

The complexity of approximating averages on bounded-degree graphs

Andreas Galanis <i>Department of Computer Science</i> <i>University of Oxford</i> <i>Oxford, UK</i> <i>Email: andreas.galanis@cs.ox.ac.uk</i>	Daniel Štefankovič <i>Department of Computer Science</i> <i>University of Rochester</i> <i>Rochester, USA</i> <i>Email: stefanko@cs.rochester.edu</i>	Eric Vigoda <i>School of Computer Science</i> <i>Georgia Institute of Technology</i> <i>Atlanta, USA</i> <i>Email: ericvigoda@gmail.com</i>
---	---	---

Abstract—We prove that, unless $P=NP$, there is no polynomial-time algorithm to approximate within some multiplicative constant the average size of an independent set in graphs of maximum degree 6. This is a special case of a more general result for the hard-core model defined on independent sets weighted by a parameter $\lambda > 0$. In the general setting, we prove that, unless $P=NP$, for all $\Delta \geq 3$, all $\lambda > \lambda_c(\Delta)$, there is no FPTAS which applies to all graphs of maximum degree Δ for computing the average size of the independent set in the Gibbs distribution, where $\lambda_c(\Delta)$ is the critical point for the uniqueness/non-uniqueness phase transition on the Δ -regular tree. Moreover, we prove that for λ in a dense set of this non-uniqueness region the problem is NP-hard to approximate within some constant factor. Our work extends to the antiferromagnetic Ising model and generalizes to all 2-spin antiferromagnetic models, establishing hardness of computing the average magnetization in the tree non-uniqueness region.

Previously, Schulman, Sinclair and Srivastava (2015) showed that it is #P-hard to compute the average magnetization exactly, but no hardness of approximation results were known. Hardness results of Sly (2010) and Sly and Sun (2014) for approximating the partition function do not imply hardness of computing averages. The new ingredient in our reduction is an intricate construction of pairs of rooted trees whose marginal distributions at the root agree but their derivatives disagree. The main technical contribution is controlling what marginal distributions and derivatives are achievable and using Cauchy’s functional equation to argue existence of the gadgets.

The full version of this paper with detailed proofs to all lemmas and theorems can be found out at <https://arxiv.org/abs/2004.09238>.

Keywords—hardness of approximation, averages, independent sets, magnetization, Ising model

I. INTRODUCTION

This paper addresses the computational problem of computing averages of simple functions over combinatorial structures of a graph. Can we estimate elementary statistics of combinatorial structures in polynomial-time? This genre of problems is nicely illustrated for the example of independent sets.

Given a graph $G = (V, E)$ can we efficiently estimate the average size of an independent set in G ? For graph $G = (V, E)$, let \mathcal{I}_G denote the set of independent sets of

G , and let μ_G denote the uniform distribution over \mathcal{I}_G . Denote the average independent set size by $\mathcal{M}(G) = \mathbf{E}_{\sigma \sim \mu} [|\sigma|]$.

Schulman *et al.* [12] (see also [13]) proved that exactly computing the average independent set size is #P-hard for bounded-degree graphs. We investigate approximation algorithms for the problem. For a constant $C > 1$ we say there is a C -approximation algorithm for the average independent set size if the algorithm outputs an estimate EST for which $\frac{1}{C} \times \mathcal{M}(G) \leq \text{EST} \leq C \times \mathcal{M}(G)$. An FPTAS is an algorithm which, for any input $G = (V, E)$ and $\epsilon > 0$, achieves a $(1 + \epsilon)$ -approximation factor in time $\text{poly}(|V|, 1/\epsilon)$.

Weitz [17] presented an FPTAS for estimating $|\mathcal{I}_G|$, the number of independent sets, in graphs of maximum degree ≤ 5 , and this immediately yields an FPTAS for the average independent set size in graphs of maximum degree ≤ 5 , see also [2], [9], [10] for new algorithmic approaches. We prove that this result is optimal. In fact, we prove that approximating the average independent set size in graphs of maximum degree 6 is hard within a constant factor.

Theorem 1. *There is a constant $C > 1$ such that, for all integers $\Delta \geq 6$, for graphs G of maximum degree Δ there is no polynomial-time C -approximation algorithm for computing the average independent-set size of G , unless $P=NP$.*

This theorem is a special case of a more general result for the hard-core model, which is a statistical physics model of particular combinatorial interest. The hard-core model is defined on independent sets weighted by a parameter $\lambda > 0$, known as the fugacity. An independent set $\sigma \in \mathcal{I}_G$ has weight $w(\sigma) = \lambda^{|\sigma|}$. For a graph G and fugacity $\lambda > 0$, the Gibbs distribution is defined as $\mu_{G;\lambda}(\sigma) = w(\sigma)/Z_{G;\lambda}$ where the partition function $Z_{G;\lambda} = \sum_{\tau \in \mathcal{I}_G} w(\tau)$.

The earlier case of unweighted independent sets corresponds to the hard-core model with $\lambda = 1$. Hence we use the same notation $\mathcal{M}(G) = \mathbf{E}_{\sigma \sim \mu} [|\sigma|]$ to denote the average independent set size in the Gibbs distribution.

On the Δ -regular tree, the hard-core model undergoes a phase transition at the critical point $\lambda_c(\Delta) = \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^\Delta}$. When $\lambda \leq \lambda_c(\Delta)$ there is a unique infinite-volume Gibbs measure on the Δ -regular tree (roughly, this corresponds to the decay of the “influence” of the leaves on the root), whereas when $\lambda > \lambda_c(\Delta)$ there is non-uniqueness, i.e., there are multiple infinite-volume Gibbs measures.

There is an interesting computational phase transition for graphs of maximum degree Δ that occurs at this same tree threshold. For all constant Δ , all $\lambda < \lambda_c(\Delta)$, all graphs of maximum degree Δ , there is an FPTAS for the partition function [17]. On the other side, for all $\Delta \geq 3$, all $\lambda > \lambda_c(\Delta)$, there is no FPRAS for approximating the partition function on graphs of maximum degree Δ , unless NP=RP [15], [16], [5]. However, hardness of computing partition functions does not imply hardness of computing averages in the Gibbs distribution; see the case of the antiferromagnetic Ising model with no external field discussed in Section I-A.

We prove that computing the average independent set size also undergoes a computational phase transition at the tree critical point $\lambda_c(\Delta)$. As before, Weitz’s algorithmic result [17] yields, for all constant Δ , all $\lambda < \lambda_c(\Delta)$, all graphs G of maximum degree Δ , an FPTAS for the average independent set size $\mathcal{M}(G)$ (more generally, an approximate sampling algorithm implies an algorithm for approximating averages). We prove that this result is optimal: when $\lambda > \lambda_c(\Delta)$ there is no FPTAS for the average independent set size.

Theorem 2. *Let $\Delta \geq 3$ be an integer and $\lambda > \lambda_c(\Delta)$. Then, for graphs G of maximum degree Δ , there is no FPTAS for computing the average independent-set size in the hard-core distribution $\mu_{G;\lambda}$, unless P=NP.*

In fact, our inapproximability result for general $\lambda > \lambda_c(\Delta)$ is stronger, it actually precludes approximation algorithms with factors of $1 \pm \frac{\delta}{\log n}$, where n is the number of the vertices of the input graph and δ is an appropriate constant, see Theorem 7 below for the precise statement. For a dense set of λ , we actually obtain constant-factor inapproximability (analogously to Theorem 1).

Theorem 3. *Let $\Delta \geq 3$ be an integer. Then, for every real $\lambda > \lambda_c(\Delta)$ and $\epsilon > 0$, there is an algebraic number $\hat{\lambda}$ with $|\hat{\lambda} - \lambda| \leq \epsilon$ and a constant $C = C(\hat{\lambda}, \Delta) > 1$ such that, for graphs G of maximum degree Δ , there is no poly-time C -approximation algorithm for computing the average independent-set size of G in the hard-core distribution $\mu_{G;\hat{\lambda}}$, unless P=NP.*

A. Results for the antiferromagnetic Ising model

Our results extend to the antiferromagnetic Ising model. Let $G = (V, E)$ be a graph. For $\beta, \lambda > 0$, let

$\mu_{G;\beta,\lambda}$ denote the Ising distribution on G with edge activity β and external field λ , i.e., for $\sigma : V \rightarrow \{0, 1\}$ we have

$$\mu_{G;\beta,\lambda}(\sigma) = \frac{\lambda^{|\sigma|} \beta^{m(\sigma)}}{Z_{G;\beta,\lambda}},$$

where $m(\sigma)$ denotes the number of monochromatic edges in G under σ , i.e., edges whose endpoints have the same spin. The model is called antiferromagnetic if $\beta \in (0, 1)$ and ferromagnetic, otherwise. We define the average magnetization of G to be the average number of vertices with spin 1, i.e.,

$$\begin{aligned} \mathcal{M}_{\beta,\lambda}(G) &= \frac{1}{Z_{G;\beta,\lambda}} \sum_{\sigma:V \rightarrow \{0,1\}} |\sigma| \lambda^{|\sigma|} \beta^{m(\sigma)} \\ &= \mathbf{E}_{\sigma \sim \mu} [|\sigma|], \end{aligned}$$

where for a configuration $\sigma : V \rightarrow \{0, 1\}$, we use $|\sigma|$ to denote $\sum_{v \in V(G)} \sigma(v)$, i.e., the total number of vertices with spin 1.

In the ferromagnetic case, there is an FPRAS for approximating the magnetization for all $\beta > 1$ and $\lambda > 0$, due to the algorithm of Jerrum and Sinclair [6]. For $\Delta \geq 3$, let $\beta_c(\Delta) = \frac{\Delta-2}{\Delta}$. It is known that the antiferromagnetic Ising model with edge activity β and external field λ has non-uniqueness on the infinite Δ -regular tree iff $\beta \in (0, \beta_c)$ and $\lambda \in (1/\lambda_c^{\text{Ising}}, \lambda_c^{\text{Ising}})$ for some explicit $\lambda_c^{\text{Ising}} = \lambda_c^{\text{Ising}}(\beta, \Delta) > 1$. For all constant Δ , in the tree uniqueness region there is an FPTAS for the partition function for graphs of maximum degree Δ [14], [7], and, once again, this implies an FPTAS for the magnetization. In the tree non-uniqueness region, for all $\Delta \geq 3$, unless NP=RP there is no FPRAS for the partition function for graphs of maximum degree Δ [16]. We prove that approximating the magnetization is also intractable in the tree non-uniqueness region, apart from the case $\lambda = 1$ (where the magnetization can be computed trivially for all graphs G since it equals $\frac{1}{2}|V(G)|$).

Theorem 4. *Let $\Delta \geq 3$ be an integer, $\beta \in (0, \beta_c(\Delta))$ and $\lambda \in (\frac{1}{\lambda_c}, \lambda_c)$ with $\lambda \neq 1$, where $\lambda_c = \lambda_c^{\text{Ising}}(\Delta, \beta)$. Then, for graphs G of maximum degree Δ , there is no FPTAS for computing the average magnetization in the Ising distribution $\mu_{G;\beta,\lambda}$, unless P=NP.*

As in Theorem 2, the inapproximability factor in Theorem 4 is in fact stronger, see Theorem 7 below for the precise statement. For a dense set of λ , we again obtain constant-factor inapproximability.

Theorem 5. *Let $\Delta \geq 3$ be an integer, rational $\beta \in (0, \beta_c(\Delta))$ and $\lambda_c = \lambda_c^{\text{Ising}}(\Delta, \beta)$. Then, for every $\lambda \in (\frac{1}{\lambda_c}, \lambda_c)$ and $\epsilon > 0$, there is an algebraic number $\hat{\lambda}$ with $|\hat{\lambda} - \lambda| \leq \epsilon$ and a constant $C = C(\beta, \hat{\lambda}, \Delta) > 1$ such that, for graphs G of maximum degree Δ , there is*

no poly-time C -approximation algorithm for computing the average magnetization $\mathcal{M}_{\beta, \hat{\lambda}}(G)$, unless $P=NP$.

B. Results for general antiferromagnetic 2-spin systems

While the hard-core model and the Ising model are the most canonical 2-spin models, the results of the previous two sections will be obtained as special cases of the following results for general antiferromagnetic 2-spin systems. This more general perspective will also allow us to give a unified proof of Theorems 2, 3, 4, 5.

Let $G = (V, E)$ be a graph. For $\beta, \gamma, \lambda > 0$, let $\mu_{G; \beta, \gamma, \lambda}$ denote the Gibbs distribution on G with edge activities β, γ and external field λ , i.e., for $\sigma : V \rightarrow \{0, 1\}$ we have

$$\mu_{G; \beta, \lambda}(\sigma) = \frac{\lambda^{|\sigma|} \beta^{m_0(\sigma)} \gamma^{m_1(\sigma)}}{Z_{G; \beta, \gamma, \lambda}},$$

where $m_0(\sigma), m_1(\sigma)$ denotes the number of edges in G whose endpoints are assigned under σ the pair of spins $(0, 0)$ and $(1, 1)$, respectively.

The parameter pair (β, γ) is called *antiferromagnetic* if $\beta\gamma \in [0, 1)$ and at least one of β, γ is non-zero, and it is called *ferromagnetic*, otherwise. Note that the hard-core model is the case $\beta = 1, \gamma = 0$ (under the convention that $0^0 \equiv 1$) whereas the antiferromagnetic Ising model is the case $0 < \beta = \gamma < 1$.

We next define the range of parameters (β, γ, λ) where our inapproximability results for the magnetization apply; these are precisely the parameters where the spin system exhibits non-uniqueness on the infinite Δ -regular tree, apart from the case of the antiferromagnetic Ising model with $\lambda = 1$, where as noted in Section I-A the magnetization can be computed trivially for all graphs G .

Definition 6. Let $\Delta \geq 3$ be an integer. We let \mathcal{U}_Δ be the set of (β, γ, λ) such that (β, γ) is antiferromagnetic, and the (unique) fixpoint $x^* > 0$ of the function $f(x) = \frac{1}{\lambda} \left(\frac{\beta x + 1}{x + \gamma} \right)^{\Delta-1}$ satisfies $|f'(x^*)| > 1$.

We let $\mathcal{U}_\Delta^* = \mathcal{U}_\Delta \setminus \bigcup_{\beta \in (0, 1)} \{(\beta, \beta, 1)\}$ be the set of parameters in \mathcal{U}_Δ other than those where computing the magnetization is trivial.

We note that Li, Lu, Yin [7] define a notion of “up-to- Δ ” uniqueness which requires (β, γ, λ) to be in uniqueness for every $d \leq \Delta$; they obtain an FPTAS for the partition function in that region. The complement of their region corresponds to non-uniqueness in the sense of Definition 6 for some $d \leq \Delta$. Our inapproximability results extend to this bigger region by applying our theorems for smaller values of Δ .

Our first inapproximability result for general antiferromagnetic 2-spin models, which is a generalization/strengthening of Theorems 2 and 4, is the following. The proof is given in Section 5.5 of the full version.

Theorem 7. Let $\Delta \geq 3$ be an integer and $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta^*$. Then, for graphs $G \in \mathcal{G}_\Delta$, there is no FPTAS for computing the average magnetization in the Gibbs distribution $\mu_{G; \beta, \gamma, \lambda}$, unless $P=NP$. In fact, there is a constant $\kappa = \kappa(\Delta, \beta, \gamma, \lambda) > 0$ such that there is no poly-time $(1 + \frac{\kappa}{\log n})$ -approximation algorithm for computing the average magnetization $\mathcal{M}_{\beta, \gamma, \lambda}(G)$, where $n = |V(G)|$.

We remark that a “no-FPTAS” result can be strengthened to an inapproximability factor of $(1 \pm \frac{1}{n^\epsilon})$ for any constant $\epsilon > 0$ via standard powering techniques; the tighter hardness factor of $(1 \pm \frac{\kappa}{\log n})$ in Theorem 7 requires a substantially more delicate argument.

Theorem 8. Let $\Delta \geq 3$ be an integer and $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta$ with β, γ rational numbers. Then, for every $\epsilon > 0$, there is an algebraic number $\hat{\lambda}$ with $|\hat{\lambda} - \lambda| \leq \epsilon$ and a constant $C = C(\beta, \gamma, \hat{\lambda}, \Delta) > 1$ such that, for graphs G of maximum degree Δ , there is no poly-time C -approximation algorithm for computing the average magnetization $\mathcal{M}_{\beta, \gamma, \hat{\lambda}}(G)$, unless $P=NP$.

II. PROOF OUTLINE

In this section, we give some of the key elements of the techniques needed to prove our inapproximability results and conclude with the proof of Theorem 1. We start by describing “field gadgets” and state the main lemmas that we will use in our reduction.

A. Gadgets with approximately equal effective fields and different averages

For a graph G and $\sigma : V \rightarrow \{0, 1\}$, we say that a vertex $v \in V$ is occupied if $\sigma(v) = 1$, and unoccupied otherwise.¹ For a subset $S \subseteq V$, we use σ_S to denote the configuration on S which is restriction of σ on S .

Definition 9. A field gadget is a rooted tree \mathcal{T} whose root ρ has degree one. For antiferromagnetic (β, γ) and $\lambda > 0$, let $\mu = \mu_{\mathcal{T}; \beta, \gamma, \lambda}$.

- 1) The effective field of the gadget, denoted by $R_{\mathcal{T}}(\lambda)$, is $1/\lambda$ times the ratio of the weight of configurations where the root is occupied to the weight of configurations where the root is unoccupied, i.e., $R_{\mathcal{T}}(\lambda) = \frac{1}{\lambda} \frac{\mu(\sigma_\rho=1)}{\mu(\sigma_\rho=0)}$. (The division by λ is to avoid double-counting the contribution of the root later on.)
- 2) The magnetization gap of the gadget, denoted by $M_{\mathcal{T}; \beta, \gamma, \lambda}$, is the expected number of occupied vertices conditioned on the root being occupied minus the expected number of occupied vertices

¹This terminology is standard for the hard-core model, but it will be convenient to use it for general 2-spin systems.

conditioned on the root being unoccupied (the root is included in the count), i.e.,

$$M_{\mathcal{T};\beta,\gamma,\lambda} = \mathbf{E}_{\sigma \sim \mu} [|\sigma| \mid \sigma(\rho) = 1] - \mathbf{E}_{\sigma \sim \mu} [|\sigma| \mid \sigma(\rho) = 0].$$

In the special case that $\lambda = \frac{1-\beta}{1-\gamma}$ with $\beta \neq \gamma$, a field gadget consists of a rooted graph obtained from a rooted tree where some of the leaves have been replaced by a distinct triangle (by identifying the leaf with a vertex of the triangle).


Note that, in the case that $\lambda = \frac{1-\beta}{1-\gamma}$, the effective field of any tree gadget can be shown to be equal to 1 and the magnetization gap to 0, so that is why we need to consider the “mildly non-tree” construction in Definition 9.

It will be useful to illustrate Definition 9 with a few examples.

Example 10. For example the rooted tree with one edge has effective field $R = \frac{1+\gamma\lambda}{\beta+\lambda}$. The degenerate example where the root has degree zero has effective field $R = 1$.

The following more interesting example in the case of independent sets will be used to prove Theorem 1; it gives a pair of field gadgets with the same effective field but different magnetization gaps.

Example 11. Consider the independent set model (with $\lambda = 1$). Consider the trees $\mathcal{T}_1, \mathcal{T}_2$ with roots ρ_1, ρ_2 below.



Then $R_{\mathcal{T}_1} = R_{\mathcal{T}_2}$ since $R_{\mathcal{T}_1} = \frac{2}{3}$ and $R_{\mathcal{T}_2} = \frac{24}{36}$, but $M_{\mathcal{T}_1} \neq M_{\mathcal{T}_2}$ since $M_{\mathcal{T}_1} = \frac{3}{2} - \frac{2}{3} = \frac{5}{6}$ and $M_{\mathcal{T}_2} = \frac{35}{12} - \frac{13}{6} = \frac{3}{4}$. These values can be either verified by enumeration and linearity of expectation, or else using the recursions of the upcoming Lemma 17.

We will be interested in finding pairs of field gadgets analogous to Example 11. Our interest in such pairs of field gadgets is justified by the following theorem.

Theorem 12. Let $\Delta \geq 3$ be an integer and $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta^*$. Suppose that there exists a pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$ with maximum degree Δ such that $R_{\mathcal{T}_1} = R_{\mathcal{T}_2}$ but $M_{\mathcal{T}_1} \neq M_{\mathcal{T}_2}$.

Then, there is constant $c = c(\Delta, \beta, \gamma, \lambda) > 1$ such that, for graphs G of maximum degree Δ , there is no poly-time c -approximation algorithm for computing the average magnetization $\mathcal{M}_{\beta,\gamma,\lambda}(G)$, unless $\text{P}=\text{NP}$.

When can we find pairs of trees as in Example 11 with the same effective fields but different magnetiza-

tion gaps? Even in the case of the hard-core model, a pair of trees either have the same effective field for only finitely many values of λ or else have the same magnetization gap for all λ . Thus we cannot hope for a universal pair of gadgets. Actually, the set of λ where the effective fields can be equal, over all pairs of trees, must be algebraic, hence measure zero.

Our main theorem for the construction of field gadgets is the following. The theorem roughly asserts that we can construct field gadgets with arbitrarily close effective fields but substantially different magnetization gaps. For a rational $r = p/q$ with integers p, q satisfying $\text{gcd}(p, q) = 1$, we let $\text{bits}(r)$ denote the total number of bits needed to represent p, q .

Theorem 13. Let (β, γ, λ) be antiferromagnetic with $(\beta, \gamma, \lambda) \neq (\beta, \beta, 1)$. There exist $\hat{R}, \hat{M}, \Xi > 0$ and an algorithm which on input a rational $r \in (0, 1/2)$ outputs in time $\text{poly}(\text{bits}(r))$ a pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$, each of maximum degree 3 and size $O(|\log r|)$, such that

$$|R_{\mathcal{T}_1} - \hat{R}|, |R_{\mathcal{T}_2} - \hat{R}| \leq r, \text{ but } |M_{\mathcal{T}_1} - M_{\mathcal{T}_2}| > \hat{M}.$$

Also, the magnetization gaps $M_{\mathcal{T}_1}, M_{\mathcal{T}_2}$ are bounded in absolute value by the constant Ξ .

The proof of Theorem 13 is given in Section III. In fact, we can bootstrap Theorem 13 to obtain pairs of trees with the same effective fields but different magnetization gaps for a dense set of (algebraic numbers) λ , and this gives a constant factor inapproximability result for the magnetization by applying Theorem 12.

Theorem 14. Let (β, γ) be antiferromagnetic. There exists a set S of algebraic numbers λ , dense in the interval $(0, \infty)$, such that for each $\lambda \in S$ the following holds. There is a pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$, each of maximum degree 3, such that $R_{\mathcal{T}_1} = R_{\mathcal{T}_2}$ but $M_1 \neq M_2$.

In the following, we sketch the proof of Theorem 12 and in Section III we will give a detailed outline of the proof of Theorems 13 and 14.

B. The reduction: using field gadgets to obtain inapproximability

Fix $\Delta \geq 3$ and $\beta, \gamma, \lambda \in \mathcal{U}_\Delta$. Our results will be based on showing that approximating MAX-CUT on 3-regular graphs H reduces to approximating the magnetization will be via a reduction from MAX-CUT. The reduction uses two types of gadgets.

The first type of gadget is a bipartite graph G which is an almost Δ -regular graph with n vertices on each side and ℓ “ports” on each side of degree $\Delta - 1$, which will be used to connect distinct copies of gadgets. These gadgets were used in previous inapproximability results for the partition function and analysing them was one of the key difficulties in those results. The main idea

in these results is to replace each vertex of H by a distinct copy of G and make appropriate connections between ports to encode the edges of H ; then due to the antiferromagnetic interaction the resulting graph settles in configurations which correspond to Max-Cut configurations of H (the Max-Cut assignment can roughly be read off by looking at the “phases” of the gadgets G , i.e., which side has the most occupied vertices, see Section II-C for details).

However, for our inapproximability results we would need to analyze the magnetization of such gadgets by taking into account the effect of conditioning the spins of the ports; this task seems even more challenging than in the previous settings since it seems to require very refined estimates.

The field gadgets, which is the second type of gadgets, give a new reduction technique to bypass the need to perform this delicate, and likely difficult, task. In particular, by using a pair of field gadgets with the same effective fields and different magnetization gaps we will be able to observe the value of $\text{MAX-CUT}(H)$ by taking the difference of the magnetizations when we append our field gadgets appropriately; crucially, the fact that the effective fields are the same implies that the underlying distribution does not change but that the difference of the magnetization gap of the two gadgets manifests itself in the magnetization.

C. The bipartite “phase-gadget” of Sly and Sun

Fix an integer $\Delta \geq 3$. Let \mathcal{G}_n^ℓ be the set of $(2n)$ -vertex bipartite graphs whose two sides are labelled with $+$, $-$ and are obtained from a Δ -regular $(2n)$ -vertex bipartite graph by deleting a matching of size ℓ . For a graph $G \in \mathcal{G}_n^\ell$, we denote its bipartition by (U^+, U^-) where U^+, U^- are vertex sets with $|U^+| = |U^-| = n$, and we denote by W^+, W^- the sets of vertices with degree $\Delta - 1$ on each side of the bipartition, i.e., the vertices incident to the edges of the matching, so that $|W^+| = |W^-| = \ell$. We will refer to set $W = W^+ \cup W^-$ as the *ports* of G .

For $\sigma : U^+ \cup U^- \rightarrow \{0, 1\}$, we define the *phase* $\mathcal{Y}(\sigma)$ of the configuration σ as the side which has the most occupied vertices under σ , i.e.,

$$\mathcal{Y}(\sigma) = \begin{cases} + & \text{if } |\sigma_{U^+}| \geq |\sigma_{U^-}|, \\ - & \text{otherwise.} \end{cases}$$

Sly and Sun [15] (see also [4]) establish that, whenever $\beta, \gamma, \lambda \in \mathcal{U}_\Delta$, for arbitrary $\epsilon > 0$ and integer $\ell \geq 1$, there exist n and $G \in \mathcal{G}_n^\ell$ such that, in the Gibbs distribution of the gadget graph G , each phase appears with probability close to $1/2 \pm \epsilon$, and that the spins of the ports of G , conditional on the phase, are roughly independent, with occupation probabilities that are asymmetric between the two sides.

The detailed properties of the gadget can be found in Lemma 33 of Section 5.1 of the full version, though the exact details will not be important at this stage.

D. The reduction

Let H be a 3-regular instance of the MAX-CUT problem. Let k be an arbitrary positive integer, and let n and $G \in \mathcal{G}_n^{3k}$ be a bipartite gadget satisfying the properties of the previous subsection (cf. Lemma 33 of the full version). Let also \mathcal{T} be a field gadget with effective field R and magnetization gap M .

Let \widehat{H}_G^k be the graph with m disconnected components obtained by replacing each vertex of H with a distinct copy of G ; for $v \in V(H)$ we denote by G_v the copy of G in \widehat{H}_G^k corresponding to the vertex v , and by U_v^\pm, W_v^\pm the vertex sets and ports of G_v in each side of the bipartition. Finally, we let $U_v = U_v^+ \cup U_v^-$ and $W_v = W_v^+ \cup W_v^-$.

Let $H_{G, \mathcal{T}}^k$ be the graph obtained from \widehat{H}_G^k as follows. For each edge $e = (u, v) \in E(H)$, pick k distinct ports from W_u^\pm, W_v^\pm and connect them using a path with three edges, for a total of $2k$ edge-disjoint paths between the gadgets G_u and G_v : k paths between $+/+$ sides and k paths between $-/-$ sides. Then, for each path P append two distinct copies of the field gadget \mathcal{T} by identifying the roots of the copies of \mathcal{T} , say ρ_1, ρ_2 , with the internal vertices of P , say t_1, t_2 .

For a configuration $\sigma : V(H_{G, \mathcal{T}}^k) \rightarrow \{0, 1\}$, we let $\widehat{\mathcal{Y}}(\sigma) : V(H) \rightarrow \{+, -\}$ be the phases over the gadgets G_v with $v \in V(H)$, i.e., $\widehat{\mathcal{Y}}(\sigma) = \{\mathcal{Y}(\sigma_{U_v})\}_{v \in V(H)}$. We define $\widehat{\mathcal{Y}}(\cdot)$ similarly for configurations on \widehat{H}_G^k, H_G^k . For a phase assignment $Y : V(H) \rightarrow \{+, -\}$, we let $\text{Cut}_H(Y) = \{(u, v) \in E(H) \mid Y_u \neq Y_v\}$ be the set of edges that are cut in H by viewing the phase assignment Y as a bipartition of $V(H)$ in the natural way. For $\mu = \mu_{H_{G, \mathcal{T}}^k; \beta, \gamma, \lambda}$, let

$$\text{AVG-CUT}_\mu(H) := \sum_{Y: V(H) \rightarrow \{+, -\}} \mu(\widehat{\mathcal{Y}}(\sigma) = Y) |\text{Cut}_H(Y)|$$

be the size of the average cut in H when phase assignments are weighted by μ .

The following lemma associates the magnetization on the graph $H_{G, \mathcal{T}}^k$ with the quantity $\text{AVG-CUT}_\mu(H)$ and, in turn, with $\text{MAX-CUT}(H)$. The main idea is that the paths of length 3 between the ports cause antiferromagnetic interaction and the spin system on $H_{G, \mathcal{T}}^k$ prefers phase configurations Y that have large $|\text{Cut}_H(Y)|$ value. By choosing k large enough we can further ensure that $\text{AVG-CUT}_\mu(H)$ arbitrarily close to $\text{MAX-CUT}(H)$.

Lemma 15. *Let $\Delta \geq 3$ and $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta$. There are rational functions $A(R), B(R), C(R)$ defined for*

all $R \geq 0$, satisfying

- (i) $A(0) = 1$, (ii) $A(R) > 1$ for $R > 0$, and
- (iii) $B(R) - C(R) = (\log A(R))'$ for $R > 0$

so that the following holds for any field gadget \mathcal{T} with effective field $R > 0$ and magnetization gap M . Let A, B, C be the values of the functions at λR and define $A' := \mathbf{E}_{\tau \sim \mu_{\mathcal{T}; \beta, \gamma, \lambda}}[|\tau| \mid \tau(\rho) = 0]$, where ρ is the root of \mathcal{T} .

Let $\epsilon \in (0, 1/10)$, $k \geq 10/\log A$ be an integer, and n and $G \in \mathcal{G}_n^{3k}$ satisfy Lemma 33 (of the full version) with $\ell = 3k$ and the given ϵ . Then, for every 3-regular graph H with sufficiently large $|V(H)|$, for $\mu = \mu_{H_G^k; \beta, \gamma, \lambda}$, we have that the average magnetization of the graph $H_{G, \mathcal{T}}^k$ satisfies

$$\mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}}^k) = 4kA'|E(H)| + \mathbf{E}_{\sigma \sim \mu}[\left| \sigma_{V(\hat{H}_G^k)} \right|] + (1 \pm 8\epsilon)kMQ$$

with $Q := (B - C)\text{AVG-CUT}_\mu(H) + C|E(H)|$, while $\text{AVG-CUT}_\mu(H)$ satisfies

$$1/K \leq \frac{\text{AVG-CUT}_\mu(H)}{\text{MAX-CUT}(H)} \leq 1 \text{ for } K := 1 + \frac{6}{k \log A}.$$

We give the proof of Lemma 15 in Section 5.3 of the full version.

Note that the expression for the magnetization of $H_{G, \mathcal{T}}^k$ accounts for the contributions of the phase gadgets G only implicitly. This is by design; when we append our pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$ which have the same effective fields, vertices in $V(\hat{H}_G^k)$ will have the same marginal distribution in both $\mu_{H_{G, \mathcal{T}_1}^k}$ and $\mu_{H_{G, \mathcal{T}_2}^k}$; on the other hand the magnetization gaps of $\mathcal{T}_1, \mathcal{T}_2$ are different. Therefore, when we take the difference of the magnetizations in H_{G, \mathcal{T}_1}^k and H_{G, \mathcal{T}_2}^k , the contributions of the phase gadgets G cancel and, in the left-over quantity, the major contribution comes from the value of $\text{AVG-CUT}_\mu(H)$.

We are almost ready to prove Theorem 12. Recall, the assumption therein is that we have a pair of field gadgets with the same effective fields and different magnetization gaps. We will need to bootstrap this slightly to obtain additional pairs of field gadgets, as stated in the following lemma.

Lemma 16. *Let $\Delta \geq 3$ be an integer and (β, γ, λ) be antiferromagnetic with $(\beta, \gamma, \lambda) \neq (\beta, \beta, 1)$. Suppose that there exists a pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$ with maximum degree Δ such that $R_{\mathcal{T}_1} = R_{\mathcal{T}_2}$ but $M_{\mathcal{T}_1} \neq M_{\mathcal{T}_2}$. Then, we can construct for $j = 0, 1, 2, \dots$ an infinite sequence of pairs of field gadgets $\mathcal{T}_{1,j}, \mathcal{T}_{2,j}$ of maximum degree Δ with effective fields $R_{1,j}, R_{2,j}$ and magnetization gaps $M_{1,j}, M_{2,j}$ such that $R_{1,j} = R_{2,j}$ but $M_{1,j} \neq M_{2,j}$; moreover the values of $R_{1,j}$, and hence of $R_{2,j}$ as well, are pairwise distinct.*

The proof of Lemma 16 is given in Section 4.5 of the full version. With Lemmas 15 and 16 at hand, we can now give the proof of Theorem 12.

Proof of Theorem 12: For $i \in \{1, 2\}$, let $R_i := R_{\mathcal{T}_i}$ be the effective field of \mathcal{T}_i , $M_i := M_{\mathcal{T}_i}$ be the magnetization gap of \mathcal{T}_i and $A'_i := \mathbf{E}_{\tau \sim \mu_{\mathcal{T}_i; \beta, \gamma, \lambda}}[|\tau| \mid \tau(\rho_i) = 0]$. Note that $R_1 = R_2$ but $M_1 \neq M_2$, and A'_1, A'_2 are absolute constants.

Let $A(R), B(R), C(R)$ be the functions in Lemma 15 and set $D(R) = B(R) - C(R)$. Let A, B, C, D denote the common values of $A(R_i), B(R_i), C(R_i), D(R_i)$, respectively, for $i \in \{1, 2\}$. We will need to ensure in our argument that $D \neq 0$. The first observation is that (1) implies that for all but finitely many values of R we have that $D(R) \neq 0$.² The second observation is that, from Lemma 16, using $\mathcal{T}_1, \mathcal{T}_2$ we can construct for $j = 0, 1, 2, \dots$ an infinite sequence of pairs of field gadgets $\mathcal{T}_{1,j}, \mathcal{T}_{2,j}$ of maximum degree Δ with effective fields $R_{1,j}, R_{2,j}$ and magnetization gaps $M_{1,j}, M_{2,j}$ such that $R_{1,j} = R_{2,j}$ but $M_{1,j} \neq M_{2,j}$; moreover the values of $R_{1,j}$ are pairwise distinct. From these two observations, it follows that we can assume without loss of generality that $D \neq 0$.

Suppose for the sake of contradiction that, for arbitrarily small $\kappa > 0$, there is a polynomial-time algorithm that, on input a graph G of maximum degree Δ , produces a $(1 + \kappa)$ -approximation of the magnetization $\mathcal{M}_{G; \beta, \gamma, \lambda}$. We will show that we can approximate MAX-CUT on 3-regular graphs within a constant factor arbitrarily close to 1, contradicting the inapproximability result of [1].

For $i \in \{1, 2\}$, consider the graph H_{G, \mathcal{T}_i}^k for some large integer $k > 1 + 10/\log A$ and small $\epsilon > 0$ to be specified later. Let μ_i denote the Gibbs distribution on H_{G, \mathcal{T}_i}^k with parameters β, γ, λ . By Lemma 15, the average magnetization of the graph H_{G, \mathcal{T}_i}^k satisfies

$$\mathcal{M}_{\beta, \gamma, \lambda}(H_{G, \mathcal{T}_i}^k) = 4kA'_i|E(H)| + \mathbf{E}_{\sigma \sim \mu_i}[\left| \sigma_{V(\hat{H}_G^k)} \right|] + (1 \pm 8\epsilon)kM_iQ_i, \quad (2)$$

with $Q_i = D \text{AVG-CUT}_{\mu_i}(H) + C|E(H)|$; note that $\text{AVG-CUT}_{\mu_i}(H)$ satisfies

$$1/K_i \leq \frac{\text{AVG-CUT}_{\mu_i}(H)}{\text{MAX-CUT}(H)} \leq 1 \text{ for } K_i := 1 + \frac{6}{k \log A}. \quad (3)$$

Note that since $R_1 = R_2$ we have that

$$\mathbf{E}_{\sigma \sim \mu_1}[\left| \sigma_{V(\hat{H}_G^k)} \right|] = \mathbf{E}_{\sigma \sim \mu_2}[\left| \sigma_{V(\hat{H}_G^k)} \right|], \quad (4)$$

$$\text{AVG-CUT}_{\mu_1}(H) = \text{AVG-CUT}_{\mu_2}(H).$$

²If $B(R)$ and $C(R)$ agree on infinitely many values of R , then since both of them are rational functions of R , it must be the case that $B(R) = C(R)$ for all $R \geq 0$. Since for all $R \geq 0$, we have that $(\log A(R))' = B(R) - C(R)$, this would give that $A(R)$ is the constant function; contradiction, since (1) asserts that $A(0) = 1$ and $A(R) > 1$ for all $R > 0$.

We will denote by $\text{AVG-CUT}_\mu(H)$ the common value of $\text{AVG-CUT}_{\mu_1}(H)$, $\text{AVG-CUT}_{\mu_2}(H)$ and by Q the common value of Q_1, Q_2 .

Let $\mathcal{D} := \mathcal{M}_{\beta,\gamma,\lambda}(H_{G,\mathcal{T}_1}^k) - \mathcal{M}_{\beta,\gamma,\lambda}(H_{G,\mathcal{T}_2}^k)$. From (2) and (4), we have that

$$\mathcal{D} = 4k(\mathcal{A}'_1 - \mathcal{A}'_2)|E(H)| + (1 \pm 8\epsilon)k(M_1 - M_2)Q.$$

Let $L := |V(\mathcal{T}_1)| + |V(\mathcal{T}_2)|$. Since H_{G,\mathcal{T}_i}^k has at most $|V(H)||V(G)| + 4k|E(H)||V(\mathcal{T}_i)| \leq X$ vertices where $X := k(|V(G)| + 4L)|E(H)|$, using the purported algorithm for the magnetization on H_{G,\mathcal{T}_i}^k we can compute an estimate of $\mathcal{M}_{\beta,\gamma,\lambda}(H_{G,\mathcal{T}_i}^k)$ that is off by at most (an additive) κX for an arbitrarily small constant $\kappa > 0$ (that we will choose later). By subtracting these estimates for $i \in \{1, 2\}$, we obtain an estimate $\widehat{\mathcal{D}}$ of \mathcal{D} satisfying

$$|\widehat{\mathcal{D}} - \mathcal{D}| \leq 2k(|V(G)| + 4L)|E(H)|\kappa.$$

Then, we have that $\widehat{\text{MC}} = \frac{\widehat{\mathcal{D}} - 4k(\mathcal{A}'_1 - \mathcal{A}'_2)|E(H)|}{k(M_1 - M_2)\mathcal{D}} - \frac{\mathcal{D}}{\mathcal{D}}|E(H)|$ satisfies

$$|\widehat{\text{MC}} - \text{AVG-CUT}_\mu(H)| \leq \left(\frac{2(|V(G)| + 4L)}{|M_1 - M_2|} \kappa + 8 \left(1 + \frac{|C|}{|D|} \right) \epsilon \right) |E(H)|, \quad (5)$$

where $\text{AVG-CUT}_\mu(H)$ denotes the common value of $\text{AVG-CUT}_{\mu_1}(H)$, $\text{AVG-CUT}_{\mu_2}(H)$ (see (4)). By choosing the integer k sufficiently large, we have from (3) that $\text{AVG-CUT}_\mu(H)$ is within a factor arbitrarily close to 1 from $\text{MAX-CUT}(H)$. We will also choose $\epsilon > 0$ to be arbitrarily close to 0. This has the potential effect of increasing the size of G , but by choosing $\kappa > 0$ to be sufficiently small we can nevertheless ensure from (5) that our approximation $\widehat{\text{MC}}$ is within a factor arbitrarily close to 1 from $\text{AVG-CUT}_\mu(H)$, and hence from $\text{MAX-CUT}(H)$ as well. This finishes the contradiction argument and completes the proof of Theorem 12. \blacksquare

With Theorem 12, we can easily conclude the inapproximability result for the average-size of an independent set. Theorems 3, 5, 8 will follow analogously, once we give the field gadget constructions of Section II-A. The proofs of these Theorems are given in Section 5.4 of the full version.

Proof of Theorem 1: By Example 11, we have two trees \mathcal{T}_1 and \mathcal{T}_2 with the same effective fields and different magnetization gaps. Also, independent sets correspond to $\beta = 1, \gamma = 0, \lambda = 1$ and it is well-known (see, e.g., [15]) that $(\beta, \gamma, \lambda) \in \mathcal{U}_\Delta^*$ iff $\Delta \geq 6$. Therefore, the desired inapproximability result follows by applying Theorem 12. \blacksquare

The proof of Theorems 2, 4, and 7 build on similar reduction ideas, but the details are slightly more technical because of the different type of gadgets that we use (cf.

Theorem 13). The details can be found in Section 5.5 of the full version.

III. FIELD GADGET CONSTRUCTION

In this section, we give the proof of Theorem 13 by first outlining the key ideas of our field-gadget constructions. We start with the following lemma which describes the effective field and magnetization gap of field gadgets built out of smaller field gadgets.

Lemma 17. *Suppose we have trees $\mathcal{T}_1, \dots, \mathcal{T}_k$ with roots ρ_1, \dots, ρ_k , effective fields R_1, \dots, R_k and magnetization gaps M_1, \dots, M_k , respectively. Let \mathcal{T} be the tree with root ρ , edge $\{\rho, u\}$ and the roots of $\mathcal{T}_1, \dots, \mathcal{T}_k$ identified with u . Let R and M be the effective field and magnetization gap for \mathcal{T} . Then*

$$R = \frac{1 + \gamma \lambda \prod_{i=1}^k R_i}{\beta + \lambda \prod_{i=1}^k R_i}, \quad M = 1 - \omega(R) \left(1 + \sum_{i=1}^k (M_i - 1) \right),$$

where

$$\omega(R) := \frac{1 + \beta\gamma - \beta R - \gamma/R}{1 - \beta\gamma}.$$

For antiferromagnetic (β, γ) and any $\lambda > 0$, it holds that $R \in (\gamma, 1/\beta)$ and $0 < \omega(R) < 1$.

The first step in the construction is to create a family of field gadgets with sufficiently dense effective fields in an interval and magnetization gaps that are in a small interval. To define the interval, consider x^* and ω^* defined as follows

$$x^* = \frac{1 + \gamma\lambda(x^*)^2}{\beta + \lambda(x^*)^2}, \quad \omega^* = \frac{1 + \beta\gamma - \beta x^* - \gamma/x^*}{1 - \beta\gamma}, \quad (6)$$

and note that $\omega^* \in (0, 1)$ from Lemma 17.

Lemma 18. *Fix $\lambda > 0$. For any $\delta > 0$ and any sufficiently small $\tau_1 > 0$ we can find $\tau \in (0, \tau_1)$ and construct a family of field gadgets $\mathcal{L} = \{\mathcal{T}_1, \dots, \mathcal{T}_k\}$ whose effective fields are in the interval $[x^* - \tau, x^* + \tau]$ and for any $x \in [x^* - \tau, x^* + \tau]$ there exist a field gadget in the family whose effective field is in the interval $[x - \tau\delta, x + \tau\delta]$.*

The proof of Lemma 18 builds on techniques from [3] and is given in Section 4.2 of the full version.

The second step of the construction is to use the field gadgets in \mathcal{L} constructed in Lemma 18 in an iterative way as follows. At time $t = 0$, we will start with some field gadget \mathcal{T}_0 . Suppose that at some stage we have constructed a rooted field gadget \mathcal{T} with effective field R and magnetization gap M . To construct a new rooted field gadget, we merge \mathcal{T} with one field gadget from \mathcal{L} using the operation from Lemma 17. We are going to analyze what pairs (effective field, magnetization gap) we can achieve with this procedure. Let R_i and M_i be

the effective field and magnetization gap for the field gadget \mathcal{T}_i in our collection. By Lemma 17, merging the field gadget \mathcal{T} with \mathcal{T}_i yields a rooted field gadget with effective field $\phi_i(R)$ and magnetization gap $\psi_i(R, M)$ where the pair of maps (ϕ_i, ψ_i) are given by

$$\begin{aligned} \phi_i(R) &= \frac{1 + \gamma\lambda RR_i}{\beta + \lambda RR_i} \quad \text{and} \\ \psi_i(R, M) &= 1 - \omega(\phi_i(R))(M + M_i - 1). \end{aligned} \quad (7)$$

For a small constant $\tau > 0$ to be specified later, let

$$I' = \left[x^* - 2\tau \frac{|\omega^*|}{1 - |\omega^*|}, x^* + 2\tau \frac{|\omega^*|}{1 - |\omega^*|} \right] \quad (8)$$

The choice of the interval I' is such that the maps ϕ_i are uniformly contracting and map the interval I' to itself. Namely, we show the following in Section 4.3 of the full version.

Lemma 19. *There exists $0 < C_{\min} < C_{\max} < 1$ and $\tau_0 > 0$ such that for all $\tau \in (0, \tau_0)$ for all $R \in I'$ and all $R_i \in [x^* - \tau, x^* + \tau]$ it holds that*

$$C_{\min} \leq |\phi'_i(R)| \leq C_{\max}, \quad \omega(R) \leq C_{\max}, \quad \phi_i(R) \in I'$$

The second step of our construction will actually take place in the following smaller sub-interval of I' :

$$I = [x^* - |\omega|\tau/2, x^* + |\omega|\tau/2] \subseteq I', \quad (9)$$

The choice of I is to ensure the following “well-covered” property for the maps ϕ_i on I which is obtained using the density of the effective fields from Lemma 18, the proof is given in Section 4.3 of the full version.

Lemma 20. *Suppose $\delta < |\omega|/100$. For every two points $x_1, x_2 \in I$ such that $|x_1 - x_2| \leq |\omega|\tau\delta/2$ there exists i such that $x_1, x_2 \in \phi_i(I)$.*

Lemmas 19 and 20 ensure that for any $x \in I$ we can construct a sequence of field gadgets whose effective field approaches x . Consider the following process $\text{Build-gadget}(x, t)$ where $x \in I$ and $t \geq 0$ is an integer. If $t = 0$ we return the degenerate tree. If $t \geq 1$ then we let ϕ_i be any map such that $x \in \phi_i(I)$. Let $y = \phi_i^{-1}(x)$ and $\mathcal{T}' = \text{Build-gadget}(y, t - 1)$. Return the tree \mathcal{T} obtained by merging \mathcal{T}' and \mathcal{T}_i using the operation of Lemma 17.

The point behind the process $\text{Build-gadget}(x, t)$ is that it allows us to construct, for arbitrary $x \in I$, a field gadget whose effective field is arbitrary close to x , with error that decays exponentially fast with t (using the contraction properties of the ϕ_i 's). This is detailed in the following lemma.

Lemma 21. *There exists $C > 0$ such that for any $x \in I$ and any $t \geq 0$ the effective field R of the field gadget*

returned by $\text{Build-gadget}(x, t)$ (for any choice of the ϕ_i 's) satisfies $|R - x| \leq CC_{\max}^t$.

Lemma 21 is proved in Section 4.4 of the full version. For the field gadgets we construct we will always maintain the effective field in the interval I' , cf. Lemma 19. Now we show that the magnetization gaps of the field gadgets constructed using our process also stay restricted to an interval J . Let

$$T = \frac{2 + \max |M_i|}{1 - C_{\max}} \quad \text{and let } J \text{ be the interval } [-T, T]. \quad (10)$$

Lemma 22. *Suppose $R \in I'$ and $M \in J$ then $\psi_i(R, M) \in J$.*

Proof: We have $\psi_i(R, M) = 1 - \omega(R)(M_i + M - 1)$. Hence $|\psi_i(R, M)| \leq 1 + C_{\max}(T + \max |M_i| + 1) \leq T$, where the last inequality follows from (10). ■

For any $x \in I$ we are going to construct a sequence of families of pairs of maps $\mathcal{F}_{x,0}, \mathcal{F}_{x,1}, \dots$ as follows. In each pair, the first map is from \mathbb{R} to \mathbb{R} , and the second map is from \mathbb{R}^2 to \mathbb{R} (similarly to (7)). Let $\mathcal{F}_{x,0}$ contain the pair of maps $(x \mapsto x, (x, y) \mapsto y)$. To construct $\mathcal{F}_{x,t+1}$ we take every ϕ_i such that that $x \in \phi_i(I)$ and every $(f, g) \in \mathcal{F}_{\phi_i^{-1}(x), t}$ and place into $\mathcal{F}_{x,t+1}$ the map

$$(R, M) \mapsto (\phi_i(f(R)), \psi_i(f(R), g(R, M))). \quad (11)$$

Every $(f, g) \in \mathcal{F}_{x,t}$ corresponds to a sequence of ϕ_i 's with length t , which in turn corresponds to a rooted field gadget built using the procedure described just above (7) for t steps. The point is that we will view the starting field gadget \mathcal{T}_0 as an “input” to the procedure. If the input has effective field R and magnetization gap M then the final field gadget will have effective field $f(R)$ and magnetization gap $g(R, M)$. In particular, with the right choice of input (i.e., $f^{-1}(x)$), we will obtain a field gadget with effective field x . We will usually view $f(R)$ and $g(R, M)$ as the effective field and the magnetization gap induced by the $\text{Build-gadget}(x, t)$ procedure when at the base step we use \mathcal{T}_0 instead of the degenerate tree (provided the choice of the ϕ_i 's in Build-gadget matches up with the sequence for the pair (f, g)). However, we will not be able to use the exact input needed to obtain the effective field x but rather some approximation of it, and the pair (f, g) will allow us to track the influence of the last t steps when building a rooted field gadget \mathcal{T}' , which intuitively have larger influence both on the effective field and magnetization gap of \mathcal{T}' (since the maps ϕ_i are contracting). In particular, once we have the value of the magnetization gap and effective field of the “input” field gadget, by applying (f, g) we know precisely where the effective field and magnetization gap will end up after applying the sequence of ϕ_i 's corresponding to (f, g) .

We are going to distinguish two possible cases for the families $\mathcal{F}_{x,t}$. In the first case we will obtain the gadgets we need immediately using the Build-gadget procedure. In the second case we will construct a continuous function from the families. Then we will argue that the function cannot satisfy a functional equation and this will yield the gadgets we require.

Lemma 23 (Case I). *Suppose there exists $x \in I$ such that for some t_1, t_2 there exist $(f_1, g_1) \in \mathcal{F}_{x,t_1}$ and $(f_2, g_2) \in \mathcal{F}_{x,t_2}$ such that $g_1(I' \times J)$ and $g_2(I' \times J)$ are disjoint. Let \hat{M} be the distance of $g_1(I' \times J)$ and $g_2(I' \times J)$.*

Then, there is an algorithm which, on input rational $r > 0$, outputs in time $\text{poly}(\text{bits}(r))$ a pair of field gadgets \mathcal{T}_1 and \mathcal{T}_2 , each of maximum degree 3 and size $O(|\log r|)$, such that

$$|R_{\mathcal{T}_1} - x|, |R_{\mathcal{T}_2} - x| \leq r, \text{ but } |M_{\mathcal{T}_1} - M_{\mathcal{T}_2}| \geq \hat{M}.$$

Proof: Note that for a fixed t the union $\bigcup_x \mathcal{F}_{x,t}$ contains finitely many functions (because there are only finitely many choices for ϕ_i in each step of the construction). If t_1, t_2 and x assumed by the lemma exist then we can find them by examining finitely many functions. For each (f_1, g_1) and (f_2, g_2) we check whether $f_1(I) \cap f_2(I) \neq \emptyset$ and $g_1(I' \times J) \cap g_2(I' \times J) = \emptyset$; if we find such a pair we take $x \in f_1(I) \cap f_2(I)$. Note the running time for this process is a constant depending on β, γ, λ . (Note that $f(I)$ and $g(I' \times J)$ is always an interval that we can find inductively.)

To construct $\mathcal{T}_i, i \in \{1, 2\}$ we will use the procedure Build-gadget following the choices of (f_i, g_i) in the first t_i steps, cf. the discussion below (11). Using Lemma 21, we run the procedure for $O(|\log r|)$ steps achieving $|R_i - x| \leq r, i \in \{1, 2\}$. We have $M_i \in g_i(I' \times J), i \in \{1, 2\}$ and hence $|M_1 - M_2| \geq \hat{M}$. ■

Lemma 24 (Case II). *Suppose that for every x , every t_1, t_2 and every two functions $(f_1, g_1) \in \mathcal{F}_{x,t_1}$ and $(f_2, g_2) \in \mathcal{F}_{x,t_2}$ we have that $f_2(I' \times J)$ and $g_2(I' \times J)$ intersect. Then there exists a continuous function $F : I \rightarrow J$ such that for every $x \in I$, every $\epsilon > 0$ there exists t_0 such that for every $t \geq t_0$ and every $(f, g) \in \mathcal{F}_{x,t}$ and every $R \in I'$ and every $M \in J$ we have*

$$|g(R, M) - F(x)| \leq \epsilon. \quad (12)$$

The proof of Lemma 24 is given in Section 4.4 of the full version.

Lemma 25. *Suppose CASE II happens, that is, the assumption of Lemma 24 is satisfied; let F be the continuous function guaranteed by Lemma 24. Suppose that there exist $x_1, x_2 \in I$ such that the following*

equation is violated.

$$F\left(\frac{1 + \gamma\lambda x_1 x_2}{\beta + \lambda x_1 x_2}\right) = 1 - \frac{(1 - \beta\gamma)\lambda x_1 x_2 (F(x_1) + F(x_2) - 1)}{(1 + \gamma\lambda x_1 x_2)(\beta + \lambda x_1 x_2)} \quad (13)$$

Then, we can find in constant time (where the constant depends on $\beta, \gamma, \lambda, k$) rational numbers $x_1, x_2 \in I$ such that (13) is violated and, moreover, for any integer $k \geq 1$, rationals $\hat{R}_1, \dots, \hat{R}_k \in I$ such that the following holds.

There is an algorithm which on input a rational $r \in (0, 1/2)$ and any $i \in [k]$ outputs in time $\text{poly}(\text{bits}(r))$ a pair of field gadgets \mathcal{T}_1 and \mathcal{T}_2 , each of maximum degree 3 and size $O(|\log r|)$, such that $|R_{\mathcal{T}_1} - \hat{R}_i| \leq r, |R_{\mathcal{T}_2} - \hat{R}_i| \leq r$ and $|M_{\mathcal{T}_1} - M_{\mathcal{T}_2}| \geq \hat{M}$.

Proof of Lemma 25: Let $k \geq 1$ be an arbitrary integer and fix $x_1, x_2 \in I$ that violate (13). Let $\delta > 0$ be the absolute value of the difference between the two sides. From the continuity of F (cf. equation (53) of the full version), there exists $\epsilon > 0$ such that for any $y_1, y_2 \in I$ with $|y_1 - x_1| \leq \epsilon$ and $|y_2 - x_2| \leq \epsilon$ we have that y_1, y_2 violate (13) with difference at least $\delta/2$ between the two sides.

Let \mathcal{C} be a finite set of pairs (y_1, y_2) which form an $\frac{\epsilon}{20k}$ -net for $I \times I$; by perturbing slightly the set of points in \mathcal{C} we can obtain an $\frac{\epsilon}{10k}$ -net for $I \times I$, say \mathcal{C}' , such that for any two pairs (y_1, y_2) and (y'_1, y'_2) it holds that $y_1 y_2 \neq y'_1 y'_2$. Now, to check whether a pair (y_1, y_2) in \mathcal{C}' violates (12), we run Build-gadget for y_1, y_2 and $y_3 := \frac{1 + \gamma\lambda y_1 y_2}{\beta + \lambda y_1 y_2}$ with the value of t given by Lemma 24 to achieve bound $\delta/100$ on the right-hand side of (12). Then, we will find at least k different pairs $(y_{1,j}, y_{2,j}), j = 1, \dots, k$ that violate (13). For $j \in [k]$, let $\hat{R}_j := y_{3,j} = \frac{1 + \gamma\lambda y_{1,j} y_{2,j}}{\beta + \lambda y_{1,j} y_{2,j}}$ and note that by the construction of \mathcal{C}' , the \hat{R}_j 's are pairwise distinct.

Now, on input $j \in [k]$, we use the Build-gadget procedure to construct $\hat{\mathcal{T}}_i$ for $y_{i,j}, i \in \{1, 2, 3\}$ for $t = O(|\log r|)$ steps, using Lemma 21. The tree \mathcal{T}_1 is obtained by merging $\hat{\mathcal{T}}_1$ and $\hat{\mathcal{T}}_2$ and the tree \mathcal{T}_2 is $\hat{\mathcal{T}}_3$. ■

Now assume that equation (13) is satisfied for all $x_1, x_2 \in I$. We are going to derive a contradiction, thereby showing that (13) must be violated. We will do this in two steps. First we use a special case of equation (13) to obtain a functional equation that constrains the possible solutions of F . Second we show that none of these solutions satisfies (13).

Suppose $x_1, x_2, x_3 \in I$ are such that $x_1 x_2 = x_3 x^*$. Plugging into (13) we obtain that F has to satisfy the following equation

$$F(x_1) + F(x_2) = F(x_3) + F(x^*). \quad (14)$$

Lemma 26. *Suppose F is a continuous function on I . Suppose that for $x_1, x_2, x_3 \in I$ such that $x_1 x_2 = x_3 x^*$ we have (14). Then there exists c such that for all $x \in I$ we have*

$$F(x) = c \log(x/x^*) + F(x^*). \quad (15)$$

Proof: We will use the following parametrization to turn (14) into Cauchy's functional equation. Let $x_1 = x^* \exp(y_1)$, $x_2 = 2x^* \exp(y_2)$, and $x_3 = x^* \exp(y_3)$. The condition $x_1 x_2 = x^* x_3$ is equivalent to $y_1 + y_2 = y_3$. Let

$$G(y) = F(x^* \exp y) - F(x^*). \quad (16)$$

Note, we have that G is defined on the interval $[\log(I_L/x^*), \log(I_R/x^*)]$ that contains 0 (since $I_L < x^* < I_R$). Equation (14) becomes

$$G(y_1) + G(y_2) = G(y_1 + y_2). \quad (17)$$

From continuity of F , we also have continuity of G . Since (17) is Cauchy's functional equation on an interval containing zero the only continuous solutions are $G(y) = cy$ for some constant c . Plugging in (16) we obtain (15). ■

Finally we show in the full version that (13) has to be violated.

Lemma 27. *A solution of the form (15) cannot satisfy (13).*

We now combine the various pieces from the previous subsection to prove Theorem 13. In fact, we will prove a slight strengthening of Theorem 13, given below, which will be used in our reduction, analogously to Lemma 16.

Theorem 28. *Let (β, γ, λ) be antiferromagnetic with $(\beta, \gamma, \lambda) \neq (\beta, \beta, 1)$. For every integer $k \geq 1$, there exist constants $\hat{M}, \Xi > 0$ and k distinct numbers $\hat{R}_1, \dots, \hat{R}_k > 0$ such that the following holds. There is an algorithm, which, on input $i \in [k]$ and a rational $r \in (0, 1/2)$, outputs in time $\text{poly}(\text{bits}(r))$ a pair of field gadgets $\mathcal{T}_1, \mathcal{T}_2$, each of maximum degree 3 and size $O(|\log r|)$, such that*

$$|R_{\mathcal{T}_1} - \hat{R}_i|, |R_{\mathcal{T}_2} - \hat{R}_i| \leq r, \text{ but } |M_{\mathcal{T}_1} - M_{\mathcal{T}_2}| \geq \hat{M}.$$

Also, the magnetization gaps $M_{\mathcal{T}_1}, M_{\mathcal{T}_2}$ are bounded in absolute value by the constant Ξ .

Proof: There are two cases to consider: either Lemma 23 or Lemma 24. In the latter case, by Lemmas 26 and 27, we obtain that (13) is violated, and hence Lemma 25 yields the desired algorithm. In the former case, we are also done, modulo that Lemma 23 only guarantees the existence of a single \hat{R} , namely the value x . Let t_1, t_2 be such that $(f_1, g_1) \in \mathcal{F}_{x, t_1}$ and $(f_2, g_2) \in \mathcal{F}_{x, t_2}$ satisfy $g_1(I' \times J) \cap g_2(I' \times J) = \emptyset$. The functions in $\mathcal{F}_{x, t}$ are continuous and defined on a

closed interval, so we have that for sufficiently small $\epsilon > 0$, for all y such that $|y - x| \leq \epsilon$, we can find $(\tilde{f}_1, \tilde{g}_1) \in \mathcal{F}_{x, t_1}$ and $(\tilde{f}_2, \tilde{g}_2) \in \mathcal{F}_{x, t_2}$ which satisfy $\tilde{g}_1(I' \times J) \cap \tilde{g}_2(I' \times J) = \emptyset$. We pick k such y 's for the values of the \hat{R}_i 's for $i = 1, \dots, k$ and produce the desired trees by running the Build-gadget procedure for $O(\log |r|)$ steps. Finally note that the magnetizations all lie in the interval J , see (10) and Lemma 22, which is bounded by absolute constants, finishing the proof. ■

Note that Theorem 13 corresponds to the case $k = 1$ in Theorem 28.

ACKNOWLEDGMENT

Daniel Štefankovič supported in part by NSF grant CCF-1563757. Eric Vigoda was supported in part by NSF grant CCF-1563838.

REFERENCES

- [1] Paola Alimonti and Viggo Kann. Hardness of approximating problems on cubic graphs. In *Algorithms and Complexity*, pages 288–298, 1997.
- [2] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. *arXiv preprint arXiv:2001.00303*, 2020.
- [3] Ivona Bezáková, Andreas Galanis, Leslie Ann Goldberg, and Daniel Štefankovič. Inapproximability of the independent set polynomial in the complex plane. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, page 1234–1240, 2018.
- [4] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability for antiferromagnetic spin systems in the tree nonuniqueness region. *J. ACM*, 62(6), 2015.
- [5] Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *Combinatorics, Probability and Computing*, 25(4):500–559, 2016.
- [6] Mark Jerrum and Alistair Sinclair. Polynomial-time approximation algorithms for the Ising model. *SIAM Journal on Computing*, 22(5):1087–1116, 1993.
- [7] Liang Li, Pinyan Lu, and Yitong Yin. Correlation decay up to uniqueness in spin systems. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 67–84, 2013.
- [8] Fabio Martinelli, Alistair Sinclair, and Dror Weitz. Fast mixing for independent sets, colorings, and other models on trees. *Random Structures & Algorithms*, 31(2):134–172, 2007.
- [9] Viresh Patel and Guus Regts. Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials. *SIAM Journal on Computing*, 46(6):1893–1919, 2017.

- [10] Han Peters and Guus Regts. On a conjecture of Sokal concerning roots of the independence polynomial. *Michigan Math. J.*, 68(1):33–55, 2019.
- [11] Ricardo Restrepo, Daniel Štefankovič, Juan C. Vera, Eric Vigoda, and Linji Yang. Phase transition for Glauber dynamics for independent sets on regular trees. *SIAM Journal on Discrete Mathematics*, 28(2):835–861, 2014.
- [12] Leonard J. Schulman, Alistair Sinclair, and Piyush Srivastava. Symbolic integration and the complexity of computing averages. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1231–1245, 2015.
- [13] Alistair Sinclair and Piyush Srivastava. Lee-Yang theorems and the complexity of computing averages. *Communications in Mathematical Physics*, 329(3):827–858, 2014.
- [14] Alistair Sinclair, Piyush Srivastava, and Marc Thurley. Approximation algorithms for two-state antiferromagnetic spin systems on bounded degree graphs. *J. Stat. Phys.*, 155(4):666–686, 2014.
- [15] Allan Sly. Computational transition at the uniqueness threshold. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 287–296, 2010.
- [16] Allan Sly and Nike Sun. Counting in two-spin models on d -regular graphs. *Ann. Probab.*, 42(6):2383–2416, 11 2014.
- [17] Dror Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 140–149, 2006.