THE ISING PARTITION FUNCTION: ZEROS AND DETERMINISTIC APPROXIMATION

JINGCHENG LIU, ALISTAIR SINCLAIR, AND PIYUSH SRIVASTAVA

Abstract. We study the problem of approximating the partition function of the ferromagnetic Ising model in graphs and hypergraphs. Our first result is a *deterministic* approximation scheme (an FPTAS) for the partition function in bounded degree graphs that is valid over the entire range of parameters β (the interaction) and λ (the external field), except for the case $|\lambda| = 1$ (the "zero-field" case). A *randomized* algorithm (FPRAS) for all graphs, and all β, λ, has long been known. Unlike most other deterministic approximation algorithms for problems in statistical physics and counting, our algorithm does not rely on the "decay of correlations" property. Rather, we exploit and extend machinery developed recently by Barvinok, and Patel and Regts, based on the location of the complex zeros of the partition function, which can be seen as an algorithmic realization of the classical Lee-Yang approach to phase transitions. Our approach extends to the more general setting of the Ising model on hypergraphs of bounded degree and edge size, where no previous algorithms (even randomized) were known for a wide range of parameters. In order to achieve this extension, we establish a tight version of the Lee-Yang theorem for the Ising model on hypergraphs, improving a classical result of Suzuki and Fisher.

Jingcheng Liu, Computer Science Division, UC Berkeley. Email: liuexp@berkeley.edu.
Alistair Sinclair, Computer Science Division, UC Berkeley. Email: sinclair@cs.berkeley.edu.
Piyush Srivastava, Tata Institute of Fundamental Research, Mumbai. Email: piyush.srivastava@tifr.res.in.
JL and AS are supported in part by NSF grant CCF-1420934.

1. Introduction

The Ising model, first studied a century ago as a model for magnetic phase transitions by Lenz and Ising [21], has since become an important tool for the modeling of interacting systems. In the Ising model, such a system is represented as a graph $G = (V, E)$, so that the individual entities comprising the system correspond to the vertices V and their pairwise interactions to the edges E. A *configuration* of the system is an assignment $\boldsymbol{\sigma} : V \to \{+, -\}$ of one of two possible values (often called "spins") to each vertex. The model then induces a probability distribution (known as a *Gibbs distribution*) over these global configurations in terms of local parameters that model the interactions of the vertices.

In our setting, it will be convenient to parameterize these interactions in terms of a *vertex activity* λ (sometimes called an "external field") that determines the propensity of a vertex to be in the $+$ configuration, and an *edge activity* $\beta \geq 0$ that models the tendency of vertices to agree with their neighbors. The model assigns to each configuration $\boldsymbol{\sigma}$ a weight

$$w(\boldsymbol{\sigma}) := \beta^{|\{\{u,v\}\in E \mid \boldsymbol{\sigma}(u)\neq\boldsymbol{\sigma}(v)\}|}\lambda^{|\{v \mid \boldsymbol{\sigma}(v)=+\}|} = \beta^{|E(S,\overline{S})|}\lambda^{|S|},$$

where $S = S(\boldsymbol{\sigma})$ is the set of vertices assigned spin + in $\boldsymbol{\sigma}$ and $E(S, \overline{S})$ is the set of edges in the cut $(S, \overline{S})$ (i.e., the number of pairs of adjacent vertices assigned opposite spins). The probability of configuration $\boldsymbol{\sigma}$ under the Gibbs distribution is then $\mu(\boldsymbol{\sigma}) := w(\boldsymbol{\sigma})/Z_G^{\beta}(\lambda)$, where the normalizing factor $Z_G^{\beta}(\lambda)$ is the *partition function* defined as

$$Z_G^{\beta}(\lambda) := \sum_{\boldsymbol{\sigma}:V\to\{+,-\}} w(\boldsymbol{\sigma}) = \sum_{S\subseteq V} \beta^{|E(S,\overline{S})|}\lambda^{|S|}. \tag{1}$$

Notice that the partition function may be interpreted combinatorially as a cut generating polynomial in the graph G.

In this paper we focus on the original *ferromagnetic* case in which $\beta < 1$, so that configurations in which a larger number of neighboring spins agree (small cuts) have higher probability. The *anti-ferromagnetic* regime $\beta > 1$ is qualitatively very different, and prefers configurations with disagreements between neighbors. We note also that all our results in this paper hold in the more general setting where there is a different interaction β_e on each edge, provided that all the β_e satisfy whatever constraints we are putting on β. (So, e.g., in the ferromagnetic case $\beta_e < 1$ for all e.) In addition, our results allow β to be negative and λ to be complex; this will be discussed in more detail below.

As in almost all statistical physics and graphical models, the partition function captures the computational complexity of the Ising model. Partition functions are #P-hard to compute in virtually any interesting case (e.g., this is true for the Ising model except in the trivial cases $\lambda = 0$ or $\beta \in \{0, 1\}$), so attention is focused on approximation. An early result in the field due to Jerrum and Sinclair [22] establishes a *fully polynomial randomized approximation scheme* for the Ising partition function, valid for all graphs G and all values of the parameters (β, λ) in the ferromagnetic regime. Like many of the first results on approximating partition functions, this algorithm is based on random sampling and the Markov chain Monte Carlo method.

For several statistical physics models on bounded degree graphs (including the anti-ferromagnetic Ising model [25, 41] and the "hard core", or independent set model [45]), fully-polynomial *deterministic* approximation schemes have since been developed, based on the decay of correlations property in those models: roughly speaking, one can estimate the local contribution to the partition function at a given vertex v by exhaustive enumeration in a neighborhood around v, using decay of correlations to truncate the neighborhood at logarithmic diameter. The range of applicability of these algorithms is of course limited to the regime in which decay of correlations holds, and indeed complementary results prove that the partition function is NP-hard to approximate outside this regime [16, 42]. Perhaps surprisingly, however, no deterministic approximation algorithm is known for the classical ferromagnetic Ising partition function that works over anything close to the full range of the randomized algorithm of [22]: to the best of our knowledge, the best deterministic algorithm, due to Zhang, Liang and Bai [47], is based on correlation decay and is applicable to graphs of maximum degree Δ only when $\beta > (\Delta - 1)/(\Delta + 1)$.

The restricted applicability of correlation decay based algorithms for the ferromagnetic Ising model arises from two related reasons: the first is that this model does not exhibit correlation decay for β sufficiently close to 0 (for any given value of the external field), so any straightforward approach based only on this property cannot be expected to work for all β. Secondly, there is a regime of parameters for which, even though decay of correlation holds, there is evidence to believe that it cannot be exploited to give an algorithm using the usual techniques: see [41, Appendix 2] for a more detailed discussion of this point.

The first goal of this paper is to supply such a deterministic algorithm which covers almost the entire range of parameters of the model except for the "zero-field" case $|\lambda| = 1$:

Theorem 1.1. *Fix any $\Delta > 0$. There is a fully polynomial time approximation scheme (FPTAS) for the Ising partition function $Z_G^\beta(\lambda)$ in all graphs G of maximum degree Δ for all edge activities $-1 \le \beta \le 1$ and all (possibly complex) vertex activities λ with $|\lambda| \neq 1$.*

Remark. Note that although λ, β are positive in the "physically relevant" range in most applications of the Ising model, the above theorem also applies more generally to $\beta \in [-1, 1]$ and complex valued λ not on the unit circle. Moreover, we can allow edge-dependent activities β_e provided all of them lie in $[-1, 1]$.

The above theorem is actually a special case of a more general theorem for the hypergraph version of the Ising model (Theorem 1.3 below). We now illustrate our approach to proving these theorems, which will also allow us to introduce and motivate our further results in the paper.

In contrast to previous algorithms based on correlation decay, our algorithm is based on isolating the complex zeros of the partition function $Z := Z_G^\beta(\lambda)$ (viewed as a polynomial in λ for a fixed value of β). This approach was introduced recently by Barvinok [6, 7] in the context of models different from the Ising model. We start with the observation that the $1 + \varepsilon$ multiplicative approximation of Z required for a FPTAS corresponds to a $O(\varepsilon)$ additive approximation of $\log Z$. Barvinok's approach considers a Taylor expansion of $\log Z$ around a point λ_0 such that $Z(\lambda_0)$ is easy to evaluate. (For the Ising model, $\lambda_0 = 0$ is such a choice.) It then seeks to evaluate the function at other points by carrying out an explicit analytic continuation. More concretely, suppose it can be shown that there are no zeros of Z in the open disk $D(\lambda_0, r)$ of radius r around λ_0. From standard results in complex analysis, it then follows that the Taylor expansion around λ_0 of $\log Z$ is absolutely convergent in $D(\lambda_0, r)$ and further, that the first k terms of this expansion evaluated at a point $\lambda \in D(\lambda_0, r)$ provide a $O\left(\frac{|\lambda - \lambda_0|^k}{r^k}\right)$ additive approximation of $\log Z(\lambda)$. We note that Barvinok's approach may be seen as an algorithmic counterpart of the traditional view of phase transitions in statistical physics in terms of analyticity of $\log Z$ [46].

To apply this approach in the case of the ferromagnetic Ising model, we may appeal to the famous Lee-Yang theorem of the 1950s [24], which establishes that the zeros of $Z(\lambda)$ all lie on the unit circle in the complex plane. This allows us to take $\lambda_0 = 0$ and $r = 1$ in the previous paragraph, and thus approximate $Z(\lambda)$ at any point λ satisfying $|\lambda| < 1$. This extends to all λ with $|\lambda| \neq 1$ via the fact that $Z(\lambda) = \lambda^n Z(\frac{1}{\lambda})$.

Remark. We note that the case $|\lambda| = 1$ is likely to require additional ideas because it is known that there exist bounded degree graphs (namely Δ-ary trees) for which the partition function $Z_G^\beta(\lambda)$ has complex zeros within distance $O(1/n)$ of $\lambda = 1$, where n is the size of the graph. In fact, the zeros are even known to become dense on the whole unit circle as n increases to infinity [4, 5]. This precludes the possibility of efficiently carrying out the analytic continuation procedure for $\log Z$ to arbitrary points on the unit circle, and to the point $\lambda = 1$ in particular.

Converting the above approach into an algorithm requires computing the first k coefficients in the Taylor expansion of $\log Z$ around λ_0. Barvinok showed that this computation can in turn be reduced to computing $O(k)$ leading coefficients of the partition function itself. In the case of the Ising model, computing k such coefficients is roughly analogous to computing k-wise correlations between the vertex spins, and doing so naively on a graph of n vertices requires time $\Omega(n^k)$. Until recently, no better ways to perform this calculation were known and hence approximation algorithms using this approach typically required quasi-polynomial time in order to achieve a $(1 + 1/\operatorname{poly}(n))$-factor multiplicative approximation of Z (equivalently, a $1/\operatorname{poly}(n)$ additive approximation of $\log Z$), since this requires the Taylor series for $\log Z$ to be evaluated to $k = \Omega(\log n)$ terms [7–9].

Recently, Patel and Regts [35] proposed a way to get around this barrier for several classes of partition functions. Their method is based on writing the coefficients in the Taylor series of $\log Z$ as linear combinations of counts of connected induced subgraphs of size up to $\Theta(\log n)$. This is already promising, since the number of connected induced subgraphs of size $O(\log n)$ of a graph G of maximum degree Δ is polynomial in the size of G when Δ is a fixed constant. Further, the count of induced copies of such a subgraph can also be computed in time polynomial in the size of G [14]. Patel and Regts built on these tools to show a way to compute the $\Theta(\log n)$ Taylor coefficients of $\log Z$ needed in Barvinok's approach for several families of partition functions, for some of which correlation decay based algorithms are still not known.

Unfortunately, for the case of the Ising model, it is not clear how to write the Taylor coefficients in terms of induced subgraph counts. Indeed, in their paper [35, Theorem 1.4], Patel and Regts apply their method to the Ising model viewed as a polynomial in β rather than λ, which allows them to use subgraph counts. However, this prevents them from using the Lee-Yang theorem, and they are consequently able to establish only a small zero-free region. As a result, they can handle only the zero-field "high-temperature" regime (where in fact the correlation decay property also holds), specifically the regime $|\beta - 1| \leq 0.34/\Delta$ and $\lambda = 1$.

In this paper, we instead propose a generalization of the framework of Patel and Regts to objects that we call *insects*. An insect is a graph in which each vertex is decorated with an additional number of dangling "half edges" attached to it: we refer to section 3.1 for precise definitions. Using the appropriate notions for counting induced sub-insects, we are then able to write the coefficients arising in the Taylor expansion of $\log Z$ for the Ising model in terms of induced sub-insect counts, and derive from there algorithms for computing $\Omega(\log n)$ such coefficients in polynomial time in graphs of bounded degree. This establishes Theorem 1.1. We note that if one is only interested in deriving Theorem 1.1, then this can also be done using the notion of *fragments*, developed by Patel and Regts [35] in the different context of edge coloring models, which turns out to be a special case of our notion of insects.

Our framework of insects, however, also allows us to extend the above approach to the more general setting where G is a hypergraph (and further, when the edge activities are edge-dependent), as studied, for example, in the work of Suzuki and Fisher [44]. We note that the Jerrum-Sinclair MCMC approach [22] apparently does not extend to hypergraphs, and there is currently no known polynomial time approximation algorithm (even randomized) for a wide range of β in this setting. For a hypergraph $H = (V, E)$, configurations are again assignments of spins to the vertices V and the partition function $Z_H^\beta(\lambda)$ is defined exactly as in (1), where the cut $E(S, \overline{S})$ consists of those hyperedges with at least one vertex in each of S and $\overline{S}$. The computation of coefficients via insects carries through as before, but the missing ingredient is a extension of the Lee-Yang theorem to hypergraphs. Suzuki and Fisher [44] prove a weak version of the Lee-Yang theorem for hypergraphs (see Theorem 4.3 in section 4), which we are able to strengthen to obtain the following optimal statement, which is of independent interest:

Theorem 1.2. *Let $H = (V, E)$ be a hypergraph with maximum hyperedge size $k \geq 3$. Then all the zeros of the Ising model partition function $Z_H^\beta(\lambda)$ lie on the unit circle if and only if the edge activity β lies in the range $-\frac{1}{2^{k-1}-1} \leq \beta \leq \frac{1}{2^{k-1}\cos^{k-1}\left(\frac{\pi}{k-1}\right)+1}$.*

Remark. Once again, we can allow edge-dependent activities β_e provided all of them lie in the range stipulated above. This extension also applies to Theorem 1.3 below.

Note that the classical Lee-Yang theorem (for the graph case $k = 2$) establishes this property for $0 \leq \beta \leq 1$ (the ferromagnetic regime). Our proof of Theorem 1.2 follows along the lines of Asano's inductive proof of the Lee-Yang theorem [3], but we need to carefully analyze the base case (where H consists of a single hyperedge) in order to obtain the above bounds on β. Since our analysis of the base case is tight, the range of β in our theorem is optimal. For a detailed comparison with the Suzuki-Fisher theorem, see the Remark following Corollary 4.5.

Combining Theorem 1.2 with our earlier algorithmic approach immediately yields the following generalization of Theorem 1.1 to hypergraphs.

Theorem 1.3. *Fix any $\Delta > 0$ and $k \geq 3$. There is an FPTAS for the Ising partition function $Z_H^\beta(\lambda)$ in all hypergraphs H of maximum degree Δ and maximum edge size k, for all edge activities β in the range of Theorem 1.2 and all vertex activities $|\lambda| \neq 1$.*

Finally, we extend our results to general ferromagnetic two-spin systems on hypergraphs, again as studied in [44]. A *two-spin system* on a hypergraph $H = (V, E)$ is specified by hyperedge activities $\varphi_e : \{+,-\}^{|e|} \to \mathbb{C}$ for $e \in E$, and a vertex activity $\psi : \{+,-\} \to \mathbb{C}$. (Note that we treat each hyperedge e as a set of vertices.) Without loss of generality, we assume $\varphi_e(-,\cdots,-) = 1$, and $\psi(-) = 1$, $\psi(+) = \lambda$. Then the partition function is defined as:

$$Z_H^{\boldsymbol{\varphi}}(\lambda) := \sum_{\boldsymbol{\sigma}:V\to\{+,-\}} \prod_{e\in E} \varphi_e\left(\boldsymbol{\sigma}\big|_e\right) \prod_{v\in V} \psi(\boldsymbol{\sigma}(v)) = \sum_{\boldsymbol{\sigma}:V\to\{+,-\}} \prod_{e\in E} \varphi_e\left(\boldsymbol{\sigma}\big|_e\right) \lambda^{|\{v:\boldsymbol{\sigma}(v)=+\}|}. \tag{2}$$

We call a hypergraph two-spin system *symmetric* if $\varphi_e(\boldsymbol{\sigma}) = \overline{\varphi_e(-\boldsymbol{\sigma})}$. Suzuki and Fisher [44] prove a Lee-Yang theorem for symmetric hypergraph two-spin systems (which is weaker than our Theorem 1.2 above when specialized to the Ising model). Combining this with our general algorithmic approach yields our final result:

Theorem 1.4. *Fix any $\Delta > 0$ and $k \geq 2$ and a family of symmetric edge activities $\boldsymbol{\varphi} = \{\varphi_e\}$ satisfying $|\varphi_e(+,\cdots,+)| \geq \frac{1}{4}\sum_{\boldsymbol{\sigma}\in\{+,-\}^V} |\varphi_e(\boldsymbol{\sigma})|$. Then there exists an FPTAS for the partition function $Z_H^{\boldsymbol{\varphi}}(\lambda)$ of the corresponding symmetric hypergraph two-spin system in all hypergraphs H of maximum degree Δ and maximum edge size k for all vertex activities $\lambda \in \mathbb{C}$ such that $|\lambda| \neq 1$.*

The remainder of the paper is organized as follows. In section 2, we spell out Barvinok's approach to approximating partition functions using Taylor series. Section 3 introduces the notion of insects and shows how to use them to efficiently compute the leading coefficients of the partition function in the general context of hypergraphs; as discussed above, this machinery applied to graphs, in conjunction with the Lee-Yang theorem, implies Theorem 1.1. Finally, in section 4 we prove our extension of the Lee-Yang theorem to the hypergraph Ising model (Theorem 1.2), and then use it and the Suzuki-Fisher theorem to prove our algorithmic results for hypergraphs, Theorems 1.3 and 1.4.

1.1. **Related work.** The problem of computing partition functions has been widely studied, not only in statistical physics but also in combinatorics, because the partition function is often a generating function for combinatorial objects (cuts, in the case of the Ising model). There has been much progress on *dichotomy theorems*, which attempt to completely classify these problems as being either #P-hard or computable (exactly) in FP (see, e.g., [15, 17]).

Since the problems are in fact #P-hard in most cases, algorithmic interest has focused largely on *approximation*, motivated also by the general observation that approximating the partition function is polynomial time equivalent to sampling (approximately) from the underlying Gibbs distribution [23]. In fact, most early approximation algorithms exploited this connection, and gave *fully-polynomial randomized approximation schemes (FPRAS)* for the partition function using Markov chain Monte Carlo (MCMC) samplers for the Gibbs distribution. In particular, for the ferromagnetic Ising model, the MCMC-based algorithm of Jerrum and Sinclair [22] is valid for all positive real values of λ and for all graphs, irrespective of their vertex degrees. (For the connection with random sampling in this case, see [36].) This was later extended to ferromagnetic two-spin systems by Goldberg, Jerrum and Paterson [18]. Similar techniques have been applied recently to the related random-cluster model by Guo and Jerrum [19].

Much detailed work has been done on MCMC for Ising spin configurations for several important classes of graphs, including two-dimensional lattices (e.g., [28, 31, 32]), random graphs and graphs of bounded degree (e.g., [34]), the complete graph (e.g., [26]) and trees (e.g., [10, 33]); we do not attempt to give a comprehensive summary of this line of work here.

In the *antiferromagnetic* regime ($\beta > 1$), *deterministic* approximation algorithms based on correlation decay have been remarkably successful for graphs of bounded degree. Specifically, for any fixed integer $\Delta \geq 3$, techniques of Weitz [45] give a (deterministic) FPTAS for the antiferromagnetic Ising partition function on graphs of maximum degree Δ throughout a region R_Δ in the (β, λ) plane (corresponding the regime of "uniqueness of the Gibbs measure on the Δ-regular tree") [25, 41]. A complementary result of Sly and Sun [42] (see also [16]) shows that the problem is NP-hard outside R_Δ. In contrast, no MCMC based algorithms are known to provide an FPRAS for the anti-ferromagnetic Ising partition function throughout R_Δ. More recently, correlation decay techniques have been extended to obtain deterministic approximation algorithms for the anti-ferromagnetic Ising partition function on hypergraphs over a range of parameters [27], as well as to several other problems not related to the Ising model. In the ferromagnetic setting, however, for reasons mentioned earlier, correlation decay techniques have had more limited success: Zhang, Liang and Bai [47] handle only the "high-temperature" regime of the Ising model, while the recent results for ferromagnetic two-spin systems of Guo and Lu [20] do not apply to the case of the Ising model.

In a parallel line of work, Barvinok initiated the study of Taylor approximation of the logarithm of the partition function, which led to quasipolynomial time approximation algorithms [6–9]. More recently, Patel and Regts [35] showed that for many models that can be written as induced subgraph sums, one can actually obtain an FPTAS from this approach. In particular, for problems such as counting independent sets with negative (or, more generally, complex valued) activities on bounded degree graphs, they were able to match the range of applicability of existing algorithms based on correlation decay, and were also able to extend the approach to Tutte polynomials and edge-coloring models (also known as Holant problems) where little is known about correlation decay.

The Lee-Yang program was initiated by Lee and Yang [46] in connection with the analysis of phase transitions. By proving the famous Lee-Yang circle theorem for the ferromagnetic Ising model [24], they were able to conclude that there can be at most one phase transition for the model. Asano [3] extended the Lee-Yang theorem to the Heisenberg model, and provided a simpler proof. Asano's work was generalized further by Suzuki and Fisher [44]. A complete characterization of Lee-Yang polynomials that are independent of the "temperature" of the model was recently obtained by Ruelle [37]. The study of Lee-Yang type theorems for other statistical physics models has also generated beautiful connections with other areas of mathematics. For example, Shearer [39] and Scott and Sokal [38] established the close connection between the location of the zeros of the independence polynomial and the Lovász Local Lemma, while the study of the zeros of generalizations of the matching polynomial was used in the recent celebrated work of Marcus, Spielman and Srivastava on the existence of Ramanujan graphs [29]. Such Lee-Yang theorems are exemplars of the more general stability theory of polynomials [11, 12], a field of study that has had numerous recent applications to theoretical computer science and combinatorics (see, e.g., [1, 2, 13, 29, 30, 40]).

2. Approximation of the log-partition function by Taylor series

In this section we present an approach due to Barvinok [7] for approximating the partition function of a physical system by truncating the Taylor series of its logarithm, as discussed in the introduction. We will work in our most general setting of symmetric two-spin systems on hypergraphs, which of course includes the Ising model (on graphs or hypergraphs) as a special case. As in (2), such a system has partition function

$$Z_H^{\boldsymbol{\varphi}}(\lambda) = \sum_{\boldsymbol{\sigma}:V\to\{+,-\}} \prod_{e\in E} \varphi_e\big(\boldsymbol{\sigma}\big|_e\big)\, \lambda^{|\{v:\boldsymbol{\sigma}(v)=+\}|}.$$

Our goal is a FPTAS for $Z_H^{\boldsymbol{\varphi}}(\lambda)$, i.e., a deterministic algorithm that, given as input H, $\{\varphi_e\}$, λ with $|\lambda| \neq 1$ and $\varepsilon \in (0,1]$, runs in time polynomial in $n = |H|$ and ε^{-1} and outputs a $1+\varepsilon$ multiplicative approximation of $Z_H^{\boldsymbol{\varphi}}(\lambda)$, i.e., a number $\hat{Z}$ satisfying

$$|\hat{Z} - Z_H^{\boldsymbol{\varphi}}(\lambda)| \leq \varepsilon |Z_H^{\boldsymbol{\varphi}}(\lambda)|. \tag{3}$$

(Note that in our setting $\hat{Z}$ and $Z_H^{\boldsymbol{\varphi}}(\lambda)$ may be complex numbers.)

For fixed H and (hyper)edge activities $\boldsymbol{\varphi}$, we will write $Z(\lambda) = Z_H^{\boldsymbol{\varphi}}(\lambda)$ for short. By the symmetry $\varphi_e(\boldsymbol{\sigma}) = \overline{\varphi_e(-\boldsymbol{\sigma})}$, we have $Z(\lambda) = \lambda^n Z(\frac{1}{\lambda})$. Thus without loss of generality we may assume $|\lambda| < 1$. Letting $f(\lambda) = \log Z(\lambda)$, using the Taylor expansion around $\lambda = 0$ we get

$$f(\lambda) = \sum_{j=0}^{\infty} f^{(j)}(0) \cdot \frac{\lambda^j}{j!},$$

where $f(0) = \log Z(0) = 0$. Note that $Z = \exp(f)$, and thus an additive error in f translates to a multiplicative error in Z. More precisely, given $\varepsilon \leq 1/4$, and $f, \widetilde{f} \in \mathbb{C}$ such that $\left| f - \widetilde{f} \right| \leq \varepsilon$, we have

$$\left| \exp\left(\widetilde{f}\right) - \exp(f) \right| = \left| \exp\left(\widetilde{f} - f\right) - 1 \right| |\exp(f)| \leq 4\varepsilon \, |\exp(f)| \,,$$

where the last inequality, valid for $\varepsilon \leq 1/4$, follows by elementary complex analysis. In other words, to achieve a multiplicative approximation of Z within a factor $1+\varepsilon$, as required by a FPTAS, it suffices to obtain an additive approximation of f within $\varepsilon/4$.

To get an additive approximation of f, we use the first m terms in the Taylor expansion. Specifically, we compute $f_m(\lambda) = \sum_{j=0}^{m} f^{(j)}(0) \cdot \frac{\lambda^j}{j!}$.

To compute $f^{(j)}(0)$, note that $f'(\lambda) = \frac{1}{Z(\lambda)} \frac{\mathrm{d}Z(\lambda)}{\mathrm{d}\lambda}$, or $\frac{\mathrm{d}Z(\lambda)}{\mathrm{d}\lambda} = f'(\lambda) Z(\lambda)$. Thus for any $m \geq 1$,

$$\frac{\mathrm{d}^m}{\mathrm{d}\lambda^m} Z(\lambda) = \sum_{j=0}^{m-1} \binom{m-1}{j} \frac{\mathrm{d}^j}{\mathrm{d}\lambda^j} Z(\lambda) \cdot \frac{\mathrm{d}^{m-j}}{\mathrm{d}\lambda^{m-j}} f(\lambda). \tag{4}$$

Given $\frac{\mathrm{d}^j}{\mathrm{d}\lambda^j} Z(\lambda) \Big|_{\lambda=0}$ for $j = 0, \cdots, m$, eq. (4) is a non-degenerate (recall that $Z(0) = 1$) triangular system of linear equations in $\left\{ f^{(j)}(0) \right\}_{j=1}^{m}$, which can be solved in $O(m^2)$ time.

We can now specify the algorithm: first compute $\left\{ \frac{\mathrm{d}^j}{\mathrm{d}\lambda^j} Z(\lambda) \Big|_{\lambda=0} \right\}_{j=0}^{m}$; next, use the system in eq. (4) to solve for $\left\{ f^{(j)}(0) \right\}_{j=1}^{m}$; and finally, compute and ouput the approximation $f_m(\lambda)$.

To quantify the approximation error in this algorithm, we need to study the locations of the complex roots $r_1, \cdots, r_n$ of Z. Throughout this paper, we will be using (some variant of) the Lee-Yang theorem to argue that, for the range of interactions $\boldsymbol{\varphi}$ we are interested in, the roots r_i all lie on the unit circle in the complex plane, i.e., $|r_i| = 1$ for all i. Then we can write $Z(\lambda) = \prod_i (1 - \frac{\lambda}{r_i})$, and the log partition function becomes

$$f(\lambda) = \log Z = \sum_{i=1}^{n} \log\left(1 - \frac{\lambda}{r_i}\right) = \sum_{i=1}^{n} \sum_{j=1}^{\infty} \frac{1}{j} \left(\frac{\lambda}{r_i}\right)^j.$$

Denoting the first m terms by $f_m(\lambda) = \sum_{i=1}^{n} \sum_{j=1}^{m} \frac{1}{j} \left(\frac{\lambda}{r_i}\right)^j$, the error due to truncation is bounded by

$$|f(\lambda) - f_m(\lambda)| \leq n \sum_{j=m+1}^{\infty} \frac{|\lambda|^j}{j} \leq \frac{n \, |\lambda|^{m+1}}{(m+1)(1-|\lambda|)},$$

recalling that by symmetry we are assuming $|\lambda| < 1$. Thus to get within $\varepsilon/4$ additive error, it suffices to take $m \geq \frac{|\lambda|}{1-|\lambda|}\left(\log(4n/\varepsilon) + \log\frac{1}{1-|\lambda|}\right)$. The following result summarizes our discussion so far.

Lemma 2.1. *Given* $0 < \varepsilon < 1/4, m \geq \frac{|\lambda|}{1-|\lambda|}\left(\log(4n/\varepsilon) + \log\frac{1}{1-|\lambda|}\right)$, *and the values of* $\left\{\frac{\mathrm{d}^j}{\mathrm{d}\lambda^j}Z(\lambda)\Big|_{\lambda=0}\right\}_{j=0}^m$, $f_m(\lambda)$ *can be computed in time* $\mathrm{poly}(n/\varepsilon)$. *Moreover, if the Lee-Yang theorem holds for the partition function* $Z(\lambda)$, *then* $|f_m(\lambda) - f(\lambda)| < \varepsilon/4$, *and thus* $\exp(f_m(\lambda))$ *approximates* $Z(\lambda)$ *within a multiplicative factor* $1+\varepsilon$.

The missing ingredient in turning Lemma 2.1 into an FPTAS is the computation of the derivatives $\left\{\frac{\mathrm{d}^j}{\mathrm{d}\lambda^j}Z(\lambda)\Big|_{\lambda=0}\right\}_{j=0}^m$, which themselves are just multiples of the leading coefficients of Z. Computing these values naively using the definition of $Z(\lambda)$ requires $n^{\Omega(m)}$ time. Since m is required to be of order $\Omega(\log(n/\varepsilon))$, this results in only a quasi-polynomial time algorithm. In the next section, we show how to compute these values in polynomial time when H is a hypergraph of bounded degree and bounded hyperedge size, which when combined with Lemma 2.1 gives an FPTAS.

3. Computing coefficients via insects

As discussed in the introduction, Patel and Regts [35] recently introduced a technique for efficiently computing the leading coefficients of a partition function using induced subgraph counts. In this section we introduce the notion of *sub-insect counts*, and show how it allows the Patel-Regts framework to be adapted to any hypergraph two-spin system with vertex activities (including the Ising model with vertex activities as a special case). We will align our notation with [35] as much as possible. From now on, unless otherwise stated, we will use G to denote a hypergraph. Recall from the introduction the partition function of a two-spin system on a hypergraph $G = (V, E)$:

(5) $$Z_G^{\boldsymbol{\varphi}}(\lambda) = \sum_{\boldsymbol{\sigma}:V\to\{+,-\}} \prod_{e\in E} \varphi_e\big(\boldsymbol{\sigma}\big|_e\big)\, \lambda^{|\{v:\boldsymbol{\sigma}(v)=+\}|}.$$

Due to the normalization $\varphi_e(-,\cdots,-) = 1$, each term in the summation depends only on the set $S = \{v : \boldsymbol{\sigma}(v) = +\}$ and the labelled induced sub-hypergraph $\big(S\cup\partial S, E[S]\cup E(S,\bar{S})\big)$, where ∂S is the boundary of S defined as $\partial S := \bigcup_{v\in S} N_G(v)\setminus S$ and $N_G(v)$ is the set of vertices adjacent to the vertex v in G. This fact motivates the induced sub-structures we will consider.

Let $\boldsymbol{\sigma}^S$ be the configuration where the set of vertices assigned $+$-spins is S, that is, $\boldsymbol{\sigma}^S(v) = +$ for $v \in S$ and $\boldsymbol{\sigma}^S(v) = -$ otherwise. We will also write $\varphi_e(S) := \varphi_e(\boldsymbol{\sigma}^S\big|_e)$ for simplicity. Thus the partition function can be written

$$Z_G^{\boldsymbol{\varphi}}(\lambda) = \sum_{S\subseteq V}\prod_{e:e\cap S\neq\emptyset}\varphi_e(S)\lambda^{|S|}.$$

We start with the standard factorization of the partition function in terms of its complex zeros $r_1, \ldots, r_n$, where $n = |V|$. Since we are assuming $\varphi_e(-,\cdots,-) = 1$, the product of the zeros is 1. The partition function can then be written as

$$Z_G^{\boldsymbol{\varphi}}(\lambda) = \prod_i (1-\lambda/r_i) = \sum_{i=0}^n (-1)^i e_i(G)\lambda^i,$$

where $e_i(G)$ is the elementary symmetric polynomial evaluated at $(\frac{1}{r_1},\cdots,\frac{1}{r_n})$.

On the other hand, we can also express the coefficients $e_i(G)$ combinatorially using the definition of the partition function:

(6) $$e_i(G) = (-1)^i \sum_{\substack{S\subseteq V\\|S|=i}}\prod_{e:e\cap S\neq\emptyset}\varphi_e(S).$$

Once we have computed the first m coefficients of Z (i.e., the values $e_i(G)$ for $i = 1, \cdots, m$), where $m = \Omega\left(\frac{|\lambda|}{1-|\lambda|}\log(n/\varepsilon)\right)$, we can use Lemma 2.1 to obtain an FPTAS as claimed in Theorems 1.1, 1.3 and 1.4.

The main result in this section will be an algorithm for computing these coefficients $e_i(G)$:

Theorem 3.1. *Let $C > 0, \Delta \in \mathbb{N}$. Given an n-vertex hypergraph G of maximum degree Δ, maximum hyperedge size k, then for any $\varepsilon > 0$ there exists a deterministic $\mathrm{poly}(n/\varepsilon)$-time algorithm to compute $e_i(G)$ for $i = 1, \cdots, m$, where $m = \lceil C\log(n/\varepsilon)\rceil$.*

3.1. **Insects in a hypergraph.** To take advantage of the fact that each term in eq. (5) only depends on the set S and the induced sub-hypergraph $\left(S \cup \partial S, E[S] \cup E(S, \bar{S})\right)$, we define the following structure.

Definition 3.2. For disjoint sets S and B, an insect $H = (S, E, B)$ is a labelled hypergraph $(S \cup B, E)$ such that each vertex in B is incident on exactly one hyperedge in E. The set S is called the *label set* of the insect H and the set B the *boundary set*.

Given an insect H, we use the notation $\underline{V}(H)$ for its label set. The *size* $|H|$ of the insect H is defined to be $|\underline{V}(H)|$. H is said to be *connected* if the induced sub-hypergraph on the label set $\underline{V}(H)$ is connected.

Remark. Note that a hypergraph $G = (V, E)$ can itself be viewed as the insect $(V, E, \emptyset)$.

In order to exploit the structure of the terms in eq. (5) alluded to above, we now define the notion of an *induced sub-insect* of an insect. Given an insect $H = (S, E, B)$ and a subset S' of S, we define $\partial^+ S'$ as follows: starting with $\partial S'$, for each vertex $v \in \partial S'$, if v is incident on d hyperedges in $E(S', \bar{S}')$, we replace v with d copies of itself labelled $v^{(1)}, \cdots, v^{(d)}$, one for each of the hyperedges in lexicographic order, so that each $v^{(i)}$ is incident on only one hyperedge. Denote the set of these new hyperedges by E'. The induced sub-insect $H^+[S']$ is then defined as $(S', E[S'] \cup E', \partial^+ S')$. Note that if $S' = \emptyset$, then $H^+[S']$ is the empty graph.

3.1.1. *Weighted sub-insect counts.* Just as graph invariants may be expressed as sums over induced subgraph counts, we will consider weighted sub-insect counts of the form $f(G) = \sum_{S \subseteq V(G)} a_{G^+[S]}$ and the functions f expressible in this way.

Let $\mathcal{G}_t^{\Delta,k}$ be the set of insects up to size t, with maximum degree Δ and maximum hyperedge size k. Note that since insects are labelled, this is an infinite set. We will fix Δ and k throughout, and write $\mathcal{G} := \bigcup_{t \geq 1} \mathcal{G}_t^{\Delta,k}$. Let $w(H, G)$ be the indicator that H is an induced sub-insect of G, that is,

$$w(H, G) = 1 \text{ if and only if there is a set } S \subseteq \underline{V}(G) \text{ such that } G^+[S] = H.$$

A weighted sub-insect count $f(G)$ of the form considered above can then also be written as $f(G) = \sum_{H \in \mathcal{G}} a_H w(H, G)$. Note that even though $\mathcal{G}$ is infinite, the above sum has only finitely many non-zero terms for any finite insect G. It is also worth noting that, as insects are labelled, $f(G)$ may also depend on the labelling of G, unlike a graph invariant where isomorphic copies of a graph are the same.

3.2. **Properties of weighted sub-insect counts.** A weighted sub-insect count f is said to be *additive* if, given any two disjoint insects G_1 and G_2, $f(G_1 \uplus G_2) = f(G_1) + f(G_2)$. Analogously to the case of graph invariants, we then have the following:

Lemma 3.3. *Let f be a weighted sub-insect count, so that f may be written as*

$$f(G) := \sum_{S \subseteq V} a_{G^+[S]} = \sum_{H \in \mathcal{G}} a_H \cdot w(H, G).$$

Then f is additive if and only if $a_H = 0$ for all insects H that are disconnected.

Proof. When H is connected, we have $w(H, G_1 \uplus G_2) = w(H, G_1) + w(H, G_2)$; thus f given in the above form is additive if $a_{H'} = 0$ for all H' that are not connected.

Conversely, suppose f is additive. By the last paragraph, we can assume without loss of generality that the sequence a_H is supported on disconnected insects (by subtracting the component of f supported on connected H). We now show that for such an f, a_H must be 0 for all disconnected H as well.

For if not, then given an f with $a_{H'} = 0$ for all connected insects H', let H be an insect with the smallest size for which $a_H \neq 0$. If H has only one vertex, then it is vacuously connected and hence $a_H = 0$, so we already have a contradiction. Otherwise, we know that $a_J = 0$ for all insects J which have strictly fewer vertices than H. Since $w(H, J) = 0$ when the number of vertices in J is strictly smaller than in H, this also implies that $f(J) = 0$ for any graph J that is strictly smaller in size than H.

Since H is disconnected, there exist non-empty insects H_1 and H_2 such that $H = H_1 \uplus H_2$. By additivity, we then have $f(H) = f(H_1) + f(H_2) = 0$, where the last equality follows since both H_1 and H_2 are strictly smaller than H. On the other hand, since H is an insect with the smallest possible number of vertices such that $a_H \neq 0$, we also have $f(H) = a_H w(H, H) = a_H$. This implies $a_H = 0$, which is a contradiction. □

Our next lemma implies that products of weighted sub-insect counts can also be expressed as a weighted sub-insect count of slightly larger insects:

Lemma 3.4. *Let $H_1 = (S_1, E_1, B_1), H_2 = (S_2, E_2, B_2)$ be arbitrary insects. Either for every G we have*

$$w(H_1, G) w(H_2, G) = 0,$$

or there exist a set $S' \supseteq S_1 \cup S_2$, a set B' disjoint from S', and an insect $H = (S', E', B')$ such that $H^+[S_1] = H_1$, $H^+[S_2] = H_2$, and for every G we have

$$w(H_1, G) w(H_2, G) = w(H, G). \tag{7}$$

Before proving this lemma, we define a notion of *compatibility* for insects.

Definition 3.5. An insect $H_1 = (S_1, E_1, B_1)$ is *compatible* with another insect $H_2 = (S_2, E_2, B_2)$ if $H_1^+[S_1 \cap S_2] = H_2^+[S_1 \cap S_2]$.

Observation 3.6. *When H_1 and H_2 are compatible, there exists a unique insect H (which we denote by $H_1 \cup H_2$ by a slight abuse of notation) with label set $\underline{V}(H_1) \cup \underline{V}(H_2)$ such that $H^+[\underline{V}(H_1)] = H_1$ and $H^+[\underline{V}(H_2)] = H_2$.*

We now proceed with the proof of Lemma 3.4.

Proof of Lemma 3.4. If H_1 is not compatible with H_2, then for every G, either H_1 is an induced sub-insect of G, or H_2 is, but not both. Thus $w(H_1, G) w(H_2, G) = 0$.

If H_1 and H_2 are compatible, then let $H = H_1 \cup H_2$ be the insect promised by Observation 3.6. We have

$$\begin{aligned} w(H_1, G) w(H_2, G) &= \sum_{T_1 \subseteq V} \sum_{T_2 \subseteq V} \left[G^+[T_1] = H_1\right] \cdot \left[G^+[T_2] = H_2\right] \\ &= \sum_{T \subseteq V} \sum_{T_1 \subseteq V} \sum_{T_2 \subseteq V} \left[G^+[T_1] = H_1\right] \cdot \left[G^+[T_2] = H_2\right] \cdot [T = T_1 \cup T_2] \\ &= \sum_{T \subseteq V} \left[G^+[T] = H_1 \cup H_2\right], \text{ using the definition of compatibility} \\ &= w(H, G). \end{aligned}$$

□

3.3. **Enumerating connected sub-insects.** Our goal is to show that $e_i(G)$ defined in eq. (6) can be written as a weighted sub-insect count. Recall that we denote by $\mathcal{G}_t^{\Delta,k}$ the set of distinct induced sub-insects of G up to t vertices with maximum degree Δ and maximum hyperedge size k. We then have

$$e_i(G) = (-1)^i \sum_{\substack{S\subseteq V\\ |S|=i}} \prod_{e:e\cap S\neq\emptyset} \varphi_e(S) = \sum_{\substack{H\in\mathcal{G}_i^{\Delta,k}\\ H=(V_H,E_H,S_H)}} \lambda_{H,i}\cdot w(H,G), \tag{8}$$

where $\lambda_{H,i} = (-1)^i \prod_{e:e\cap S_H\neq\emptyset} \varphi_e(S_H)$ if $|H| = i$, and 0 otherwise. Note that $\lambda_{H,i}$ is easily computable in time $\mathrm{poly}(|H|)$; however, as discussed in the introduction, the number of $H\in\mathcal{G}_i^{\Delta,k}$ such that $w(H,G)\neq 0$ will be $\Omega(n^i)$, and a naive computation of $e_i(G)$ using eq. (8) would be inefficient. To prove Theorem 3.1, we consider the set of *connected* insects, denoted by $\mathcal{C}_i^{\Delta,k}$, instead of $\mathcal{G}_i^{\Delta,k}$. We will show in this subsection that $\mathcal{C}_i^{\Delta,k}$ can be efficiently enumerated, and then in the following subsection reduce the above summation over $\mathcal{G}_i^{\Delta,k}$ to enumerations of $\mathcal{C}_i^{\Delta,k}$.

Following Patel and Regts [35], we use the following result based on Borgs, Chayes, Kahn and Lovász [14, Lemma 2.1 (c)].

Lemma 3.7. *Let G be a multigraph with maximum degree Δ (counting edge multiplicity) and let $v\in V(G)$. Then the number of subtrees of G with t vertices containing the vertex v is at most $\frac{(e\Delta)^{t-1}}{2}$.*

Proof. Consider the infinite rooted Δ-ary tree T_Δ. The number of subtrees with t vertices starting from the root is $\frac{1}{t}\binom{t\Delta}{t-1} < \frac{(e\Delta)^{t-1}}{2}$. (See also [43, Theorem 5.3.10].)

Next note that every subtree of the multigraph G containing the vertex v, induces injectively a subtree of T_Δ containing the root. Thus the number of subtrees of G containing vertex v is also at most $\frac{(e\Delta)^{t-1}}{2}$. ☐

Corollary 3.8. *Let G be a hypergraph with maximum degree Δ and maximum hyperedge size k, and let $v\in V(G)$. Then the number of connected induced sub-insects of G with t vertices containing the vertex v is at most $\frac{(e\Delta k)^{t-1}}{2}$.*

Proof. Consider the multigraph H obtained by replacing every hyperedge of size r in G by an r-clique. For any two distinct connected induced sub-insects A and B, let S_A and S_B be their sets of spanning trees. Then since the label sets of A and B are different, $S_A\cap S_B=\emptyset$. Thus the number of connected subtrees is an upperbound on the number of connected induced sub-insects.

In this multigraph the maximum degree is Δk, so by Lemma 3.7 the number of subtrees is at most $\frac{(e\Delta k)^{t-1}}{2}$. Thus the number of connected induced sub-insects of G with t vertices containing the vertex v is also at most $\frac{(e\Delta k)^{t-1}}{2}$. ☐

As a consequence we can efficiently enumerate all connected induced sub-insects of logarithmic size in a bounded degree graph. This follows from a similar reduction to a multigraph, and then applying [35, Lemma 3.4]. However, for the sake of completeness we also include a direct proof.

Lemma 3.9. *For a hypergraph of maximum degree Δ and maximum hyperedge size k, there exists an efficient enumerator for connected induced sub-insects of size t, that runs in time $\widetilde{O}(nt^3(e\Delta k)^{t+2})$. Here $\widetilde{O}$ hides factors of the form* $\mathrm{polylog}(n)$, $\mathrm{polylog}(\Delta k)$ *and* $\mathrm{polylog}(t)$.

Proof. Let $\mathcal{T}_t$ be the set of $S\subseteq V(G)$ such that $|S|\le t$ and $G[S]$ is connected. Note that given $S\in\mathcal{T}_t$, $G^+[S]$ will be a sub-insect of size t, and this clearly enumerates all of them. Also by Corollary 3.8, $|\mathcal{T}_t|\le O(n(e\Delta k)^t)$. Thus it remains to give an algorithm to construct $\mathcal{T}_t$ in about the same amount of time.

We construct $\mathcal{T}_t$ inductively. For $t=1$, $\mathcal{T}_1 := V(G)$. Then given $\mathcal{T}_{t-1}$, let

$$\mathcal{T}_t^* := \mathcal{T}_{t-1}\cup\{S\cup\{v\} : S\in\mathcal{T}_{t-1} \text{ and } v\in N_G(S)\setminus S\}.$$

Since $|N_G(S)| < t\Delta k$, thus one can compute the set $N_G(S) \setminus S$ in time $O(t\Delta k)$, and construct $\mathcal{T}_t^*$ in time $O(|\mathcal{T}_{t-1}|\, t\Delta k) = O(nt(e\Delta k)^{t+1})$. Finally we remove duplicates in $\mathcal{T}_t^*$ to get $\mathcal{T}_t$ (e.g., by sorting the sets $S \in \mathcal{T}_t^*$, where each can be represented by a string of length $\widetilde{O}(t)$), in time $\widetilde{O}(nt^3(e\Delta k)^{t+1})$.

Starting from $\mathcal{T}_1$, inductively we perform t iterations to construct $\mathcal{T}_t$. Thus the overall running time is $\sum_{i=1}^{t} \widetilde{O}(ni^3(e\Delta k)^{i+1}) = \widetilde{O}(nt^3(e\Delta k)^{t+2})$. □

3.4. **Proof of Theorem 3.1.** The previous subsection allows us to efficiently enumerate connected sub-insects. To prove Theorem 3.1, it remains to reduce the sum over all (possibly disconnected) H in eq. (8) to a sum over *connected* H. Consider the t-th power sum:

$$p_t = \sum_{i=1}^{n} \frac{1}{r_i^t}.$$

Now by Newton's identities (which relate power sums to elementary symmetric polynomials), we have

$$p_t = \sum_{i=1}^{t-1} (-1)^{i-1} p_{t-i} e_i + (-1)^{t-1} t e_t. \tag{9}$$

Recall from eq. (8) that e_i is a weighted sub-insect count, and also from Lemma 3.4 that the product of weighted sub-insect counts is also a weighted sub-insect count. Therefore, p_k is also a weighted sub-insect count:

$$p_t(G) = \sum_{H \in \mathcal{G}_t^{\Delta,k}} a_{H,t} w(H,G), \tag{10}$$

for some $a_{H,t}$ to be determined.

Next note that $Z_{G_1 \uplus G_2}(\lambda) = Z_{G_1}(\lambda) \cdot Z_{G_2}(\lambda)$ is multiplicative, and hence the t-th power sum p_t is additive. Hence by Lemma 3.3, its coefficients are supported on connected insects:

$$p_t(G) = \sum_{H \in \mathcal{C}_t^{\Delta,k}} a_{H,t} w(H,G), \tag{11}$$

where $\mathcal{C}_t^{\Delta,k}$ is the set of distinct connected sub-insects with up to t vertices and maximum degree Δ, maximum hyperedge size k. Notice that by Corollary 3.8, there are at most $n(e\Delta k)^t$ non-zero terms (where $w(H,G) \neq 0$).

Now we show how to compute the coefficients $a_{H,t}$ efficiently, following the same approach as in [35, Lemma 3.6].

Lemma 3.10. *There is a* $\mathrm{poly}(n/\varepsilon)$*-time algorithm to compute all the coefficients* $a_{H,t}$ *in eq. (11), for* $t \leq O(\log(n/\varepsilon))$*.*

Proof. By Lemma 3.9, we compute $\mathcal{T}_t$, consisting of all $S \subseteq V(G)$ such that $|S| \leq t$ and $G[S]$ is connected. As we have removed duplicates, this is exactly $\mathcal{C}_t^{\Delta,k}$.

By eq. (9), for $t = 1$ we have $p_1 = e_1$, so by eq. (8) we can read off the coefficients $a_{H,1}$ from $e_1(G)$. Next suppose we have computed $a_{H',t'}$ for $|H'| \leq t' < t$, and we want to compute $a_{H,t}$ for fixed H. Again by Newton's identities, it suffices to compute the coefficients of $w(H,G)$ in $p_{t-i}e_i$. By eqs. (7), (8) and (11), the coefficient of $w(H,G)$ in $p_{t-i}e_i$ is given by:

$$\sum_{\substack{H_1 \in \mathcal{G}_i, H_2 \in \mathcal{C}_{(t-i)} \\ H_1 \text{ compatible with } H_2 \\ H_1 \cup H_2 = H}} a_{H_2,(t-i)} \lambda_{H_1,i} = \sum_{\substack{(S_1,S_2) \\ S_1 \cup S_2 = V(H)}} a_{H[S_2],(t-i)} \lambda_{H[S_1],i}. \tag{12}$$

Since $t \leq O(\log(n/\varepsilon))$, the second sum involves at most $4^t = \mathrm{poly}(n/\varepsilon)$ terms. Moreover, due to Corollary 3.8, there are at most $tn(e\Delta k)^t = \mathrm{poly}(n/\varepsilon)$ previously computed $a_{H',t'}$, where H' is a connected sub-insect of G and $|H'| \le t' < t$. And to look up $a_{H[T],(t-i)}$, one can do a linear scan, which also takes time $\mathrm{poly}(n/\varepsilon)$ for $t \le O(\log(n/\varepsilon))$.

To conclude, because $t \le O(\log(n/\varepsilon))$, eq. (11) only contains $\mathrm{poly}(n/\varepsilon)$ terms. And for each term, $a_{H,t}$ can be computed using the above dynamic programming scheme in $\mathrm{poly}(n/\varepsilon)$ time. ☐

Finally, now that we can compute $a_{H,t}$ efficiently, by eq. (11) we can compute p_k using the sub-insect enumerator in Lemma 3.9, and we can then compute e_k using Newton's identities as in eq. (9), which completes the proof of Theorem 3.1.

3.5. **Proofs of main theorems.** Our first main result in the introduction, the FPTAS for the Ising model on graphs throughout the ferromagnetic regime with non-zero field stated in Theorem 1.1, now follows by combining Theorem 3.1 with Lemma 2.1 and the Lee-Yang theorem [24] (also stated as Theorem 4.2 in the next section). Recall from the introduction that the Lee-Yang theorem ensures that the partition function has no zeros inside the unit disk.

Similarly, Theorem 1.4, the FPTAS for two-spin systems on hypergraphs, follows by combining Theorem 3.1 with Lemma 2.1 and the Suzuki-Fisher theorem [44] (also stated as Theorem 4.3 in the next section). Again, the Suzuki-Fisher theorem ensures that there are no zeros inside the unit disk, under the condition on the hyperedge activities stated in Theorem 1.4.

To establish our final main algorithmic result, Theorem 1.3, we first need to prove a new Lee-Yang theorem for the hypergraph Ising model as stated in Theorem 1.2 in the introduction. This will be the content of the next and final section of our paper. Once we have that, Theorem 1.3 follows immediately by the same route as above.

4. A Lee-Yang Theorem for Hypergraphs

In this section we prove a tight Lee-Yang theorem for the hypergraph Ising model (Theorem 1.2 in the introduction). We start by extending the definition of the hypergraph Ising model to the multivariate setting, where each vertex and each hyperedge is allowed to have a different activity. As before, we have an underlying hypergraph $G = (V, E)$ with $|V| = n$ vertices. Given vertex activities $\lambda_1, \lambda_2, \ldots, \lambda_n$ and hyperedge activities $\boldsymbol{\beta} = (\beta_e)$, we define

$$Z_G^{\boldsymbol{\beta}}(\lambda_1, \cdots, \lambda_n) = \sum_{S \subseteq V} \prod_{e \in E(S, \overline{S})} \beta_e \prod_{i \in S} \lambda_i \,,$$

where for a subset $S \subseteq V$, $E(S, \overline{S})$ is the set of hyperedges with at least one vertex in each of S and $\overline{S}$. Note that

$$Z_G^{\boldsymbol{\beta}}(\lambda_1, \cdots, \lambda_n) = \prod_{i=1}^{n} \lambda_i \cdot Z_G^{\boldsymbol{\beta}}\left(\frac{1}{\lambda_1}, \cdots, \frac{1}{\lambda_n}\right). \tag{13}$$

We recall the definition of the Lee-Yang property (see, e.g., [37]).

Definition 4.1 (Lee-Yang property). Let $P(z_1, z_2, \ldots, z_n)$ be a multilinear polynomial with real coefficients. P is said to have the *Lee-Yang property* (sometimes written as "P is LY") if for any complex numbers $\lambda_1, \cdots, \lambda_n$ such that $|\lambda_1| \ge 1, \cdots, |\lambda_n| \ge 1$, and $|\lambda_i| > 1$ for some i, it holds that $P(\lambda_1, \cdots, \lambda_n) \ne 0$.

Then the seminal Lee-Yang theorem [24] can be stated as follows:

Theorem 4.2. *Let G be a connected undirected graph, and suppose $0 < \beta < 1$. Then the Ising partition function $Z_G^{\beta}(\lambda_1, \cdots, \lambda_n)$ has the Lee-Yang property.*

The following extension of the Lee-Yang theorem to general two-spin systems on hypergraphs is due to Suzuki and Fisher [44]. Again the theorem is stated in the multivariate setting, where in the two-spin partition function in eq. (5) each vertex i has a distinct activity λ_i.

Theorem 4.3. *Consider any symmetric hypergraph two-spin system, with a connected hypergraph G and edge interactions $\{\varphi_e\}$. Then the partition function $Z_G^{\boldsymbol{\varphi}}(\lambda_1, \cdots, \lambda_n)$ has the Lee-Yang property if $|\varphi_e(+, \cdots, +)| \geq \frac{1}{4}\sum_{\boldsymbol{\sigma}\in\{+,-\}^V} |\varphi_e(\boldsymbol{\sigma})|$ for every hyperedge e.*

Theorem 4.3 is not tight for the important special case of the Ising model on hypergraphs. Our goal in this section is to prove a tight analog of the original Lee-Yang theorem for this case. Specifically, we will prove the following:

Theorem 4.4. *Let $G = (V, E)$ be a connected hypergraph, and $\boldsymbol{\beta} = (\beta_e)_{e\in E}$ be the vector of real valued hyperedge activities so that the activity of edge $e \in E$ is β_e. Then $Z_G^{\boldsymbol{\beta}}$ has the Lee-Yang property if and only if the following condition holds for every hyperedge e: let $k > 1$ be the size of e,*

- *if $k = 2$, then $-1 < \beta_e < 1$;*
- *if $k \geq 3$, then $-\frac{1}{2^{k-1}-1} < \beta_e < \frac{1}{2^{k-1}\cos^{k-1}(\frac{\pi}{k-1})+1}$.*

Note that the case $k = 2$ is just the original Lee-Yang theorem (Theorem 4.2).

The following corollary for the univariate polynomial $Z_G^{\boldsymbol{\beta}}(\lambda)$ follows immediately via eq. (13) and the fact that, by Hurwitz's theorem, the zeros of $Z_G^{\boldsymbol{\beta}}(\lambda)$ are continuous functions of $\boldsymbol{\beta}$ and thus remain on the unit circle after taking the limit in the range of each β_e.

Corollary 4.5. *Let $G = (V, E)$ be a connected hypergraph, and $\boldsymbol{\beta} = (\beta_e)_{e\in E}$ be the vector of real valued hyperedge activities so that the activity of edge $e \in E$ is β_e. Then, all complex zeros of the univariate partition function $Z_G^{\boldsymbol{\beta}}(\lambda)$ lie on the unit circle if and only if the following condition holds for every hyperedge e: let $k > 1$ be the size of e,*

- *if $k = 2$, then $-1 \leq \beta_e \leq 1$;*
- *if $k \geq 3$, then $-\frac{1}{2^{k-1}-1} \leq \beta_e \leq \frac{1}{2^{k-1}\cos^{k-1}(\frac{\pi}{k-1})+1}$.*

This establishes Theorem 1.2 in the introduction, and hence also Theorem 1.3 as explained at the end of the previous section.

Remark. As a comparison, the Suzuki and Fisher result, which we restated in Theorem 4.3, implies that a sufficient condition for the Lee-Yang property of $Z_G^{\boldsymbol{\beta}}(\lambda)$ is

$$-\frac{1}{2^{k-1}-1} \leq \beta_e \leq \frac{1}{2^{k-1}-1}.$$

Note that while the lower bound on β_e is the same as ours, our (tight) upper bound is always better, and significantly so for the more interesting case of small k. E.g., for $k = 3$, our result gives the optimal range $-\frac{1}{3} \leq \beta_e \leq 1$, while the Suzuki-Fisher theorem gives $-\frac{1}{3} \leq \beta_e \leq \frac{1}{3}$. Similarly, for $k = 4$, the respective ranges are $[-1/7, 1/2]$ (for ours) and $[-1/7, 1/7]$ (for Suzuki-Fisher).

We turn now to the proof of Theorem 4.4. We will make use of the following criterion for the Lee-Yang property.

Lemma 4.6. *Given a multilinear polynomial $P(z_1, z_2, \ldots, z_n)$ with real coefficients define, for each $1 \leq j \leq n$, multilinear polynomials A_j and B_j in the variables $z_1, \ldots, z_{j-1}, z_{j+1}, \ldots, z_n$ such that*

$$P = A_j z_j + B_j$$

If P has the Lee-Yang property then, for every j such that the variable z_j has positive degree in P, it holds that $A_j(z_1, \ldots, z_{j-1}, z_{j+1}, \ldots, z_n) \neq 0$ when $|z_i| \geq 1$ for all $i \neq j$. In particular, A_j itself is LY.

Proof. Without loss of generality, we assume that $j = 1$. Note that since z_1 has positive degree in P, A_1 is a non-zero polynomial. Suppose that, in contradiction of the claim of the lemma, there exist complex numbers $\lambda_2, \ldots, \lambda_n$ satisfying $|\lambda_i| \geq 1$ such that $A_1(\lambda_2, \ldots, \lambda_n) = 0$. Since P is LY, it follows that

$B_1(\lambda_2, \ldots, \lambda_n) \neq 0$ (for otherwise, we get a contradiction to the Lee-Yang property by choosing z_1 to be an arbitrary value outside the closed unit disk).

By continuity, this implies that $|B_1|$ is positive in any small enough neighborhood of $(\lambda_2, \ldots, \lambda_n)$ in $\mathbb{C}^{n-1}$. In particular, let S_ϵ be the open set

$$S_\epsilon := \{(y_2, \ldots, y_n) \mid \ |y_i - \lambda_i| < \epsilon \text{ and } |y_i| > 1 \text{ for } 2 \leq i \leq n\ \}.$$

Then there exist positive δ_0 and ϵ_0 such that $|B_1|$ is at least δ_0 in the open set S_ϵ when $\epsilon < \epsilon_0$.

Now, since A_1 is a non-zero multilinear polynomial, it cannot vanish identically on any open set. In particular, it cannot vanish identically in S_ϵ for any $\epsilon > 0$. On the other hand, since A_1 vanishes at $(\lambda_2, \ldots, \lambda_n)$ it follows from continuity that for $\epsilon < \epsilon_0$ small enough, $|A_1| \leq \delta_0/2$ in S_ϵ. Since A_1 does not vanish identically on S_ϵ, there must exist a point $(y_2, \ldots, y_n)$ in S_ϵ such that $0 < |A_1(y_1, y_2, \ldots, y_n)| < \delta_0/2$. Since $|B_1(y_2, \ldots, y_n)| \geq \delta_0$ by the choice of ϵ_0, it follows that if we define $y_1 = -B_1(y_2, \ldots, y_n)/A_1(y_2, \ldots, y_n)$ then $2 < |y_1| < \infty$. However, we then have $P(y_1, y_2, \ldots, y_n) = 0$ even though $|y_1| > 1$ and $|y_i| \geq 1$ for all i. This contradicts the Lee-Yang property of P. ☐

By iterating the above lemma, we get the following corollary.

Corollary 4.7. *Let $P(z_1, z_2, \ldots, z_n)$ be a multilinear polynomial with non-zero coefficients, i.e.,*

$$P(z_1, \ldots, z_n) = \sum_{S \subseteq [n]} p_S \prod_{i \in S} z_i,$$

where $p_S \neq 0$ for all $\subseteq [n]$. Then, for every subset S of $[n]$, the polynomial A_S defined by the equation

$$P(z_1, \ldots, z_n) = A_S((z_i)_{i \notin S}) \prod_{i \in S} z_i + \sum_{T \supsetneq S} p_T \prod_{i \in T} z_i$$

has the property that $A_S((z_i)_{i \notin S}) \neq 0$ when $|z_i| \geq 1$ for all $i \notin S$. In particular, A is LY.

We next show that Lemma 4.6 has a partial converse for symmetric multilinear functions.

Lemma 4.8. *Let $P(z_1, z_2, \ldots, z_n)$ be a symmetric multilinear polynomial with non-zero real coefficients, i.e.,*

$$P(z_1, \ldots, z_n) = \sum_{S \subseteq [n]} p_S \prod_{i \in S} z_i,$$

where $p_S \neq 0$ for all $\subseteq [n]$ and $p_S = p_{\overline{S}}$. Assume further that the polynomials A_j as defined in Lemma 4.6 all have the property that they are non-zero when all their arguments z_i satisfy $|z_i| \geq 1$. Then P is LY.

Proof. We first show that, under our assumptions, if all but one of the z_j lie on the unit circle, then P can only vanish if the remaining z_j is also on the unit circle. Without loss of generality we set $j = 1$, that is, we will show that if $|z_i| = 1$ for $i \geq 2$, then any root $z_1 = \zeta_1$ in the equation $A_1 z_1 + B_1 = 0$ satisfies $|\zeta_1| = 1$. (Here A_1 and B_1 are in the notation of Lemma 4.6.)

Since $A_1 = \sum_{S \subseteq [2,n]} p_{S \cup \{1\}} \prod_{i \in S} z_i$ does not vanish with this setting of the z_i, we have

$$|\zeta_1| = \left|\frac{B_1}{A_1}\right| = \left|\frac{\sum_{S \subseteq [2,n]} p_S \prod_{i \in S} z_i}{\sum_{S \subseteq [2,n]} p_{S \cup [1]} \prod_{i \in S} z_i}\right| = \left|\left(\prod_{i \in [2,n]} z_i\right) \frac{\sum_{S \subseteq [2,n]} p_S \prod_{i \notin S, i \neq 1} (1/z_i)}{\sum_{S \subseteq [2,n]} p_{S \cup [1]} \prod_{i \in S} z_i}\right|$$

$$\stackrel{(\star)}{=} \left|\frac{\sum_{S \subseteq [2,n]} p_S \prod_{i \notin S, i \neq 1} \overline{z_i}}{\sum_{S \subseteq [2,n]} p_{S \cup [1]} \prod_{i \in S} z_i}\right| \stackrel{(\dagger)}{=} \left|\frac{\sum_{S \subseteq [2,n]} p_{S \cup [1]} \prod_{i \in S} \overline{z_i}}{\sum_{S \subseteq [2,n]} p_{S \cup [1]} \prod_{i \in S} z_i}\right| = 1. \tag{14}$$

Here $(\star)$ uses the fact that $|z_i| = 1$ for $i \geq 2$ and $(\dagger)$ uses the symmetry of P. We have thus shown that if $(z_1, z_2, \ldots, z_n)$ is a zero of P such that $|z_i| \geq 1$ for all i then it is impossible for only one z_i to lie outside the closed unit disk.

We now show that if there are $k \geq 2$ values of i for which z_i lies outside the closed unit disk, then we can find another zero $(\zeta_1, \zeta_2, \zeta_3, \ldots, \zeta_n)$ of P such that $|\zeta_i| \geq 1$ for all i, and exactly $k-1$ of the ζ_i lie outside the closed open disk. We can then iterate this process to reduce k to 1, in which case the observation in the previous paragraph leads to a contradiction.

By re-numbering the indices if needed, we can assume that $|z_1|, |z_2| > 1$ and $|z_i| \geq 1$ for $i \geq 3$. We can then write

$$P(z_1, \ldots, z_n) = \alpha_{12} z_1 z_2 + \alpha_1 z_1 + \alpha_2 z_2 + \alpha_\emptyset,$$

where $\alpha_{12}, \alpha_1, \alpha_2$ and $\alpha_\emptyset$ are non-zero polynomials in $z_3, \ldots, z_n$. Further, the hypotheses of the lemma imply that $A_1 = \alpha_{12} z_2 + \alpha_1$ and $A_2 = \alpha_{12} z_1 + \alpha_2$ both have the Lee-Yang property. Thus, by Lemma 4.6, $\alpha_{12}(z_3, \ldots, z_n) \neq 0$, whence $|z_i| \geq 1$ for $i \geq 3$. Now, $A_2 \neq 0$ when $|z_1|$ and $|z_3|, \ldots, |z_n|$ are at least 1, while $z_1 = \frac{\alpha_2(z_3, \ldots, z_n)}{\alpha_{12}(z_3, \ldots, z_n)}$ gives $A_2 = 0$. Thus, we must have that

$$\frac{|\alpha_2(z_3, \ldots, z_n)|}{|\alpha_{12}(z_3, \ldots, z_n)|} < 1 \text{ when } |z_i| \geq 1 \text{ for } i \geq 3. \tag{15}$$

We now set $\zeta_i = z_i$ for $i \geq 3$, and consider z_1 as a function of z_2. The equality $P(z_1, z_2, \zeta_3, \ldots, \zeta_n) = 0$ is now equivalent to

$$z_1 = -\frac{\alpha_2 z_2 + \alpha_\emptyset}{\alpha_{12} z_2 + \alpha_1}. \tag{16}$$

Further, the hypotheses of the lemma imply that the denominator (which is $A_1(z_2, \zeta_3, \ldots, \zeta_n)$) is non-zero when $|z_2| \geq 1$. We thus see that

$$\lim_{z_2 \to \infty} |z_1| = \frac{|\alpha_2|}{|\alpha_{12}|} < 1. \tag{17}$$

Initially, both z_1 and z_2 lie outside the closed unit disk. Thus, by eq. (17) and continuity, we can take z_2 large enough in absolute value such that z_1 as defined in eq. (16) lies on the unit circle. We now choose ζ_1 and ζ_2 to be these values of z_1 and z_2 respectively, so that we have $P(\zeta_1, \ldots, \zeta_n) = 0$ and the number of the ζ_i lying on the unit circle is exactly one less than the number of the z_i lying on the unit circle, as required. □

The first step in our proof of Theorem 4.4 is to derive conditions under which the Ising partition function of a hypergraph consisting of a *single hyperedge* has the Lee-Yang property; it will turn out that this is actually the determining case for the full theorem. We will require the following technical lemma.

Lemma 4.9. *Let m be any integer, and k a positive integer such that $2|m| \leq k$. Consider the maximization problem*

$$\begin{aligned} &\max \prod_{i=1}^{k} \cos\theta_i \\ \textit{subject to} \quad & -\frac{\pi}{2} \leq \theta_i \leq \frac{\pi}{2}, \\ & \sum_{i=1}^{k} \theta_i = m\pi. \end{aligned}$$

The maximum is $\cos^k\left(\frac{m\pi}{k}\right)$, and is attained when $\theta_i = \frac{\pi}{k}$ for all i.

Proof. We can assume without loss of generality that $\theta_i \in (-\pi/2, \pi/2)$ at any maximum (for otherwise the objective value is 0). Now, consider the function $f(x) = \log\cos x$ defined on the interval $(-\pi/2, \pi/2)$.

Since $f'(x) = -\tan x$ is a decreasing function, $f(x)$ is concave for $x \in (-\frac{\pi}{2}, \frac{\pi}{2})$. Thus by Jensen's inequality,

$$\log \prod_{i=1}^{k} \cos\theta_i = \sum_{i=1}^{k} f(\theta_i) \leq k f\left(\frac{\sum_{i=1}^{k}\theta_i}{k}\right) \leq k \log\cos\left(\frac{m\pi}{k}\right),$$

and equality holds when $\theta_i = \frac{m\pi}{k}$ for all i. Note that these θ_i are in $(-\pi/2, \pi/2)$ since $2\,|m| \leq k$. ☐

We are now ready to tackle the case of a single hyperedge.

Lemma 4.10. *Fix an integer $k \geq 2$ and a hyperedge activity $\beta \in \mathbb{R}$. Let $G = (V = \{v_1, v_2, \ldots, v_k\}, E = \{\{v_1, v_2, \ldots, v_k\}\})$ be a hypergraph consisting of a single hyperedge of size k and activity β. If $k = 2$ and $\beta \in (-1, 1)$, or $k \geq 3$ and β satisfies*

$$-\frac{1}{2^{k-1}-1} < \beta < \frac{1}{2^{k-1}\cos^{k-1}\left(\frac{\pi}{k-1}\right)+1},$$

then the partition function Z_G^{β} has the Lee-Yang property.

Remark. Note that the condition on β imposed above is monotone in k: i.e., if β is such that the partition function of a hyperedge of size $k \geq 2$ is LY, then for the same β the partition function of a hyperedge of size $k' < k$ is also LY.

Proof. For $k = 2$, the lemma is a special case of the Lee-Yang theorem [24] (although it also follows by specializing the argument below). We therefore assume $k \geq 3$.

Since the Ising partition function is symmetric and all terms in the polynomial appear with positive coefficients, Lemma 4.8 applies and it suffices to verify that the coefficients A_j do not vanish when $|z_i| \geq 1$ for $i \neq j$. Without loss of generality we fix $j = 1$. We then have,

$$A_1 = \beta \prod_{\substack{i=1\\ i\neq j}}^{k}(1+z_i) + (1-\beta)\prod_{\substack{i=1\\ i\neq j}}^{k} z_i.$$

Thus $A_1 = 0$ for $|z_i| \geq 1$ is equivalent to

$$\frac{1}{\beta} = 1 - \prod_{\substack{i=1\\ i\neq j}}^{k}\left(1+\frac{1}{z_i}\right). \tag{18}$$

To establish the lemma, we therefore only need to show that for the claimed values of β, eq. (18) has no solutions when $|z_i| \geq 1$ for all $i \geq 2$. We now proceed to establish this by analyzing the product on the right hand side of eq. (18).

The map $z \mapsto 1 + 1/z$ is a bijection from the complement of the open unit disk to the closed disk D of radius 1 centered at 1. Any $y \in D$ can be written as $y = r\exp(\iota\theta)$ for $\theta \in [-\pi/2, \pi/2]$ and $0 \leq r \leq 2\cos\theta$. Consider now the set $\mathbb{R} \cap \left\{\prod_{i=1}^{k-1} y_i \mid y_i \in D \text{ for all } i\right\}$ for $k \geq 3$. We show that this set is exactly the interval $[-\tau_0, \tau_1]$ where $\tau_0 = 2^{k-1}\cos^{k-1}(\pi/(k-1))$ and $\tau_1 = 2^{k-1}$. The claim of the lemma then follows since for the given values of β, $1 - 1/\beta$ lies outside $[-\tau_0, \tau_1]$ and hence eq. (18) cannot hold.

Recalling that each $y \in D$ can be written in the form $r\exp(\iota\theta)$ where $\theta \in [-\pi/2, \pi/2]$ and $0 \leq r \leq 2\cos\theta$, we find that the values τ_0 and τ_1 are defined by the following optimization problems (note that since $k \geq 3$, both programs are feasible):

$$\tau_0 = 2^{k-1} \max \prod_{i=1}^{k-1} \cos\theta_i \qquad\qquad \tau_1 = 2^{k-1} \max \prod_{i=1}^{k-1} \cos\theta_i$$

$$\text{subject to} \quad -\frac{\pi}{2} \le \theta_i \le \frac{\pi}{2}, \qquad\qquad \text{subject to} \quad -\frac{\pi}{2} \le \theta_i \le \frac{\pi}{2},$$

$$\sum_{i=1}^{k-1} \theta_i = \pi + 2n\pi, \qquad\qquad \sum_{i=1}^{k-1} \theta_i = 2n\pi.$$

$$\text{for some } n \in \mathbb{Z} \qquad\qquad \text{for some } n \in \mathbb{Z}$$

$$\text{such that } |2n+1| \le (k-1)/2. \qquad\qquad \text{such that } |n| \le (k-1)/4.$$

Using Lemma 4.9, we then see that $\tau_0 = 2^{k-1}\cos^{k-1}(\pi/(k-1))$ and $\tau_1 = 2^{k-1}$, as required. □

We now proceed to an inductive proof of Theorem 4.4, using Lemma 4.10 as the base case.

Proof of Theorem 4.4. The case $k = 2$ is a special case of the Lee-Yang theorem [24] (though, as with the proof of Lemma 4.10, the argument below can again be specialized to directly establish this). We assume therefore that $k \ge 3$.

The proof uses the inductive method of Asano [3]. When the hypergraph consists of a single hyperedge of size $k' \le k$, it follows from Lemma 4.10 and the remark following it that the partition function is LY for the claimed values of edge activity β. For the induction, we use the fact that the Lee-Yang property of the partition function is preserved under the following two operations:

(1) **Adding a hyperedge**: In this operation, a new hyperedge of size $k' \le k$ and activity β as claimed in the statement of the theorem, is added to a connected hypergraph in such a way that exactly one of its k' vertices already exists in the starting hypergraph, while the other $k' - 1$ vertices are new. Note that this operation keeps the hypergraph connected. We assume that the partition functions of both the original hypergraph as well as the newly added edge separately have the Lee-Yang property: this follows from the induction hypothesis (for the hypergraph) and Lemma 4.10 (for the new hyperedge).

(2) **Asano contraction**: In this operation, two vertices u', u'' in a connected hypergraph that are not both included in any one hyperedge are merged so that the new merged vertex u is incident on all the hyperedges incident on u' or u'' in the original graph. Note that this operation keeps the hypergraph connected and does not change the size of any of the hyperedges.

Any connected non-empty hypergraph G can be constructed by starting with any arbitrary hyperedge present in G and performing a finite sequence of the above two operations: to add a new hyperedge e with activity β_e as present in G, one first uses operation 1 to add a hyperedge which has the same activity β_e and has new copies of all but one of the incident vertices of e, and then uses operation 2 to merge these new copies with their counterparts, if any, in the starting hypergraph. Note that in this process, a hyperedge e can be added only when at least one of its vertices is already included in the current hypergraph. However, since G is assumed to be connected, its hyperedges can be ordered so that all of them are added by the above process. Thus, assuming that the above two operations preserve the Lee-Yang property, it follows by induction on the number of hyperedges that the partition functions of all connected hypergraphs of hyperedge size at most k, and edge activities β_e as claimed in the theorem, have the Lee-Yang property.

The proof of the fact, first proved by Asano [3], that these two operations preserve the Lee-Yang property are by now standard, and can be found, e.g., in [44, Propositions 1, 2]. We include a proof here for completeness.

Consider first operation 1. Let G be the original hypergraph and H the new hyperedge (with $k' \le k$ vertices) being added, and assume, by renumbering vertices if required, that the vertices being merged are v_1 in G and u_1 in H respectively. Let $P(z_1, z_2, \ldots, z_n) = A(z_2, \ldots, z_n)z_1 + B(z_2, \ldots, z_n)$ and

$Q(y_1, y_2, \ldots, y_{k'}) = C(y_2, \ldots, y_{k'})y_1 + D(y_2, \ldots, y_{k'})$ be the Ising partition functions of G and H, respectively, where z_1 and y_1 are the variables corresponding to v_1 and u_1 respectively. Both P and Q are LY by the hypothesis of the operation. The partition function R of the new graph can be written as

$$R(z, z_2, \ldots, z_n, y_2, \ldots, y_{k'}) = A(z_2, \ldots, z_n)C(y_2, \ldots, y_{k'})z + B(z_2, \ldots, z_n)D(y_2, \ldots, y_{k'}),$$

where z is a new variable corresponding to the new vertex created by the merger of u and v. Let $\lambda_2, \ldots, \lambda_n, \mu_2, \ldots, \mu_{k'}$ be complex numbers lying outside the open unit disk. In order to prove that R is LY, we need to show that (i) $R(z, \lambda_2, \ldots, \lambda_n, \mu_2, \ldots, \mu_{k'}) = 0$ implies that $|z| \leq 1$; and (ii) when at least one of these complex numbers lies strictly outside the *closed* unit disk then $R(z, \lambda_2, \ldots, \lambda_n, \mu_2, \ldots, \mu_{k'}) = 0$ implies that $|z| < 1$. Now, since P and Q are assumed to be LY, Lemma 4.6 implies that $A = A(\lambda_2, \ldots, \lambda_n)$ and $C = C(\mu_2, \ldots, \mu_{k'})$ are both non-zero. Thus, $R = 0$ implies that

$$|z| = |B/A| \cdot |D/C| \,, \tag{19}$$

where $B = B(\lambda_2, \ldots \lambda_n)$ and $D = D(\mu_2, \ldots, \mu_{k'})$. Since all the λ_i and μ_i lie outside the open unit disk and P and Q are LY, $|B/A|\,, |D/C| \leq 1$, so that from eq. (19) $|z| \leq 1$. Further, when at least one of the λ_i lies strictly outside the closed unit disk, then again, since P is LY, $|B/A| < 1$. Similarly,$|D/C| < 1$ when one of the μ_i lies outside the closed unit disk. Thus, when at least one of the λ_i and the μ_i lies outside the closed unit disk, it follows from eq. (19) that $|z| < 1$. Thus, the conditions (i) and (ii) needed to establish that R is LY are both satisfied.

We now consider operation 2. By renumbering vertices if necessary, let v_1 and v_2 be the vertices to be merged. The partition function P of the original graph (where v_1 and v_2 are not merged) can be written as

$$P(z_1, z_2, z_3, \ldots, z_n) = A(z_3, \ldots, z_n)z_1z_2 + B(z_3, \ldots, z_n)z_1 + C(z_3, \ldots, z_n)z_2 + D,$$

and is LY by the hypothesis of the operation. The partition function R after the merger is then given by

$$R(z, z_3, \ldots, z_n) = A(z_3, \ldots, z_n)z + D,$$

where z is a new variable corresponding to the new vertex created by the merger of v_1 and v_2. Now, let $\lambda_3, \ldots, \lambda_n$ be complex numbers lying outside the open unit disk. Corollary 4.7 then implies that $A = A(\lambda_3, \ldots, \lambda_n) \neq 0$. Thus, $R(z, \lambda_3, \ldots, \lambda_n) = 0$ implies that

$$|z| = |D/A| = |D(\lambda_3, \ldots, \lambda_n)/A(\lambda_3, \ldots, \lambda_n)| \,. \tag{20}$$

Now, since P is LY, both zeros of the quadratic equation $P(x, x, \lambda_3, \ldots, \lambda_n) = 0$ satisfy $|x| \leq 1$, and indeed, $|x| < 1$ when at least of the λ_i lies strictly outside the closed unit disk. Thus, the product D/A of its zeros also satisfies $|D/A| \leq 1$, and further satisfies the stronger inequality $|D/A| < 1$ in case at least one of the λ_i lies strictly outside the closed unit disk. Eq. (20) then implies that $|z| \leq 1$ in the first case and $|z| < 1$ in the second case, which establishes that R is LY.

This concludes the proof of Theorem 4.4. □

References

[1] Nima Anari and Shayan Oveis Gharan. The Kadison-Singer problem for strongly Rayleigh measures and applications to Asymmetric TSP. In *Proceedings of the 56th IEEE Symposium on Foundations of Computer Science (FOCS)*, October 2015. Available at http://arxiv.org/abs/1412.1143.

[2] Nima Anari and Shayan Oveis Gharan. A generalization of permanent inequalities and applications in counting and optimization. *arXiv:1702.02937 [cs, math]*, February 2017.

[3] Taro Asano. Lee-Yang Theorem and the Griffiths Inequality for the Anisotropic Heisenberg Ferromagnet. *Physical Review Letters*, 24(25):1409–1411, June 1970.

[4] J. C. A. Barata and D. H. U. Marchetti. Griffiths' singularities in diluted Ising models on the Cayley tree. *J. Stat. Phys.*, 88(1-2):231–268, July 1997.

[5] João C. A. Barata and Pedro S. Goldbaum. On the distribution and gap structure of Lee–Yang zeros for the Ising model: Periodic and aperiodic couplings. *J. Stat. Phys.*, 103(5-6):857–891, June 2001.

[6] Alexander Barvinok. Computing the partition function for cliques in a graph. *Theory of Computing*, 11(13):339–355, 2015.

[7] Alexander Barvinok. Computing the permanent of (some) complex matrices. *Found. Comput. Math.*, 16(2):329–342, January 2015. arXiv:1405:1303.

[8] Alexander Barvinok and Pablo Soberón. Computing the partition function for graph homomorphisms. *Combinatorica*, May 2016. Available online at http://dx.doi.org/10.1007/s00493-016-3357-2.

[9] Alexander Barvinok and Pablo Soberón. Computing the partition function for graph homomorphisms with multiplicities. *J. Combin. Theory, Ser. A*, 137:1–26, January 2016.

[10] Noam Berger, Claire Kenyon, Elchanan Mossel, and Yuval Peres. Glauber dynamics on trees and hyperbolic graphs. *Probability Theory and Related Fields*, 131(3):311–340, 2005.

[11] Julius Borcea and Petter Brändén. The Lee-Yang and Pólya-Schur programs. I. Linear operators preserving stability. *Invent. Math.*, 177(3):541–569, 2009.

[12] Julius Borcea and Petter Brändén. The Lee-Yang and Pólya-Schur programs. II. Theory of stable polynomials and applications. *Comm. Pure Appl. Math.*, 62(12):1595–1631, 2009.

[13] Julius Borcea, Petter Brändén, and Thomas Liggett. Negative dependence and the geometry of polynomials. *J. AMS*, 22(2):521–567, 2009.

[14] Christian Borgs, Jennifer Chayes, Jeff Kahn, and László Lovász. Left and Right Convergence of Graphs with Bounded Degree. *Random Struct. Algorithms*, 42(1):1–28, January 2013.

[15] Jin-Yi Cai, Xi Chen, and Pinyan Lu. Graph homomorphisms with complex values: A dichotomy theorem. In *Automata, Languages and Programming*, volume 6198 of *Lecture Notes in Computer Science*, pages 275–286. Springer Berlin/Heidelberg, 2010.

[16] Andreas Galanis, Qi Ge, Daniel Štefankovič, Eric Vigoda, and Linji Yang. Improved inapproximability results for counting independent sets in the hard-core model. *Random Struct. Algorithms*, 45(1):78–110, 2014.

[17] Leslie Ann Goldberg, Martin Grohe, Mark Jerrum, and Marc Thurley. A complexity dichotomy for partition functions with mixed signs. *SIAM Journal on Computing*, 39(7):3336–3402, January 2010.

[18] Leslie Ann Goldberg, Mark Jerrum, and Mike Paterson. The computational complexity of two-state spin systems. *Random Structures & Algorithms*, 23(2):133–154, 2003.

[19] Heng Guo and Mark Jerrum. Random cluster dynamics for the Ising model is rapidly mixing. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1818–1827. SIAM, 2017.

[20] Heng Guo and Pinyan Lu. Uniqueness, spatial mixing, and approximation for ferromagnetic 2-spin systems. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 2016.

[21] E. Ising. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift fur Physik*, 31:253–258, February 1925.

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

[23] Mark Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theor. Comput. Sci.*, 43:169–188, 1986.

[24] T. D. Lee and C. N. Yang. Statistical theory of equations of state and phase transitions. II. Lattice gas and Ising model. *Phys. Rev.*, 87(3):410–419, August 1952.

[25] Liang Li, Pinyan Lu, and Yitong Yin. Correlation decay up to uniqueness in spin systems. In *Proc. 24th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 67–84. SIAM, 2013.

[26] Yun Long, Asaf Nachmias, Weiyang Ning, and Yuval Peres. A power law of order 1/4 for critical mean-field swendsen-wang dynamics. *Memoirs of the AMS*, 232(1092), 2014.

[27] Pinyan Lu, Kuan Yang, and Chihao Zhang. FPTAS for Hardcore and Ising models on hypergraphs. In *33rd Symposium on Theoretical Aspects of Computer Science*, 2016.

[28] Eyal Lubetzky and Allan Sly. Critical ising on the square lattice mixes in polynomial time. *Communications in Mathematical Physics*, 313(3):815–836, 2012. Communicated by F. Toninelli.

[29] Adam Marcus, Daniel Spielman, and Nikhil Srivastava. Interlacing families I: Bipartite Ramanujan graphs of all degrees. *Ann. Math.*, pages 307–325, July 2015.

[30] Adam W Marcus, Daniel A Spielman, and Nikhil Srivastava. Interlacing families II: Mixed characteristic polynomials and the Kadison-Singer problem. *Ann. Math.*, 182:327–350, 2015.

[31] F. Martinelli and E. Olivieri. Approach to equilibrium of Glauber dynamics in the one phase region. I. The attractive case. *Comm. Math. Phys.*, 161(3):447–486, 1994.

[32] F. Martinelli and E. Olivieri. Approach to equilibrium of Glauber dynamics in the one phase region: II. The general case. *Comm. Math. Phys.*, 161:487–514, 1994.

[33] Fabio Martinelli, Alistair Sinclair, and Dror Weitz. Glauber dynamics on trees: Boundary conditions and mixing time. *Communications in Mathematical Physics*, 250(2):301–334, 2004.

[34] Elchanan Mossel and Allan Sly. Exact thresholds for Ising-Gibbs samplers on general graphs. *Ann. Probab.*, 41(1):294–328, 2013.

[35] Viresh Patel and Guus Regts. Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials. *arXiv:1607.01167 [cs, math]*, July 2016. arXiv: 1607.01167.

[36] Dana Randall and David Wilson. Sampling spin configurations of an Ising system. In *Proc. 10th ACM-SIAM Symp. Discret. Algorithms (SODA)*, SODA '99, pages 959–960, Philadelphia, PA, USA, 1999. SIAM.
[37] David Ruelle. Characterization of Lee-Yang polynomials. *Ann. Math.*, 171(1):589–603, March 2010.
[38] Alex Scott and Alan Sokal. The Repulsive Lattice Gas, the Independent-Set Polynomial, and the Lovász Local Lemma. *J. Stat. Phys.*, 118(5-6):1151–1261, 2004.
[39] James B. Shearer. On a problem of Spencer. *Combinatorica*, 5(3), 1985.
[40] Alistair Sinclair and Piyush Srivastava. Lee-Yang theorems and the complexity of computing averages. *Comm. Math. Phys.*, 329(3):827–858, August 2014.
[41] Alistair Sinclair, Piyush Srivastava, and Marc Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. *J. Stat. Phys.*, 155(4):666–686, 2014.
[42] Allan Sly and Nike Sun. Counting in two-spin models on d-regular graphs. *Ann. Probab.*, 42(6):2383–2416, November 2014.
[43] R.P. Stanley and S. Fomin. *Enumerative Combinatorics:*. Number v. 2 in Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1999.
[44] Masuo Suzuki and Michael E. Fisher. Zeros of the partition function for the Heisenberg, Ferroelectric, and general Ising models. *J. Math. Phys.*, 12(2):235–246, 1971.
[45] Dror Weitz. Counting independent sets up to the tree threshold. In *Proc. 38th ACM Symp. on Theory of Computing (STOC)*, pages 140–149. ACM, 2006.
[46] C. N. Yang and T. D. Lee. Statistical theory of equations of state and phase transitions. I. Theory of condensation. *Phys. Rev.*, 87(3):404–409, August 1952.
[47] Jinshan Zhang, Heng Liang, and Fengshan Bai. Approximating partition functions of the two-state spin system. *Inf. Process. Lett.*, 111(14):702–710, July 2011.