weight with spin 0 are both normalized to 1. The partition function of the special case $\beta = 1$, $\gamma = 0$, and $\lambda = 1$ counts the number of independent sets. We define two notions, *nearly-independent phase-correlated spins* and *symmetry breaking*. We prove that it is #BIS-hard to approximate the partition function of any two-spin system on bipartite graphs supporting these two notions.

As a consequence, we show that #BIS on graphs of degree at most 6 is as hard to approximate as #BIS without degree bound. The degree bound 6 is the best possible as Weitz presented an FPTAS to count independent sets on graphs of maximum degree 5. This result extends to the hard-core model and to other anti-ferromagnetic two-spin models. In particular, for all antiferromagnetic two-spin systems, namely those satisfying $\beta\gamma < 1$, we prove that when the infinite $(\Delta - 1)$ -ary tree lies in the non-uniqueness region then it is #BIS-hard to approximate the partition function on bipartite graphs of maximum degree Δ , except for the case $\beta = \gamma$ and $\lambda = 1$. The exceptional case is precisely the antiferromagnetic Ising model without an external field, and we show that it has an FPRAS on bipartite graphs. Our inapproximability results match the approximability results of Li et al., who presented an FPTAS for general graphs of maximum degree Δ when the parameters lie in the uniqueness region.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Spin systems, approximate counting, complexity, #BIS-hardness, phase transition

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2014.582

* Full version [8].

1 Introduction

There has been great progress in classifying the complexity of counting problems recently. One important success is for counting constraint satisfaction problems ($\#CSP$), where a sweeping complexity dichotomy is proved [1, 11, 3]. While the landscape of exact counting becomes clearer, the complexity of approximate counting remains mysterious. Two typical classes of problems have been identified: 1) those that have a fully polynomial-time randomized approximation scheme (FPRAS), and 2) those that are $\#SAT$ -hard with respect to approximation preserving reductions (AP-reductions) [9]. If $NP \neq RP$ then $\#SAT$ admits no FPRAS¹ [32], and therefore neither does any problem in the second class. These two classes are analogous to P-time tractable vs. NP-hard decision or optimization problems.

Interestingly, in approximate counting, there has emerged a third distinct class of natural problems, which seems to be of intermediate complexity. It is conjectured [9] that the problems in this class do not have an FPRAS but that they are not as hard as $\#SAT$ to approximate. A canonical problem in this class has been identified, which is to count the number of independent sets in a bipartite graph ($\#BIS$). Despite many attempts, nobody has found an FPRAS for $\#BIS$ or an AP-reduction from $\#SAT$ to $\#BIS$. The conjecture is that neither exists. Mossel et al. [27] showed that the Gibbs sampler for sampling independent sets in bipartite graphs mixes slowly even on bipartite graphs of degree at most 6. Another interesting attempted Markov Chain for $\#BIS$ by Ge and Stefankovic [15] was also shown later to be slowly mixing by Goldberg and Jerrum [18].

$\#BIS$ plays an important role in classifying counting problems with respect to approximation. A trichotomy theorem is shown for the complexity of approximately solving unweighted Boolean counting CSPs, where in addition to problems that are solvable by FPRASes and those that are AP-reducible from $\#SAT$, there is the intermediate class of problems which are equivalent to $\#BIS$ [10]. Many counting problems are shown to be $\#BIS$ -hard and hence are conjectured to have no FPRAS [2, 7], including estimating the partition function of the the ferromagnetic Potts model [19]. Moreover, under AP-reductions $\#BIS$ is complete in a logically defined class of problems, called $\#RHH_1$, to which an increasing variety of problems have been shown to belong. Other typical complete problems in $\#RHH_1$ include counting the number of downsets in a partially ordered set [9] and computing the partition function of the ferromagnetic Ising model with local external fields [17].

The problem of counting independent sets ($\#IS$) can be viewed as a special case in the general framework of spin systems, which originated from statistical physics to model interactions between neighbors on graphs. In this paper, we focus on two-state spin systems. In general such a system is parameterized by edge weights $\beta, \gamma \geq 0$ and a vertex weight $\lambda > 0$. An instance is a graph $G = (V, E)$. A configuration σ is a mapping $\sigma : V \rightarrow \{0, 1\}$ from vertices to (two) spins. The weight $w(\sigma)$ of a configuration σ is given by

$$w(\sigma) = \beta^{m_0(\sigma)} \gamma^{m_1(\sigma)} \lambda^{n_1(\sigma)} \quad (1)$$

where $m_0(\sigma)$ is the number of $(0, 0)$ edges given by the configuration σ , $m_1(\sigma)$ is the number of $(1, 1)$ edges, and $n_1(\sigma)$ is the number of vertices assigned 1. We are interested in computing the partition function, which is defined by

$$Z_G(\beta, \gamma, \lambda) = \sum_{\sigma: V \rightarrow \{0, 1\}} w(\sigma). \quad (2)$$

¹ In fact, Zuckerman proves a stronger result—there is no FPRAS for the logarithm of the number of satisfying assignments unless $NP=RP$.

The partition function is the normalizing factor of the Gibbs distribution, which is the distribution in which a configuration σ is drawn with probability $\Pr_{G;\beta,\gamma,\lambda}(\sigma) = \frac{w(\sigma)}{Z_G(\beta,\gamma,\lambda)}$. The spin system is called *ferromagnetic* if $\beta\gamma > 1$ and *antiferromagnetic* if $\beta\gamma < 1$. In particular, when $\beta = \gamma$, such a system is the famous *Ising model*, and when $\beta = 1$ and $\gamma = 0$, it is the *hard-core gas model*, the partition function of which counts independent sets when $\lambda = 1$. The external field λ is typically referred to as the activity or fugacity of the hard-core model.

Approximating the partition function of the hard-core model is especially well studied. We now know that the complexity transition from easy to hard corresponds exactly to the uniqueness of the Gibbs measure in infinite $(\Delta - 1)$ -ary trees \mathbb{T}_Δ (for details of these notions, see [16]). Notice that $(\Delta - 1)$ -ary trees are graphs of maximum degree Δ , hence our use of the notation \mathbb{T}_Δ . On the algorithmic side, Weitz presented a fully polynomial-time approximation scheme (FPTAS) for the hard-core gas model on graphs of maximum degree Δ when uniqueness holds [31]. On the other hand, Sly showed that the approximation problem is $\#\text{SAT}$ -hard for a small interval beyond the uniqueness threshold [29]. Building on their work, it is now established that for all antiferromagnetic 2-spin systems there is an FPTAS for graphs of maximum degree Δ up to the uniqueness threshold [25] (see also [24, 28]), whereas non-uniqueness implies $\#\text{SAT}$ -hardness under AP-reductions on Δ -regular graphs [30] (see also [4, 13]). The only place that remains unclear is exactly at the uniqueness threshold.

A key feature of spin systems in the antiferromagnetic non-uniqueness region is the ability to support a gadget with many vertices whose spins are highly correlated with the phase of the gadget (which is either $+$ or $-$), but are nearly independent among themselves conditioned on that phase. Such a gadget was used by Sly [29] to show inapproximability of the partition function of the hard-core model when λ is just above the uniqueness threshold. A different gadget with similar properties was used by Goldberg et al. [20] to show inapproximability on a planar graph when λ is much larger. We abstract this notion of nearly-independent phase-correlated spins. It is this feature that enables us to reduce from $\#\text{SAT}$ to approximating the partition function of antiferromagnetic two-spin systems in the non-uniqueness region.

Restricted to bipartite graphs, it appears that supporting nearly-independent phase-correlated spins alone is not enough to imply $\#\text{BIS}$ -hardness. It was shown that Sly's gadget is applicable to the antiferromagnetic Ising model without an external field by Galanis et al. [13]. However, such a system has an FPRAS on bipartite graphs. The reason is that this system is perfectly symmetric on bipartite graphs and therefore can be translated into a ferromagnetic Ising system, whose partition function can be approximated using the FPRAS of Jerrum and Sinclair [22] (see Corollary 13 in the full version [8] for details). To get around this perfectly symmetric case, we introduce the second key concept called symmetry breaking. Symmetry breaking does not refer to whether the parameters of the model are symmetric, but rather whether a gadget can be constructed with a distinguished degree 1 vertex that has a certain asymmetry. Formal definitions of the two notions – nearly-independent phase-correlated spins and symmetry breaking – can be found in Section 3. Our main technical theorem is the following.

► **Theorem 1.** *Suppose a tuple of parameters $(\beta, \gamma, \lambda, \Delta)$ with $\beta\gamma \neq 1$ and $\Delta \geq 3$ supports nearly-independent phase-correlated spins and symmetry-breaking. Then the partition function (2) of two-spin systems (β, γ, λ) is $\#\text{BIS}$ -hard to approximate on bipartite graphs with maximum degree Δ .*

Previous hardness proofs for the problem $\#\text{IS}$ and for the problem of estimating the partition function of antiferromagnetic 2-spin systems typically reduce from MAX-CUT or from the problem of counting certain types of cuts [21, 29, 30]. However such a technique

sheds little light in the bipartite setting as cut problems are trivial on bipartite graphs. Reductions between #BIS-equivalent problems typically involve transformations that “blow up” vertices and edges into sets of vertices that are completely connected, so they do not apply to bounded-degree graphs either.

A key property of Sly’s gadget is that either phase occurs with probability bounded below by an inverse polynomial. This bound is sufficient in Sly’s setting to reduce from MAX-CUT, but it is not enough to construct AP-reductions for our use. We resolve this issue by introducing a balancing construction. The construction takes two copies of a gadget with nearly-independent phase-correlated spins, and produces a new gadget with similarly-correlated spins, but in the new gadget the two phases occur with probability close to 1/2.

The proof of Theorem 1 utilizes an intermediate problem, that is, computing the partition function of antiferromagnetic Ising systems with non-uniform external fields on bipartite graphs. A non-uniform external field means that the instance specifies a subset of vertices on which the external field acts. A 2-spin system with a non-uniform external field is very similar to a Boolean #CSP with one binary symmetric non-negative valued function (corresponding to edge weights) and one unary non-negative valued function (corresponding to vertex weights) (see, for example [5]).

Our reduction implements an external field, and this is where symmetry breaking comes into play. As discussed earlier, the partition function of Ising model without an external field has an FPRAS, so the symmetry breaking gadget seems necessary. In fact, we show that symmetry breaking holds for all 2-spin systems except for the Ising model without an external field or degenerate systems (i. e., systems satisfying $\beta\gamma = 1$). We also prove that all antiferromagnetic 2-spin systems support nearly-independent phase-correlated spins in the non-uniqueness region. Finally, applying Theorem 1 yields our main result:

► **Theorem 2.** *For all tuples of parameters $(\beta, \gamma, \lambda, \Delta)$ with $\Delta \geq 3$ and $\beta\gamma < 1$, except for the case $(\beta = \gamma, \lambda = 1)$, if the infinite Δ -regular tree \mathbb{T}_Δ is in the non-uniqueness region then approximating the partition function (2) on bipartite graphs with maximum degree Δ is #BIS-equivalent under AP-reductions.*

Let us now survey the approximability picture that this theorem helps establish. For general antiferromagnetic 2-spin models with soft constraints (i. e., $\beta\gamma > 0$), non-uniqueness holds if and only if $\sqrt{\beta\gamma} < \frac{\Delta-2}{\Delta}$ and $\lambda \in (\lambda_1, \lambda_2)$ for some critical values λ_1 and λ_2 depending on β, γ , and Δ (see [25, Lemma 5]). Hence, for all $\beta, \gamma > 0$ where $\beta\gamma < 1$, and all $\Delta \geq 3$ the following holds:

1. If $\sqrt{\beta\gamma} > \frac{\Delta-2}{\Delta}$, for all λ , there is an FPTAS to approximate the partition function for Δ -regular graphs [28, 25] (this extends to graphs of maximum degree Δ in an appropriate sense, see [25] for details).
2. If $\sqrt{\beta\gamma} < \frac{\Delta-2}{\Delta}$, then there exists $0 < \lambda_1 < \lambda_2$ so that:
 - a. For all $\lambda \notin [\lambda_1, \lambda_2]$, there is an FPTAS to approximate the partition function for Δ -regular graphs [28, 25] (this again extends in an appropriate sense to graphs of maximum degree Δ [25]).
 - b. For all $\lambda \in (\lambda_1, \lambda_2)$, it is #SAT-hard to approximate the partition function on Δ -regular graphs [30].
 - c. For all $\lambda \in (\lambda_1, \lambda_2)$, it is #BIS-hard to approximate the partition function on bipartite graphs of maximum degree Δ (Theorem 2 in this paper).

For the particular case of the hard-core model the critical value (i. e., critical activity $\lambda_c(\Delta)$) is more easily stated. For the hard-core model (i. e., $\beta = 0$ and $\gamma = 1$) Kelly [23]

showed that non-uniqueness holds on \mathbb{T}_Δ if and only if $\lambda > \lambda_c(\Delta) := \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^\Delta}$. As a consequence we get the following corollary for the hard-core model.

► **Corollary 3.** *For all $\Delta \geq 3$, all $\lambda > \lambda_c(\Delta) = \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^\Delta}$, it is #BIS-hard to approximate the partition function of the hard-core model on bipartite graphs of maximum degree Δ .*

We also get a corollary concerning the more general partition function as long as β and γ are less than 1 and the degree bound Δ is sufficiently large. For β and γ satisfying $0 < \beta, \gamma < 1$ and any $\lambda > 0$, there exists a Δ such that (β, γ, λ) is in the non-uniqueness region of \mathbb{T}_Δ [25, Lemma 21.2]. This implies the following corollary.

► **Corollary 4.** *For every $0 < \beta, \gamma < 1$ and $\lambda > 0$, there exists a Δ such that the 2-spin system with parameters β, γ and with uniform or non-uniform external field λ on bipartite graphs with degree bound Δ is #BIS-equivalent under AP-reductions, except when $\beta = \gamma$ and $\lambda = 1$, in which case it has an FPRAS.*

More generally, for antiferromagnetic 2-spin systems we get the following picture for the complexity of approximating the partition function on general graphs. As usual there is a difficulty classifying the complexity of approximating the partition function at the boundary between uniqueness and non-uniqueness. To address this issue, for parameters $(\beta, \gamma, \lambda, \Delta)$, [25] define a notion of *up-to- Δ unique* which is equivalent to the parameters lying in the interior of the uniqueness region for the infinite $(d-1)$ -ary tree \mathbb{T}_d for all $3 \leq d \leq \Delta$ (see Definition 7 in [25]). Moreover, the parameters (β, γ, λ) satisfy *∞ -strict-uniqueness* if it is up-to- ∞ unique.² On the other side, we say the parameters (β, γ, λ) satisfy *∞ -non-uniqueness* if for some $\Delta \geq 3$ the tree \mathbb{T}_Δ has non-uniqueness. The only gap between the notions of ∞ -strict-uniqueness and ∞ -non-uniqueness is the case when the parameters (β, γ, λ) are at the uniqueness/non-uniqueness threshold of \mathbb{T}_Δ for some Δ .

The following result detailing the complexity for general graphs is now established.

► **Corollary 5.** *For all tuples of parameters (β, γ, λ) with $\beta\gamma < 1$, the following holds:*

1. *If the parameters satisfy ∞ -strict-uniqueness then there is a FPTAS for the partition function for all graphs [25, Theorem 2].*
2. *If the parameters satisfy ∞ -non-uniqueness then:*
 - a. *it is #SAT-hard to approximate the partition function on graphs [30].*
 - b. *it is #BIS-hard to approximate the partition function on bipartite graphs (Theorem 2 in this paper).*

A recent paper of Liu et al. [26] shows that our Theorem 1 can also be used to analyse the complexity of *ferromagnetic* partition functions (where $\beta\gamma > 1$). In particular, it uses Theorem 1 to show #BIS-hardness for approximating the partition function for ferromagnetic 2-spin systems when $\beta \neq \gamma$ for sufficiently large external field λ . An interesting problem that remains open is to prove #BIS-hardness for the entire non-uniqueness region for ferromagnetic 2-spin systems with $\beta \neq \gamma$.

2 Approximation-Preserving Reductions and #BIS

We are interested in the complexity of approximate counting. Let Σ be a finite alphabet. We want to approximate the value of a function $f : \Sigma^* \rightarrow \mathbb{R}$. A *randomized approximation*

² To be precise, the notion of ∞ -strict-uniqueness is called *universally unique* in [25, Definition 7]).

scheme is an algorithm that takes an instance $x \in \Sigma^*$ and a rational error tolerance $\varepsilon > 0$ as inputs, and outputs a rational number z such that, for every x and ε , $\Pr[e^{-\varepsilon}f(x) \leq z \leq e^\varepsilon f(x)] \geq \frac{3}{4}$. A *fully polynomial randomized approximation scheme* (FPRAS) is a randomized approximation scheme which runs in time bounded by a polynomial in $|x|$ and ε^{-1} . Note that the quantity $\frac{3}{4}$ can be changed to any value in the interval $(\frac{1}{2}, 1)$ or even $1 - 2^{-n^c}$ for a problem of size n without changing the set of problems that have fully polynomial randomized approximation schemes since the higher accuracy can be achieved with only polynomial delay by taking a majority vote of multiple samples.

Dyer *et al.* [9] introduced the notion of approximation-preserving reductions. Suppose f and g are two functions from Σ^* to \mathbb{R} . An *approximation-preserving reduction* (AP-reduction) from f to g is a randomized algorithm \mathcal{A} to approximate f using an oracle for g . The algorithm \mathcal{A} takes an input $(x, \varepsilon) \in \Sigma^* \times (0, 1)$, and satisfies the following three conditions: (i) every oracle call made by \mathcal{A} is of the form (y, δ) , where $y \in \Sigma^*$ is an instance of g , and $0 < \delta < 1$ is an error bound satisfying $\delta^{-1} \leq \text{poly}(|x|, \varepsilon^{-1})$; (ii) the algorithm \mathcal{A} meets the specification for being a randomized approximation scheme for f whenever the oracle meets the specification for being a randomized approximation scheme for g ; (iii) the run-time of \mathcal{A} is polynomial in $|x|$ and ε^{-1} .

If an AP-reduction from f to g exists, we write $f \leq_{\text{AP}} g$, and say that f is *AP-reducible* to g . If $f \leq_{\text{AP}} g$ and $g \leq_{\text{AP}} f$, then we say that f and g are *AP-interreducible* or *AP-equivalent*, and write $f \equiv_{\text{AP}} g$. The problem #BIS is defined as follows.

Name. #BIS.

Instance. A bipartite graph B .

Output. The number of independent sets in B .

In this paper, we are interested in 2-spin systems over bounded degree bipartite graphs parametrized by a tuple (β, γ, λ) . We say a real number z is *efficiently approximable* if there is an FPRAS for the problem of computing z . Throughout the paper we only deal with non-negative real parameters that are efficiently approximable. For efficiently approximable non-negative real numbers β, γ, λ and a positive integer Δ , we define the problem of computing the partition function of the 2-spin system (β, γ) with external field λ on bipartite graphs of bounded degree Δ , as follows.

Name. BI-(M-)2-SPIN $(\beta, \gamma, \lambda, \Delta)$.

Instance. A bipartite (multi)graph $B = (V, E)$ with degree bound Δ .

Output. The quantity

$$Z_B(\beta, \gamma, \lambda) = \sum_{\sigma: V \rightarrow \{0,1\}} \lambda^{\sum_{v \in V} \sigma(v)} \prod_{(v,u) \in E} \beta^{(1-\sigma(v))(1-\sigma(u))} \gamma^{\sigma(v)\sigma(u)}.$$

Notice that we also introduced a multigraph version of the same problem. It will be useful later. We drop the parameter Δ when there is no degree bound, that is, BI-2-SPIN (β, γ, λ) is the same as BI-2-SPIN $(\beta, \gamma, \lambda, \infty)$.

We also found the notion of non-uniform external field useful in the reductions. The following problems are introduced as intermediate problems. We also introduce a multigraph version, but as intermediate problems we do not need the bounded degree variant.

Name. BI-(M-)NONUNIFORM-2-SPIN (β, γ, λ) .

Instance. A bipartite (multi)graph $B = (V, E)$ and a subset $U \subseteq V$.

Output. The quantity

$$Z_{B,U}(\beta, \gamma, \lambda) = \sum_{\sigma: V \rightarrow \{0,1\}} \lambda^{\sum_{v \in U} \sigma(v)} \prod_{(v,u) \in E} \beta^{(1-\sigma(v))(1-\sigma(u))} \gamma^{\sigma(v)\sigma(u)}.$$

3 Key Properties of the Gadget

In this section we define two key concepts: nearly-independent phase-correlated spins and symmetry breaking.

We first describe the basic setup of a certain gadget. For positive integers Δ , t and n where n is even and is at least $2t$, let T^- and T^+ be disjoint vertex sets of size t and let V^- be a size- $n/2$ superset of T^- and V^+ be a size- $n/2$ superset of T^+ which is disjoint from V^- . Let $T = T^- \cup T^+$ and $V(t, n) = V^- \cup V^+$. Let $\mathcal{G}(t, n, \Delta)$ be the set of bipartite graphs with vertex partition (V^-, V^+) in which every vertex has degree at most Δ and every vertex in T has degree at most $\Delta - 1$. We refer to the vertices in T as “terminals”. Vertices in T^+ are “positive terminals” and vertices in T^- are “negative terminals”.

When the gadget G is drawn from $\mathcal{G}(t, n, \Delta)$, we use the notation $T(G)$ to refer to the set of terminals. Each configuration $\sigma: V(t, n) \rightarrow \{0, 1\}$ is assigned a unique phase $Y(\sigma) \in \{-, +\}$. Roughly in our applications of the definitions below the phase of a configuration σ is π if V^π contains more vertices with spin 1 than does $V^{-\pi}$.

- We define measures Q^+ and Q^- . Fix some $0 < q^- < q^+ < 1$. For any positive integer t ,
- Q^+ is the distribution on configurations $\tau: T \rightarrow \{0, 1\}$ such that, for every $v \in T^+$, $\tau(v) = 1$ independently with probability q^+ and, for every $v \in V^-$, $\tau(v) = 1$ independently with probability q^- , and
 - Q^- is the distribution on configurations $\tau: T \rightarrow \{0, 1\}$ such that, for every $v \in T^-$, $\tau(v) = 1$ independently with probability q^+ and, for every $v \in T^+$, $\tau(v) = 1$ independently with probability q^- .

To give a rough sense for the values q^- and q^+ they will correspond to the marginal probabilities of the root of an infinite tree obtained by taking limits of finite trees with appropriate boundary conditions, see Section 7 of the full version [8] for more details.

To prove the #BIS-hardness we need a gadget where the spins of the terminals are drawn from distributions close to Q^+ or Q^- conditioned on the phase $+$ or $-$.

► **Definition 6.** A tuple of parameters $(\beta, \gamma, \lambda, \Delta)$ **supports nearly-independent phase-correlated spins** if there are efficiently-approximable values $0 < q^- < q^+ < 1$ such that the following is true. There are functions $n(t, \varepsilon)$, $m(t, \varepsilon)$, and $f(t, \varepsilon)$, each of which is bounded from above by a polynomial in t and ε^{-1} , and for every t and ε there is a distribution on graphs in $\mathcal{G}(t, n(t, \varepsilon), \Delta)$ such that a gadget $G = (V, E)$ with terminals T can be drawn from the distribution within $m(t, \varepsilon)$ time, and the probability that the following inequalities hold is at least $3/4$:

1. The phases are roughly balanced, i. e.,

$$\Pr_{G; \beta, \gamma, \lambda}(Y(\sigma) = +) \geq \frac{1}{f(t, \varepsilon)} \text{ and } \Pr_{G; \beta, \gamma, \lambda}(Y(\sigma) = -) \geq \frac{1}{f(t, \varepsilon)}. \quad (3)$$

2. For a configuration $\sigma: V \rightarrow \{0, 1\}$ and any $\tau: T \rightarrow \{0, 1\}$,

$$\left| \frac{\Pr_{G; \beta, \gamma, \lambda}(\sigma|_T = \tau \mid Y(\sigma) = +)}{Q^+(\tau)} - 1 \right| \leq \varepsilon \text{ and } \left| \frac{\Pr_{G; \beta, \gamma, \lambda}(\sigma|_T = \tau \mid Y(\sigma) = -)}{Q^-(\tau)} - 1 \right| \leq \varepsilon. \quad (4)$$

In fact, given a gadget with the above property, one can construct a gadget where the phases are (nearly) uniformly distributed as detailed in the following definition.

► **Definition 7.** We say that the tuple of parameters $(\beta, \gamma, \lambda, \Delta)$ supports **balanced** nearly-independent phase-correlated spins if Definition 6 holds with (3) replaced by:

$$\Pr_{G;\beta,\gamma,\lambda}(Y(\sigma) = +) \geq \frac{1-\varepsilon}{2} \text{ and } \Pr_{G;\beta,\gamma,\lambda}(Y(\sigma) = -) \geq \frac{1-\varepsilon}{2}, \tag{5}$$

where ε is quantified as in Definition 6.

In Section 5 of the full version [8], we prove the following lemma, which shows that for essentially all 2-spin systems, Definition 6 implies Definition 7. The lemma is proved by constructing a balanced gadget by combining two unbalanced ones.

► **Lemma 8.** *If the parameter tuple $(\beta, \gamma, \lambda, \Delta)$ with $\beta\gamma \neq 1$ supports nearly-independent phase-correlated spins, then it supports balanced nearly-independent phase-correlated spins.*

The main technical result for proving #BIS-hardness for 2-spin antiferromagnetic systems in the tree non-uniqueness region is the following lemma, which is proved in Section 7 of the full version [8]. The proof is rather technical, and is based on a detailed analysis of Sly’s gadget [29], using ideas from [14]. Once a gadget is constructed, it can be balanced using Lemma 8.

► **Lemma 9.** *For all $\Delta \geq 3$, all $\beta, \gamma, \lambda > 0$ where $\beta\gamma < 1$, if the infinite Δ -regular tree \mathbb{T}_Δ is in the non-uniqueness region then the tuple of parameters $(\beta, \gamma, \lambda, \Delta)$ supports balanced nearly-independent phase-correlated spins.*

The second property of the gadget is the notion of symmetry breaking which is relatively simple.

► **Definition 10.** We say that a tuple of parameters $(\beta, \gamma, \lambda, \Delta)$ **supports symmetry-breaking** if there is a bipartite graph H whose vertices have degree at most Δ which has a distinguished degree-1 vertex v_H such that $\Pr_{H;\beta,\gamma,\lambda}(\sigma_{v_H} = 1) \notin \{0, \lambda/(1 + \lambda), 1\}$.

We will prove in Section 6 of the full version [8] that symmetry breaking holds for all 2-spin models except in two cases (where we have tractability).

► **Lemma 11.** *Assume $\Delta \geq 3$. The parameters $(\beta, \gamma, \lambda, \Delta)$ support symmetry breaking unless (i) $\beta\gamma = 1$ or (ii) $\beta = \gamma$ and $\lambda = 1$.*

Having Lemma 9 and Lemma 11, Theorem 2 is a straightforward consequence of Theorem 1.

4 General Reduction

In this section we prove Theorem 1. We first show how the two notions of “nearly-independent phase-correlated spins” and “symmetry-breaking” lead to #BIS-hardness.

4.1 An Intermediate Problem

The goal of this section is to show that it is #BIS-hard to approximate the partition function of antiferromagnetic Ising models with non-uniform non-trivial external fields on bipartite graphs.

► **Lemma 12.** *For any $0 < \alpha < 1$, $\lambda > 0$ and $\lambda \neq 1$, #BIS \leq_{AP} BI-M-NONUNIFORM-2-SPIN(α, α, λ).*

Proof. By flipping 0 to 1 and 1 to 0 for each configuration σ , we see that BI-M-NONUNIFORM-2-SPIN(α, α, λ) is in fact the same as BI-M-NONUNIFORM-2-SPIN($\alpha, \alpha, 1/\lambda$). Hence we may assume $\lambda < 1$.

Let $M = \begin{pmatrix} \alpha & 1 \\ 1 & \alpha \end{pmatrix}$, and $\begin{pmatrix} \rho_0 \\ \rho_1 \end{pmatrix} = M \begin{pmatrix} 1 \\ \lambda \end{pmatrix} = \begin{pmatrix} \alpha + \lambda \\ 1 + \alpha\lambda \end{pmatrix}$. Note that $\alpha < 1$ and $\lambda < 1$, so $\rho_1 > \rho_0$. Let $B = (V, E)$ be an input to #BIS with $n = |V|$ and $m = |E|$. Let I_B be the number of independent sets of B . Let ε be the desired accuracy of the reduction. We will construct an instance $B' = (V', E')$ with a specified vertex subset $U \subset V'$ for BI-M-NONUNIFORM-2-SPIN(α, α, λ) such that $\exp(-\frac{\varepsilon}{2}) I_B \leq Z_{B',U}(\alpha, \alpha, \lambda)/C \leq \exp(\frac{\varepsilon}{2}) I_B$, where C is a quantity that is easy to approximate. Therefore it suffices to call oracle BI-M-NONUNIFORM-2-SPIN(α, α, λ) on B' with the specified subset U with accuracy $\frac{\varepsilon}{4}$ and approximate C within $\frac{\varepsilon}{4}$.

The construction of B' involves two positive integers t_1 and t_2 . Let t_1 be the least positive integer (depending on n and ε) satisfying the first equation in (6) and let t_2 be the least positive integer depending on n, ε and t_1 satisfying the second equation in (6).

$$\alpha^{2t_1} \leq \frac{\varepsilon}{6 \cdot 2^n} \quad \text{and} \quad \left(\frac{\rho_0}{\rho_1}\right)^{t_2} \leq \frac{\alpha^{t_1 m} \cdot \varepsilon}{6 \cdot 2^{2t_1 m + n}}. \quad (6)$$

Note that both t_1 and t_2 are bounded from above by a polynomial in n and ε^{-1} . Given the integers t_1 and t_2 , the graph B' is constructed as follows. Let $W_v = \{w_{v,j} \mid 1 \leq j \leq t_1 \deg(v)\}$ for each $v \in V$ where $\deg(v)$ is the degree of v in B . Let $U_{v,j} = \{u_{v,j,k} \mid 1 \leq k \leq t_2\}$ for any $v \in V$ and $1 \leq j \leq t_1 \deg(v)$. Let $W = \bigcup_{v \in V} W_v$ and $U = \bigcup_{v \in V} \bigcup_{1 \leq j \leq t_1 \deg(v)} U_{v,j}$. The vertex set of B' is $V' = V \cup U \cup W$. Note that $|W| = 2t_1 m$ and $|U| = 2t_1 t_2 m$.

We add t_1 parallel edges in B' between u and v for each $(u, v) \in E$ and add edges between v and every vertex in W_v , and between $w_{v,j}$ and every vertex in $U_{v,j}$ for each $v \in V$ and $1 \leq j \leq t_1 \deg(v)$. Formally the edge set of B' is $E' = \left(\biguplus_{1 \leq i \leq t_1} E\right) \cup \bigcup_{v \in V} E_v \cup \bigcup_{\substack{v \in V \\ 1 \leq j \leq t_1 \deg(v)}} E_{v,j}$, where \biguplus denotes a disjoint union as a multiset of t_1 copies of E , $E_v = \{(v, w) \mid w \in W_v\}$ and $E_{v,j} = \{(w_{v,j}, u) \mid u \in U_{v,j}\}$ for each v and j .

Let $C = \rho_1^{2t_1 t_2 m} \alpha^{t_1 m}$ and $N = \begin{pmatrix} 1 \\ \alpha^{2t_1} \end{pmatrix}$. For each $\sigma : V \cup W \rightarrow \{0, 1\}$, let $w(\sigma)$ be the contribution to $Z_{B',U}(\alpha, \alpha, \lambda)$ of configurations that are consistent with σ . First consider configurations σ such that $\sigma(w) = 1$ for all $w \in W$. Denote by Σ the set of all such configurations on $V \cup W$. Then for $\sigma \in \Sigma$,

$$w(\sigma) = \rho_1^{t_2 |W|} \prod_{(u,v) \in E} (M_{1,\sigma(u)} M_{\sigma(u),\sigma(v)} M_{\sigma(v),1})^{t_1} = C \prod_{(u,v) \in E} N_{\sigma(u),\sigma(v)}.$$

Let $\Sigma^{ind} \subset \Sigma$ be the subset of configurations which induce an independent set on the vertices V and Z^{ind} be its contribution to $Z_{B',U}(\alpha, \alpha, \lambda)$. Let $\Sigma^{bad} = \Sigma \setminus \Sigma^{ind}$ and Z^{bad} be its contribution. If $\sigma \in \Sigma^{ind}$ then $w(\sigma) = C$. Otherwise, $w(\sigma) \leq \alpha^{2t_1} C$. It implies

$$Z^{ind} = I_B \cdot C \quad \text{and} \quad Z^{bad} \leq 2^n \alpha^{2t_1} C \leq \frac{\varepsilon}{6} \cdot C, \quad (7)$$

since t_1 satisfies Eq. (6). Next consider configurations σ on $V \cup W$ such that $\sigma(w) = 0$ for at least one $w \in W$. Denote this set by Σ' and its contribution by Z^{small} . Then for $\sigma \in \Sigma'$,

$$w(\sigma) \leq \left(\rho_0 \rho_1^{|W|-1}\right)^{t_2} \leq \left(\frac{\rho_0}{\rho_1}\right)^{t_2} \rho_1^{t_2 |W|} = \left(\frac{\rho_0}{\rho_1}\right)^{t_2} \frac{C}{\alpha^{t_1 m}}.$$

It implies $Z^{small} \leq 2^{2t_1 m + n} \left(\frac{\rho_0}{\rho_1}\right)^{t_2} \frac{C}{\alpha^{t_1 m}} \leq \frac{\varepsilon}{6} \cdot C$, since $|\Sigma'| \leq 2^{2t_1 m + n}$ and t_2 satisfies Eq. (6). Using this with Eq. (7) we have

$$Z_{B',U}(\alpha, \alpha, \lambda) = Z^{ind} + Z^{bad} + Z^{small} \leq I_B \cdot C + \frac{\varepsilon}{6} \cdot C + \frac{\varepsilon}{6} \cdot C \leq \exp\left(\frac{\varepsilon}{3}\right) I_B \cdot C,$$

and clearly $Z_{B',U}(\alpha, \alpha, \lambda) \geq I_B \cdot C$. It is also clear that C can be approximated accurate enough given FPRAS's for λ and α . This finishes our proof. \blacktriangleleft

4.2 Simulating the Antiferromagnetic Ising Model

In this section we prove the following lemma.

► Lemma 13. *Suppose β, γ and λ are efficiently approximable reals satisfying $\beta, \gamma \geq 0, \lambda > 0$ and $\beta\gamma \neq 1$. Suppose that Δ is either an integer that is at least 3 or $\Delta = \infty$ (indicating that we do not have a degree bound). If $(\beta, \gamma, \lambda, \Delta)$ supports nearly-independent phase-correlated spins and symmetry breaking, then there exist efficiently approximable $0 < \alpha < 1$ and $\lambda' > 0$ such that $\lambda' \neq 1$ and $\text{BI-M-NONUNIFORM-2-SPIN}(\alpha, \alpha, \lambda') \leq_{AP} \text{BI-2-SPIN}(\beta, \gamma, \lambda, \Delta)$.*

Proof. We prove the antiferromagnetic case first, that is, $\beta\gamma < 1$. α and λ' are chosen as follows. Recall that $M = \begin{pmatrix} \beta & 1 \\ 1 & \gamma \end{pmatrix}$ and $M^+ = \begin{pmatrix} 1-q^- & q^- \\ 1-q^+ & q^+ \end{pmatrix}$. Let $N = M^+M(M^+)^T = \begin{pmatrix} N_{--} & N_{-+} \\ N_{+-} & N_{++} \end{pmatrix}$. Then $\det(N) = (\beta\gamma - 1)(q^+ - q^-)^2 < 0$. Therefore $N_{--}N_{++} < N_{-+}N_{+-}$ and let $\alpha = \frac{N_{--}N_{++}}{N_{-+}N_{+-}} < 1$. Moreover, suppose H is the symmetry breaking gadget with distinguished vertex v_H . Let $\rho = \begin{pmatrix} \rho_0 \\ \rho_1 \end{pmatrix}$ where ρ_i denote $\Pr_{H;\beta,\gamma,\lambda}(\sigma_{v_H} = i)$ for spin $i \in \{0, 1\}$ and $\rho_0 + \rho_1 = 1$. Let $\rho' = \begin{pmatrix} \rho'_0 \\ \rho'_1 \end{pmatrix} = M^+ \begin{pmatrix} \rho_0 \\ \rho_1/\lambda \end{pmatrix}$, and $\lambda' = \frac{\rho'_1}{\rho'_0}$. It is easy to verify that $\lambda' \neq 1$ as $\rho_1 \neq \lambda/(1 + \lambda)$ by the symmetry breaking assumption.

Given $0 < \varepsilon < 1$ and a bipartite multigraph $B = (V, E)$ with a subset $U \subseteq V$ where $|V| = n, |E| = m$, and $|U| = n'$, our reduction first constructs a bipartite graph B' with degree at most Δ . The construction of B' involves a gadget G . Since $(\beta, \gamma, \lambda, \Delta)$ supports nearly-independent phase-correlated spins, by Lemma 8 $(\beta, \gamma, \lambda, \Delta)$ also supports balanced nearly-independent phase-correlated spins. Therefore we draw $G \sim \mathcal{G}(t, n(t, \varepsilon'), \Delta)$ such that Eq. (5) and Eq. (4) hold with probability at least $3/4$, where $t = m + 1$ and $\varepsilon' = \frac{\varepsilon}{8n}$. Assume G satisfies them and otherwise the reduction fails. We will construct B' such that

$$\exp\left(-\frac{\varepsilon}{2}\right) Z_{B,U}(\alpha, \alpha, \lambda') \leq \frac{Z_{B'}}{(N_{+-}N_{-+})^m (\rho'_0 Z_H)^{n'} \left(\frac{Z_G}{2}\right)^n} \leq \exp\left(\frac{\varepsilon}{2}\right) Z_{B,U}(\alpha, \alpha, \lambda'),$$

where we use the abbreviated expressions $Z_{B'} = Z_{B'}(\beta, \gamma, \lambda), Z_H = Z_H(\beta, \gamma, \lambda)$, and $Z_G = Z_G(\beta, \gamma, \lambda)$. The lemma follows by one oracle call for $Z_{B'}$ with accuracy $\frac{\varepsilon}{6}$, one oracle call for Z_G with accuracy $\frac{\varepsilon}{6n}$, and an approximation of other terms in the denominator with accuracy $\frac{\varepsilon}{6}$ using FPRASes for q^-, q^+, β, γ and λ .

The graph B' is constructed as follows. For each vertex $v \in V$ we introduce a copy of G , denoted by G_v with vertex set $V(G_v)$. Moreover, for each vertex $u \in U$ we introduce a copy of H , denoted by H_u . Whenever a terminal vertex is used in the construction once, we say it is occupied. For each $(u, v) \in E$, we connect one currently unoccupied positive (and respectively negative) terminal of G_u to one currently unoccupied positive (and respectively negative) terminal of G_v . Denote by E' all these edges between terminals. For each $u \in U$, we identify an unoccupied positive terminal of G_u with the distinguished vertex v_{H_u} of H_u . We denote this terminal by t_u . The resulting graph is B' . It is clear that B' is bipartite and has bounded degree Δ .

Let $\tilde{\sigma}: V \rightarrow \{-, +\}$ be a configuration of the phases of the G_v 's. Let $Z_{B'}(\tilde{\sigma})$ be the contribution to $Z_{B'}$ from the configurations σ that are consistent with $\tilde{\sigma}$ in the sense that, for each $v \in V, Y(\sigma_{V(G_v)}) = \tilde{\sigma}(v)$. Then $Z_{B'} = \sum_{\tilde{\sigma}} Z_{B'}(\tilde{\sigma})$. Let T be the set of all terminals $T = \cup_{v \in V} T(G_v)$ and $\tau: T \rightarrow \{0, 1\}$ be a configuration on T . Let $\tau_{T(G_v)}$ be the configuration τ restricted to $T(G_v)$. Recall that for $\pi \in \{-, +\}, Z_{G_v}^\pi(\tau_{T(G_v)})$ is the

contribution to Z_{G_v} from configurations that have phase π and are consistent with $\tau_{T(G_v)}$. Also, $\Pr_{G_v;\beta,\gamma,\lambda}(\tau_{T(G_v)} \mid Y(\sigma_{V(G_v)}) = \pi) = Z_{G_v}^\pi(\tau_{T(G_v)})/Z_{G_v}^\pi$. Moreover, for each $u \in U$ and each spin $i \in \{0,1\}$, let $Z_{H_u}(i)$ be the contribution to Z_{H_u} from configurations σ with $\sigma(t_u) = i$. Hence, $\rho_i = \Pr_{H_u;\beta,\gamma,\lambda}(\sigma(t_u) = i) = \frac{Z_{H_u}(i)}{Z_{H_u}}$. We express $Z_{B'}(\tilde{\sigma})$ as $Z_{B'}(\tilde{\sigma}) = \sum_{\tau: T \rightarrow \{0,1\}} w_{E'}(\tau) \prod_{v \in V} Z_{G_v}^{\tilde{\sigma}(v)}(\tau_{T(G_v)}) \prod_{u \in U} (Z_{H_u}(\tau(t_u))/\lambda^{\tau(t_u)})$, where $w_{E'}(\tau)$ is the contribution of edges in E' given configuration τ . Notice that we divide the last factor by λ when $\tau(t_u) = 1$ because we counted the vertex weight twice in that case. Define $\widetilde{Z}_{B'}(\tilde{\sigma})$ to be an approximation version of the partition function where on each $T(G_v)$ the spins are chosen exactly according to $Q^{\tilde{\sigma}(v)}$. That is,

$$\begin{aligned} \widetilde{Z}_{B'}(\tilde{\sigma}) &= \sum_{\tau: T \rightarrow \{0,1\}} w_{E'}(\tau) \prod_{v \in V} Z_{G_v}^{\tilde{\sigma}(v)} Q^{\tilde{\sigma}(v)}(\tau_{T(G_v)}) \prod_{u \in U} \frac{Z_{H_u}(\tau(t_u))}{\lambda^{\tau(t_u)}} \\ &= \left(\prod_{v \in V} Z_{G_v}^{\tilde{\sigma}(v)} \right) \cdot \left(\sum_{\tau: T \rightarrow \{0,1\}} w_{E'}(\tau) \prod_{v \in V} Q^{\tilde{\sigma}(v)}(\tau_{T(G_v)}) \prod_{u \in U} \frac{Z_{H_u}(\tau(t_u))}{\lambda^{\tau(t_u)}} \right). \end{aligned} \quad (8)$$

Let $\widetilde{Z}_{B'} = \sum_{\tilde{\sigma}} \widetilde{Z}_{B'}(\tilde{\sigma})$. Eq. (4) implies that $Z_{B'}(\tilde{\sigma})$ and $\widetilde{Z}_{B'}(\tilde{\sigma})$ are close, that is,

$$(1 - \varepsilon')^n \leq \frac{Z_{B'}(\tilde{\sigma})}{\widetilde{Z}_{B'}(\tilde{\sigma})} \leq (1 + \varepsilon')^n. \quad (9)$$

Moreover Eq. (5) implies that

$$\left(\frac{1 - \varepsilon'}{2} \right)^n \leq \frac{\prod_{v \in V} Z_{G_v}^{\tilde{\sigma}(v)}}{(Z_G)^n} \leq \left(\frac{1 + \varepsilon'}{2} \right)^n. \quad (10)$$

Notice that here Z_{G_v} is the same for any $v \in V$ as the G_v 's are identical copies of G . Now, given $\tilde{\sigma}$, we calculate $\sum_{\tau: T \rightarrow \{0,1\}} w_{E'}(\tau) \prod_{v \in V} Q^{\tilde{\sigma}(v)}(\tau_{T(G_v)}) \prod_{u \in U} Z_{H_u}(\tau(t_u))/\lambda^{\tau(t_u)}$. As the measure $Q^{\tilde{\sigma}(v)}$ is i.i.d., we may count the weight of each edge in E' independently. Notice that $N_{\pi_1 \pi_2}$ is the edge contribution when one end point is chosen with probability q^{π_1} and the other q^{π_2} . For an edge $(u, v) \in V$, if u and v are assigned the same phase $+$, then an edge in E' connecting one $+$ terminal of G_u and one $+$ terminal of G_v gives a weight of N_{++} and an edge connecting two $-$ terminals gives N_{--} . The total weight is $\mu_1 = N_{++}N_{--}$. Similarly if u and v are assigned the same phase $-$, the total weight is μ_1 as well. On the other hand if u and v are assigned distinct phases $+$ and $-$, the total weight is $\mu_2 = N_{+-}N_{-+}$. Recall that $\alpha = \frac{\mu_1}{\mu_2}$. Moreover, for each $u \in U$, if $\tilde{\sigma}(u) = +$, then the contribution of H_u is $\rho'_1 Z_{H_u}$ and otherwise $\rho'_0 Z_{H_u}$. Notice that here Z_{H_u} is the same for any $u \in U$ as the H_u 's are identical copies of H . Recall that $\lambda' = \frac{\rho'_1}{\rho'_0}$. Plugging these calculations into Eq. (8), we have

$$\begin{aligned} \widetilde{Z}_{B'}(\tilde{\sigma}) &= \left(\prod_{v \in V} Z_{G_v}^{\tilde{\sigma}(v)} \right) \cdot \left(\mu_1^{m_+(\tilde{\sigma})} \mu_2^{m-m_+(\tilde{\sigma})} (\rho'_1 Z_H)^{n_+(\tilde{\sigma})} (\rho'_0 Z_H)^{n'-n_+(\tilde{\sigma})} \right) \\ &= \mu_2^m (\rho'_0 Z_H)^{n'} \left(\prod_{v \in V} Z_{G_v}^{\tilde{\sigma}(v)} \right) \cdot \left(\alpha^{m_+(\tilde{\sigma})} (\lambda')^{n_+(\tilde{\sigma})} \right), \end{aligned} \quad (11)$$

where $m_+(\tilde{\sigma})$ denotes the number of edges of which the two endpoints are of the same phase given $\tilde{\sigma}$, and $n_+(\tilde{\sigma})$ denotes the number of vertices in U that are assigned $+$ given $\tilde{\sigma}$. Apply Eq.(10) to Eq.(11),

$$(1 - \varepsilon')^n \left(\alpha^{m_+(\tilde{\sigma})} (\lambda')^{n_+(\tilde{\sigma})} \right) \leq \frac{\widetilde{Z}_{B'}(\tilde{\sigma})}{\mu_2^m (\rho'_0 Z_H)^{n'} \left(\frac{Z_G}{2} \right)^n} \leq (1 + \varepsilon')^n \left(\alpha^{m_+(\tilde{\sigma})} (\lambda')^{n_+(\tilde{\sigma})} \right). \quad (12)$$

Then we sum over $\tilde{\sigma}$ in Eq. (12),

$$(1 - \varepsilon')^n \left(\sum_{\tilde{\sigma}} \alpha^{m_+(\tilde{\sigma})} (\lambda')^{n_+(\tilde{\sigma})} \right) \leq \frac{\widetilde{Z_{B'}}}{\mu_2^m (\rho'_0 Z_H)^{n'} \left(\frac{Z_G}{2}\right)^n} \leq (1 + \varepsilon')^n \left(\sum_{\tilde{\sigma}} \alpha^{m_+(\tilde{\sigma})} (\lambda')^{n_+(\tilde{\sigma})} \right). \tag{13}$$

However notice that $Z_{B,U}(\alpha, \alpha, \lambda') = \sum_{\tilde{\sigma}} \alpha^{m_+(\tilde{\sigma})} (\lambda')^{n_+(\tilde{\sigma})}$ by just mapping $-$ to 0 and $+$ to 1 in each configuration $\tilde{\sigma}$. Combine Eq.(9), and Eq.(13),

$$(1 - \varepsilon')^{2n} Z_{B,U}(\alpha, \alpha, \lambda') \leq \frac{Z_{B'}}{\mu_2^m (\rho'_0 Z_H)^{n'} \left(\frac{Z_G}{2}\right)^n} \leq (1 + \varepsilon')^{2n} Z_{B,U}(\alpha, \alpha, \lambda').$$

Recall that $\varepsilon' = \frac{\varepsilon}{8n}$ and we get the desired bounds.

The other case is ferromagnetic, that is, $\beta\gamma > 1$. Notice that in this case $\det(N) = (\beta\gamma - 1)(q^+ - q^-)^2 > 0$, So we choose $\alpha = \frac{N_{+-}N_{-+}}{N_{++}N_{--}} < 1$ and λ' to be the same as the antiferromagnetic case. The construction of B' is similar to the previous case, with the following change. For each $(u, v) \in E$, we connect one unoccupied positive terminal of G_u to one unoccupied negative terminal of G_v , and vice versa. The rest of the construction is the same. With this change, given a configuration $\tilde{\sigma}: V \rightarrow \{-, +\}$, if two endpoints are assigned the same spin, the contribution is $N_{+-}N_{-+}$ and otherwise $N_{++}N_{--}$. Therefore the effective edge weight is $\alpha < 1$ when the spins are the same, after normalizing the weight to 1 when the spins are distinct. The rest of the proof is the same. ◀

4.3 Completing the Proof of Theorem 1

Proof of Theorem 1. #BIS-hardness in Theorem 1 follows directly from Lemma 12 and Lemma 13. The other direction, #BIS-easiness, follows fairly directly from Theorem 47 of [6] (the full version of [7]). An edge in the instance graph can be viewed as a constraint of arity 2. If $\beta\gamma > 1$, then the constraint on the edge is “weakly log-supermodular” and the vertex weight can be viewed as a unary constraint, which is taken as given in a “conservative” CSP. If $\beta\gamma \leq 1$, then reverse the interpretation of 0 and 1 on one side of the bipartition of the instance graph, so that the effective interaction along an edge is given by the matrix $\begin{pmatrix} 1 & \beta \\ \gamma & 1 \end{pmatrix}$. This constraint is also “weakly log-supermodular” since $1 \cdot 1 \geq \beta\gamma$. After the reversing there are two vertex weights λ and λ^{-1} , which are also allowed for “conservative” CSP instances. ◀

Acknowledgement. Jin-Yi Cai and Heng Guo are supported by NSF grant CCF-1217549. Heng Guo is also supported by a 2013 Simons award for graduate students in theoretical computer science. Andreas Galanis and Eric Vigoda are supported by NSF grant CCF-1217458. Daniel Štefankovič is supported by NSF grant CCF-1318374.

The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) ERC grant agreement no. 334828. The paper reflects only the authors’ views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.

References

- 1 Andrei A. Bulatov. The complexity of the counting constraint satisfaction problem. In *ICALP*, pages 646–661. Springer-Verlag, 2008.
- 2 Andrei A. Bulatov, Martin Dyer, Leslie Ann Goldberg, Mark Jerrum, and Colin McQuillan. The expressibility of functions on the Boolean domain, with applications to counting CSPs. *J. ACM*, 60(5):32:1–32:36, October 2013.
- 3 Jin-Yi Cai and Xi Chen. Complexity of counting CSP with complex weights. In *STOC*, pages 909–920. ACM, 2012.
- 4 Jin-Yi Cai, Xi Chen, Heng Guo, and Pinyan Lu. Inapproximability after uniqueness phase transition in two-spin systems. In *COCOA*, pages 336–347, 2012.
- 5 Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holant problems and counting CSP. In *STOC*, pages 715–724. ACM, 2009.
- 6 Xi Chen, Martin E. Dyer, Leslie Ann Goldberg, Mark Jerrum, Pinyan Lu, Colin McQuillan, and David Richerby. The complexity of approximating conservative counting CSPs. Preprint, 2012. Available from the arXiv at: <http://arxiv.org/abs/1208.1783>
- 7 Xi Chen, Martin E. Dyer, Leslie Ann Goldberg, Mark Jerrum, Pinyan Lu, Colin McQuillan, and David Richerby. The complexity of approximating conservative counting CSPs. In *STACS*, pages 148–159, 2013.
- 8 Jin-Yi Cai, Andreas Galanis, Leslie Ann Goldberg, Heng Guo, Mark Jerrum, Daniel Štefankovič, and Eric Vigoda. #BIS-Hardness for 2-Spin Systems on Bipartite Bounded Degree Graphs in the Tree Non-uniqueness Region. Preprint, 2014. Available from the arXiv at: <http://arxiv.org/abs/1311.4451>
- 9 Martin E. Dyer, Leslie Ann Goldberg, Catherine S. Greenhill, and Mark Jerrum. The relative complexity of approximate counting problems. *Algorithmica*, 38(3):471–500, 2003.
- 10 Martin E. Dyer, Leslie Ann Goldberg, and Mark Jerrum. An approximation trichotomy for Boolean #CSP. *J. Comput. Syst. Sci.*, 76(3-4):267–277, 2010.
- 11 Martin E. Dyer and David Richerby. An effective dichotomy for the counting constraint satisfaction problem. *SIAM J. Comput.*, 42(3):1245–1274, 2013.
- 12 Andreas Galanis, Qi Ge, Daniel Štefankovič, Eric Vigoda, and Linji Yang. Improved inapproximability results for counting independent sets in the hard-core model. In *APPROX-RANDOM*, pages 567–578, 2011.
- 13 Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. Preprint, 2012. Available from the arXiv at: <http://arxiv.org/abs/1203.2226v3>
- 14 Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability for antiferromagnetic spin systems in the tree non-uniqueness region. Preprint, 2013. To appear in *STOC*, 2014. Available from the arXiv at: <http://arxiv.org/abs/1305.2902v2>
- 15 Qi Ge and Daniel Štefankovič. A graph polynomial for independent sets of bipartite graphs. In *FSTTCS*, pages 240–250, 2010.
- 16 Hans-Otto Georgii. *Gibbs Measures and Phase Transitions, 2nd edition*. Walter de Gruyter, 2011.
- 17 Leslie Ann Goldberg and Mark Jerrum. The complexity of ferromagnetic Ising with local fields. *Combinatorics, Probability & Computing*, 16(1):43–61, 2007.
- 18 Leslie Goldberg and Mark Jerrum. A counterexample to rapid mixing of the Ge-Stefankovic process. *Electron. Commun. Probab.*, 17:no. 5, 1–6, 2012.
- 19 Leslie Ann Goldberg and Mark Jerrum. Approximating the partition function of the ferromagnetic Potts model. *J. ACM*, 59(5):25, 2012.
- 20 Leslie Ann Goldberg, Mark Jerrum, and Colin McQuillan. Approximating the partition function of planar two-state spin systems. Preprint, 2012. Available from the arXiv at: <http://arxiv.org/abs/1208.4987>

- 21 Leslie Ann Goldberg, Mark Jerrum, and Mike Paterson. The computational complexity of two-state spin systems. *Random Struct. Algorithms*, 23(2):133–154, 2003.
- 22 Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM J. Comput.*, 22(5):1087–1116, 1993.
- 23 Frank Kelly. Stochastic models of computer communication systems. *Journal of the Royal Statistical Society. Series B (Methodological)*, 47(3):pp. 379–395, 1985.
- 24 Liang Li, Pinyan Lu, and Yitong Yin. Approximate counting via correlation decay in spin systems. In *SODA*, pages 922–940, 2012.
- 25 Liang Li, Pinyan Lu, and Yitong Yin. Correlation decay up to uniqueness in spin systems. In *SODA*, pages 67–84, 2013. Available from the arXiv at: <http://arxiv.org/pdf/1111.7064v2.pdf>
- 26 Jingcheng Liu, Pinyan Lu, and Chihao Zhang. The complexity of ferromagnetic two-spin systems with external fields. Preprint, 2014. Available from the arXiv at: <http://arxiv.org/abs/1402.4346>
- 27 Elchanan Mossel, Dror Weitz, and Nicholas Wormald. On the hardness of sampling independent sets beyond the tree threshold. *Prob. Theory Related. Fields*, 143:401–439, 2009.
- 28 Alistair Sinclair, Piyush Srivastava, and Marc Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. In *SODA*, pages 941–953, 2012.
- 29 Allan Sly. Computational transition at the uniqueness threshold. In *FOCS*, pages 287–296, 2010.
- 30 Allan Sly and Nike Sun. The computational hardness of counting in two-spin models on d -regular graphs. In *FOCS*, pages 361–369, 2012.
- 31 Dror Weitz. Counting independent sets up to the tree threshold. In *STOC*, pages 140–149, 2006.
- 32 David Zuckerman. On unapproximable versions of NP-complete problems. *SIAM J. Comput.*, 25(6):1293–1304, 1996.