

Symbolic integration and the complexity of computing averages

Leonard J. Schulman*, Alistair Sinclair†, Piyush Srivastava*

**Computing and Mathematical Sciences*
California Institute of Technology
Pasadena, USA

Email: {schulman|piyushs}@caltech.edu

†*Computer Science Division*
University of California, Berkeley
Berkeley, USA

Email: sinclair@cs.berkeley.edu

Abstract

We study the computational complexity of several natural problems arising in statistical physics and combinatorics. In particular, we consider the following problems: the mean magnetization and mean energy of the Ising model (both the ferromagnetic and the anti-ferromagnetic settings), the average size of an independent set in the hard core model, and the average size of a matching in the monomer-dimer model. We prove that for all non-trivial values of the underlying model parameters, exactly computing these averages is #P-hard.

In contrast to previous results of Sinclair and Srivastava [1] for the mean magnetization of the *ferromagnetic* Ising model, our approach does not use any Lee-Yang type theorems about the complex zeros of partition functions. Indeed, it was due to the lack of suitable Lee-Yang theorems for models such as the anti-ferromagnetic Ising model that some of the problems we study here were left open in [1]. In this paper, we instead use some relatively simple and well-known ideas from the theory of automatic symbolic integration to complete our hardness reductions.

Keywords

Computational Complexity; Statistical Mechanics; Counting Problems; #P-hardness

I. INTRODUCTION

A. Background

Let $G = (V, E)$ be graph and Ω a set of combinatorial structures defined on the graph. We assume that these configurations have positive rational *weights* $w : \Omega \rightarrow \mathbb{Q}$, along with a positive integer valued *observable* $f : \Omega \rightarrow \mathbb{Z}$ defined on the configurations. Our goal in this paper is to study the computational complexity of the mean value $\langle f \rangle$ of such observables, which is defined as:

$$\langle f \rangle := \frac{1}{Z} \sum_{\sigma \in \Omega} w(\sigma) f(\sigma).$$

Here $Z := \sum_{\sigma \in \Omega} w(\sigma)$ is the *partition function*. Two of the simplest examples of such problems include computing the average size of independent sets and matchings in graphs: in this case Ω is the set of all independent sets (or matchings) of G , $w(\sigma) = 1$ for all $\sigma \in \Omega$, and $f(\sigma)$ is the size of the independent set σ (or the matching σ). Note that in both cases, for all $\sigma \in \Omega$, it is easy to compute both the weight $w(\sigma)$ as well as the observable $f(\sigma)$, but the partition function Z is #P-hard to compute.

Although both of these examples are naturally motivated combinatorial problems, we will find it fruitful to study them in the more general context of *spin systems*. These originated in the early twentieth century in statistical physics as a tool for studying phase transitions in magnetism, but have since been studied extensively in combinatorics and computer science as well, and have been applied to the modeling of large systems in a variety of settings. Computational problems of the form outlined above appear naturally in the study of these

systems (starting with the original applications in statistical physics, as we shall see below), and spin systems provide a general framework for modeling and studying them. We now proceed to briefly describe the general properties of spin systems followed by three well-studied classical examples. We will see that all the problems that we study in this paper, including the computation of average sizes of matchings and independent sets, find a natural expression in the language of spin systems.

In a spin system, the configurations $\sigma : V \rightarrow \{+, -\}$ are assignments of $\{+, -\}$ spins to the vertices of G (in general the number of allowed spins can be larger, but all the systems we consider in this paper require only two spins). Further, there is a small number of observables f_1, f_2, \dots, f_k and positive rational parameters c_1, c_2, \dots, c_k such that the weight $w(\sigma)$ can be written as

$$w(\sigma) = \prod_{i=1}^k c_i^{f_i(\sigma)}.$$

Viewing the partition function Z as a function of the parameters c_i , we then see that there is a simple analytical relationship between $Z(c_1, c_2, \dots, c_k)$ and the mean values $\langle f_i \rangle$. Formally, we have

$$\langle f_i \rangle := \frac{1}{Z} \sum_{\sigma \in \Omega} w(\sigma) f_i(\sigma) = c_i \frac{\partial}{\partial c_i} \log Z(c_1, c_2, \dots, c_k) = c_i \frac{Z'(c_1, c_2, \dots, c_k)}{Z(c_1, c_2, \dots, c_k)},$$

where Z' denotes the derivative with respect to c_i (the parameter with respect to which the derivative is being taken will always be clear from the context). Thus, we see that the definition of spin systems encodes mean values of various observables as derivatives of the logarithm of the partition function: this correspondence between derivatives and mean values of observables is standard in the study of spin systems in statistical physics, and also turns out to be very important for our goal of understanding the computational complexity of these mean quantities.

We now give examples of some well studied spin systems, all of which will be relevant to this paper.

Example I.1 (The hard core model). This is a distribution over the independent sets of G . Formally, the configurations are assignments of $\{+, -\}$ spins to the vertices of G , with any configuration that assigns spin $+$ to any pair of adjacent vertices having weight zero. (The vertices with $+$ spin correspond to the elements of the independent set.) The natural observable $f(I)$ here is the size $|I|$ of the independent set (i.e., the number of $+$ spins). The weight $w_H(I)$ and the partition function Z_H are then given by

$$w_H(I) = \lambda^{|I|} \quad \text{and} \quad Z_H(G, \lambda) = \sum_{I \in \Omega} \lambda^{|I|},$$

where the subscript H denotes the name of the model. Here $\lambda > 0$ is a model parameter that is known as the *vertex activity* or *fugacity*. The natural mean observable is the average size $M_H(G, \lambda)$ of the independent sets, and is given by

$$M_H(G, \lambda) = \lambda \frac{d}{d\lambda} \log Z_H(G, \lambda). \tag{1}$$

Example I.2 (Ising model). Historically, this was the first spin system to be studied, and was proposed by Lenz as a model for magnetism in bulk matter (the model first appeared, however, in a paper of Ising [2]). Here, unlike the case of the hard core model, all configurations $\sigma : V \rightarrow \{+, -\}$ are permissible, and we have two observables, the *energy* $d(\sigma)$, which is defined as the number of edges $\{u, v\}$ in G such that $\sigma(u) \neq \sigma(v)$, and the *magnetization* $m(\sigma)$ which is defined as the number of vertices v such that $\sigma(v) = +$. The weight $w_I(\sigma)$ and the partition function $Z_I(G, \beta, \lambda)$ are then given by

$$w_I(\sigma) = \beta^{d(\sigma)} \lambda^{m(\sigma)} \quad \text{and} \quad Z_I(G, \beta, \lambda) = \sum_{\sigma: V \rightarrow \{+, -\}} w_I(\sigma),$$

where β is a model parameter called the *edge activity* and λ is a model parameter called the *vertex activity*. (Note that in the statistical physics literature, it is customary to parametrize the models in terms of an *inverse temperature* proportional to $\log \beta$ and an *external field* proportional to $\log \lambda$, but we will not use this terminology in this paper.) There are two natural and well studied mean quantities in this setting: the mean magnetization $M_{IM}(G, \beta, \lambda)$ and the mean energy $M_{IE}(G, \beta, \lambda)$, where

$$M_{IM}(G, \beta, \lambda) = \lambda \frac{\partial}{\partial \lambda} \log Z_I(G, \beta, \lambda) \quad \text{and} \quad M_{IE}(G, \beta, \lambda) = \beta \frac{\partial}{\partial \beta} \log Z_I(G, \beta, \lambda). \quad (2)$$

Note that the *energy* $d(\sigma)$ is exactly the size of the cut induced by the labeling σ . Thus, when $\lambda = 1$, the mean energy is just the average size of the cut induced by the Ising model. Of course, this is trivial to compute when $\beta = 1$. Similarly, the mean magnetization is $|V|/2$ when $\lambda = 1$, but as we show later, both these quantities are #P-hard to compute at all other values of the parameters.

The Ising model has very different qualitative behavior in the two settings $0 < \beta < 1$ and $\beta > 1$. In the first case, known as the *ferromagnetic Ising model*, configurations with smaller cuts get larger weight, so that neighboring vertices tend to have similar spins in a configuration sampled from the model. On the other hand, in the *anti-ferromagnetic Ising model*, which corresponds to the setting $\beta > 1$, disagreements between neighboring vertices are favored. The contrast is most visible if we consider the case $\lambda = 1$: the maximum weight configurations in the ferromagnetic case then are those in which all vertices have the same spin, while the maximum weight configurations in the latter case are the max-cuts of the graph.

Example I.3 (Monomer-dimer model [3]). In this model, the configurations in Ω are matchings of the graph G . (Thus, strictly speaking, this is a spin system not on the graph G , but on the *line graph* of G . However, this distinction will not be important for our purposes.) The natural observable is the size $|M|$ of a matching, but for technical reasons, we consider instead the equivalent observable $u(M)$ which counts the number of vertices left unmatched by the matching M . The weight $w_D(M)$ and the partition function $Z_M(D, \lambda)$ are then given by

$$w_D(M) = \lambda^{u(M)} \quad \text{and} \quad Z_D(G, \lambda) = \sum_{M \in \Omega} w_D(M),$$

where λ is a model parameter known as the *vertex activity* or *monomer activity*. The mean number of monomers $M_D(G, \lambda)$ is then given by

$$M_D(G, \lambda) = \lambda \frac{d}{d\lambda} \log Z_D(G, \lambda).$$

Note that the average size of matchings can be written as $\frac{1}{2}(|V| - M_D(G, \lambda))$. It is also possible to add edge activities to the model, but for notational simplicity we do not do so here. (However, our hardness result, Theorem I.4 stated later, continues to hold for the more general model as well.)

In applications of spin systems, the average quantities mentioned above have important operational meanings. For example, the “mean magnetization” in the classical Ising model of magnetism is so named because its behavior indeed models the (physical) magnetization. Similarly, when the Ising model is used to model binary preferences in social networks, the mean magnetization models the popularity of one of the preferences. For these reasons, the question of algorithmically computing such averages is of great importance. As a result, and in the absence of efficient algorithms for exact computation, the *approximation* problem for such mean quantities is quite well studied. A large body of work addresses the problem of approximately *sampling* configurations from the probability distribution (called the *Gibbs distribution*) induced by the weights of a given system. Since the mean quantities mentioned above all have variances (with respect to the Gibbs distribution) which are polynomially bounded in the size of the graph, such a sampling procedure, when available, can then be used to give efficient randomized additive approximation procedures for the mean quantities. This approach has been shown to work, e.g., for the *ferromagnetic Ising model* and the *monomer-dimer model* (see, e.g., the papers by Jerrum and Sinclair [4] and Randall and Wilson [5] for the Ising model, and by Jerrum and Sinclair [6] for the monomer-dimer model).

However, in spite of the significant progress made in designing approximation algorithms for these mean quantities, it was not known until recently if any of these quantities are actually hard to compute *exactly*. Although this is an interesting natural problem, the first work directly addressing it appears to be a relatively recent paper of two of the present authors [1], where the #P-hardness of computing the mean magnetization of the Ising model and mean monomer count in the monomer-dimer model on multigraphs is proved. In the next section, we begin by briefly outlining the approach taken in that paper, and the problems it faces in handling other spin systems such as the anti-ferromagnetic Ising and hard core models, and then give a brief outline of our approach which surmounts these problems. We use the mean magnetization of the Ising model with a fixed edge activity $\beta \neq 1$ as a running example in this discussion. Finally, at the end of the section, we state our complexity results.

B. Techniques and results

The starting point of [1] is eq. (2), which shows that $M_{IM}(\beta, \lambda) = \frac{\lambda \frac{\partial Z_I(G, \beta, \lambda)}{\partial \lambda}}{Z_I(G, \beta, \lambda)}$ is a rational function of λ . Now, suppose we have a hypothetical algorithm that for a fixed β , efficiently computes the values $M_{IM}(G, \beta, x)$ for any inputs G and x . The time-honored way to proceed then would be to use these evaluations as an input to an *interpolation* procedure which would determine the coefficients of the unreduced form of the rational function $M_{IM}(G, \beta, \lambda)$. These coefficients include, in particular, the coefficients of the partition function $Z_I(G, \beta, \lambda)$. Since the partition function itself is known to be #P-hard to compute, guaranteed success of this process would ensure that the problem of computing $M_{IM}(G, \beta, \lambda)$ is #P-hard as well.

Unfortunately, we are not done yet since one cannot *a priori* guarantee that the rational interpolation succeeds in retrieving the coefficients of Z_I , as Z_I and $\frac{\partial Z_I}{\partial \lambda}$ may have common factors, and in that case any interpolation procedure will only be able to recover a reduced form of $M_{IM}(G, \beta, \lambda)$ whose coefficients can be quite different from those of Z_I . To get around this, Sinclair and Srivastava [1] showed that such cancellations *cannot* occur when one is considering the *ferromagnetic* Ising model on connected graphs at non-trivial values of the parameters. This required them to prove an extension of the Lee-Yang theorem [7] on the complex zeros of the partition function Z_I of the ferromagnetic Ising model, which showed that Z_I , when viewed as a polynomial in λ , cannot have repeated complex zeros under this setting (which in turn implies that it cannot share factors with its derivative). Their companion result in [1] for the monomer-dimer model uses a similar strategy, appealing to a classical result of Heilmann and Lieb [3] on the non-degeneracy of zeros of the partition function Z_D on Hamiltonian graphs.

While their results on zeros of the partition function remain of independent interest, as pointed out in [1], their approach necessarily fails for models such as the anti-ferromagnetic Ising model and the hard core model because similar strong results for the non-degeneracy of zeros of the partition function are not true for these models. Indeed, it is not hard to construct simple examples of connected graphs where the partition functions of these models indeed have degenerate zeros.

In this paper, we follow a different approach where we do not try to guarantee that the interpolation will succeed in determining Z_I , but instead view eq. (2) as a differential equation for Z_I . We then attempt to use the hypothetical algorithm for computing the mean magnetization M_{IM} to integrate this differential equation. Again, one might attempt to do this numerically, by trying to argue that if evaluations of $M_{IM}(G, \beta, x)$ at sufficiently many distinct values of x were available, then one could determine $\log Z_I(G, \beta, \lambda)$ (for some appropriate λ where the partition function is known to hard to compute) to sufficient accuracy via a suitable numerical quadrature rule. However, since we are trying to prove #P-hardness and hence need to evaluate $Z_I(G, \beta, \lambda)$ exactly, the required accuracy is likely to be too high for any numerical quadrature rule to be efficient. (This problem is further exacerbated by the fact that the differential equation is for the transcendental function $\log Z_I$ and not for Z_I .)

Instead, we attempt to integrate the differential equation *symbolically*. Our first step is to use rational interpolation to compute the coefficients of $M_{IM}(G, \beta, \lambda)$. However, unlike in [1], we admit the possibility that this may not yield the coefficients of $Z_I(G, \beta, \lambda)$ due to cancellations of common factors between Z_I and its derivative. To get around this problem, we instead use classical ideas from the theory of automatic symbolic integration (going back at least to the work of Horowitz [8] and Tobey [9], but likely to be even older) to symbolically integrate

the reduced form of $M_{IM}(G, \beta, \lambda)$ in order to obtain $\log Z_I(G, \beta, \lambda)$ symbolically. The symbolic nature of the computation means that the transcendental nature of this function is no longer a problem. We now proceed to state our complexity results.

Remark I.1. Note that in all the results below, #P-hardness continues to hold even when the model parameters (such as the edge activity β or the vertex activity λ) are *not* part of the input but fixed in advance. However, in the technical discussion at the beginning of this section, we discussed proving hardness for the case where at least one of the model parameters was actually part of the input. (For example, in the outline of the reduction, we assumed that we have a hypothetical algorithm that computes $M_{IM}(G, \beta, x)$ for a fixed β and G, x as inputs.) However, as is also done in [1], it is easy to simulate different values of the model parameters by modifying the input graph, and such simulations can be used to prove hardness even for the restricted problem where the model parameters are not input to the algorithm but fixed in advance. These simulations are identical in spirit to those used in [1] and we therefore defer their discussion to Sections IV and V.

Our first result extends the main complexity result of Sinclair and Srivastava [1] for the ferromagnetic Ising model to the anti-ferromagnetic setting. As discussed above, the methods used in [1] face a fundamental obstacle when confronted with this model, since the analog of the extended Lee-Yang theorem proved in [1] is actually false for the anti-ferromagnetic Ising model.

Theorem I.1. *Fix positive rational $\beta > 1$ and $\lambda \neq 1$. Then, the problem of computing the magnetization of the Ising model with edge activity β and vertex activity λ is #P-hard. Further, the problem remains #P-hard even when the input is restricted to graphs of degree at most 4.*

Note that the problem is in P when $\lambda = 1$: the magnetization is then $|V|/2$ by symmetry.

Sinclair and Srivastava [1] left open the computational complexity of the mean energy of the Ising model. Using our methods, we are able to resolve this problem.

Theorem I.2. *Fix a positive rational $\beta \neq 1$. Then, the problem of computing the mean energy of the Ising model with edge activity β (and vertex activity $\lambda = 1$) is #P-hard. Further, the problem remains hard even when the input is restricted to graphs of degree at most 3.*

Note that the above result holds for both the ferromagnetic and anti-ferromagnetic Ising models, and also that the problem is trivial when $\beta = 1$ (the graph is effectively completely disconnected in this case). Our next result concerns the hard core model, which was also left open in [1].

Theorem I.3. *Fix rational $\lambda > 0$. Then, the problem of computing the average size of an independent set in the hard core model with vertex activity λ is #P-hard. Further, the problem remains #P-hard even when the input is restricted to graphs of degree at most 4.*

By taking $\lambda = 1$ in the above result, we see that computing the average size of independent sets in a graph is #P-hard. Note that a weaker version of this result, where λ is restricted to be the square of a rational number and where the input graph has degree at most 10, can be obtained as a corollary of Theorem I.4 below by using the equivalence between the monomer-dimer model on a graph G and the hard core model on the line graph of G . However, since our direct reduction given in Section IV does not require these restrictions, we omit the (easy) details of this fact.

Finally, we use our methods to improve upon and simplify the main complexity results of Sinclair and Srivastava [1] for the monomer-dimer model. In particular, we prove the following.

Theorem I.4. *Fix a positive rational λ . Then, the problem of computing the average number of matchings (equivalently, the mean number of monomers) in the monomer-dimer model with vertex activity λ is #P-hard. Further, the problem remains #P-hard even when the input is restricted to graphs of degree at most 5.*

By taking $\lambda = 1$ in the above result we see that computing the average number of matchings in a graph is #P-hard.

In contrast, the hardness result proved in [1] could only achieve this for multigraphs.

Remark I.2. We observe that our results in this paper are obtained without appealing to deep facts about the location of zeros of partition functions (such as the Lee-Yang or Heilmann-Lieb theorems or their extensions). On the other hand, we should also note that our results do not have any consequences for the study of zeros of partition functions, which is also of interest for reasons beyond computational complexity theory, and has connections, for example, with the theory of stable polynomials (see, e.g., [10], [11]). In particular, our techniques here cannot prove results on non-degeneracy of zeros of partition functions, as was done in the main technical result of [1].

II. PRELIMINARIES

In this section we review some standard facts and observations about polynomials with rational coefficients.

A. Polynomials in $\mathbb{Q}[x]$

We start by recalling some standard terms for future reference. The *description length* of a polynomial $p(x) \in \mathbb{Q}[x]$ is the number of bits required to write down all its coefficients, where each coefficient is written as a quotient of two integers written in binary. We will often shorten this to *length* if the meaning is clear from the context. We denote the degree of a polynomial $p(x) \in \mathbb{Q}[x]$ as $\deg(p(x))$ and the g.c.d. of polynomials $p(x)$ and $q(x)$ as $\gcd(p(x), q(x))$. We assume that the g.c.d. of two polynomials is normalized to be *primitive* as defined below. The polynomials $p(x)$ and $q(x)$ are *co-prime* if their g.c.d. is 1.

We call a polynomial $p(x) \in \mathbb{Q}[x]$ of positive degree *primitive* if all its coefficients are integers (i.e., $p(x) \in \mathbb{Z}[x]$), its leading coefficient is positive, and the g.c.d. of its coefficients is 1. Note that given a polynomial $q(x) \in \mathbb{Q}[x]$, there is a unique primitive polynomial $p(x)$ and a unique $c \in \mathbb{Q}$ such that $p(x) = c \cdot q(x)$. Further, given $q(x)$, both c and $p(x)$ can be determined efficiently (in time polynomial in the description length of $q(x)$). We will need the following standard fact about primitive polynomials.

Lemma II.1 (Gauss, [12, Lemma 3.10.1]). *If $p(x), q(x) \in \mathbb{Q}[x]$ are primitive then so is their product $p(x)q(x)$.*

We call $p(x) \in \mathbb{Q}[x]$ *irreducible* if $p(x)$ cannot be written as a product of two polynomials in $\mathbb{Q}[x]$ both of which are of positive degree. If $p(x) \in \mathbb{Q}[x]$ is *irreducible*, it is also *prime*, i.e., if $p(x)$ divides $g(x)h(x)$ (where $g(x), h(x) \in \mathbb{Q}[x]$), then it must divide at least one of $g(x)$ or $h(x)$. The following well known result of Lenstra, Lenstra and Lovász [13] is used to establish the algorithmic efficiency of a part of our main technical lemma in Section III.

Theorem II.2 (Efficient polynomial factoring [13]). *There exists an algorithm which, given a polynomial $p(x) \in \mathbb{Q}[x]$, outputs in time polynomial in the description length of p its unique primitive irreducible factorization, i.e., a rational number c , distinct primitive irreducible polynomials $p_i(x)$, and positive integers d_i such that*

$$f(x) = c \cdot \prod_{i=1}^k p_i(x)^{d_i},$$

where k is the number of distinct irreducible factors of $p(x)$.

B. Rational functions

We will also need to consider *rational functions* of the form $\frac{p(x)}{q(x)}$ where $p(x), q(x) \in \mathbb{Q}[x]$. We will view these both as elements of the field of fractions over $\mathbb{Q}[x]$, and as functions from \mathbb{Q} to \mathbb{Q} . However, we define equality of rational functions by viewing them as elements of the field of fractions over $\mathbb{Q}[x]$ and not as functions. Formally, we say that $\frac{p(x)}{q(x)} = \frac{r(x)}{s(x)}$ if and only if

$$p(x)s(x) - r(x)q(x) \text{ is the zero polynomial.}$$

Note that this definition is slightly weaker than that obtained by considering $\frac{p(x)}{q(x)}$ and $\frac{r(x)}{s(x)}$ as functions: for example, $\frac{1}{x}$ and $\frac{x-1}{x(x-1)}$ are equal under the above definition, but as functions one of them is not defined at $x = 1$ while the other is. However, equality of two rational functions in the field of fractions over $\mathbb{Q}[x]$ *does* imply equality of their evaluations at all points where both are well defined as functions (i.e., at points where both have a non-zero denominator). We will therefore be careful to ensure this when using the above notion of equality to argue about evaluations of rational functions.

We also note here the standard fact that the reduced form of a rational function is essentially unique.

Fact II.3. *Let $a(x), b(x), c(x), d(x) \in \mathbb{Q}[x]$ such that $\gcd(a(x), b(x)) = 1 = \gcd(c(x), d(x))$. Then, $\frac{a(x)}{b(x)} = \frac{c(x)}{d(x)}$ implies that there exists a non-zero $s \in \mathbb{Q}$ such that $a(x) = s \cdot c(x)$ and $b(x) = s \cdot d(x)$.*

The following theorem is folklore.

Theorem II.4 (Uniqueness of partial fraction expansions). *Let $a(x)$ and $b(x)$ be polynomials in $\mathbb{Q}[x]$, where $b(x)$ is square-free and primitive. Let $p_i(x)$, $i = 1, \dots, k$ be the distinct primitive irreducible factors of $b(x)$, so that $b(x) = \prod_{i=1}^k p_i(x)$. Then, there exists at most one sequence of polynomials $q_i(x) \in \mathbb{Q}[x]$ satisfying $\deg(q_i(x)) < \deg(p_i(x))$ for $i = 1, \dots, k$ and*

$$\frac{a(x)}{b(x)} = \sum_{i=1}^k \frac{q_i(x)}{p_i(x)}. \quad (3)$$

In fact, when $\deg(a(x)) < \deg(b(x))$, it can also be shown that such a sequence q_i always exists, but we will not need to invoke this fact. We defer the proof of this standard result to the appendix.

C. Rational interpolation

We will need the following standard fact about interpolation of rational functions.

Theorem II.5 (Rational Interpolation [14]). *Let $R(x) := \frac{p(x)}{q(x)}$ be a rational function with rational coefficients such that $\deg(p(x)) = n$, $\deg(q(x)) = m > n$ are known. Suppose we have evaluations $R(y_i)$ of R at $m + n + 2$ distinct points $y_1, y_2, \dots, y_{m+n+2} \in \mathbb{Q}$ at which the $q(y_i)$ are non-zero. Then, the evaluation pairs $(y_i, R(y_i))_{i=1}^{m+n+2}$ uniquely determine polynomials $a(x), b(x) \in \mathbb{Q}[x]$ satisfying the following conditions:*

- 1) $\gcd(a(x), b(x)) = 1$, $\deg(a(x)) \leq n$, $\deg(b(x)) \leq m$, and $b(x)$ is a primitive polynomial.
- 2) $R(x) = \frac{a(x)}{b(x)}$.

Remark II.1. Note that given the evaluation pairs $(y_i, R(y_i))_{i=1}^{m+n+2}$, the polynomials $a(x)$ and $b(x)$ can be determined in time polynomial in the description length of the pairs. To see why this is true, observe that $R(x) = a(x)/b(x)$ implies that we must have $p(x) = s(x)a(x)$ and $q(x) = s(x)b(x)$ for some polynomial $s(x) \in \mathbb{Q}[x]$ (and indeed $s(x) = c \gcd(p(x), q(x))$ for some rational c). Now, suppose that the degree $k \leq m$ of $s(x)$ is known. Then, the degrees of a and b are $n - k$ and $m - k$ respectively, and we can write down a system of linear equations for the coefficients of $a(x)$ and $b(x)$ by using the evaluations. This system can be solved in polynomial time using Gaussian elimination to yield $a(x)$ and $b(x)$ up to a constant factor. This factor can then be determined using the condition that $b(x)$ must be primitive.

The above argument assumes that we know k , the degree of $\gcd(p(x), q(x))$. However, this is not a problem, since we can try all values of k from m to 0 in decreasing order, and the first value for which Gaussian elimination actually finds a solution gives the right answer.

III. SYMBOLIC INTEGRATION OF RATIONAL FUNCTIONS

In this section we prove the following lemma, which is the main tool used in our reductions.

Lemma III.1. *Let $p(x)$ be an (unknown) polynomial with rational coefficients, and let $a(x)$ and $b(x)$ be co-prime polynomials in $\mathbb{Q}[x]$ such that $b(x)$ is primitive, and such that*

$$\frac{d}{dx} \log p(x) = \frac{p'(x)}{p(x)} = \frac{a(x)}{b(x)}. \quad (4)$$

Then, given the polynomials $a(x)$ and $b(x)$, and one non-zero coefficient of $p(x)$, the polynomial $p(x)$ can be determined in time polynomial in the length of $a(x)$ and $b(x)$.

The ideas behind the proof of this theorem go back at least to Horowitz [8], who in turn cites earlier work of Tobey [9], and are likely to be folklore. Note that the algorithm promised by the lemma essentially outputs an anti-derivative of the function $\frac{a(x)}{b(x)}$.

Proof: As a warm-up, we first consider the simple case where p is square-free so that $\gcd(p(x), p'(x)) = 1$.¹ In this case, by Fact II.3, we must have $a(x) = cp'(x)$ and $b(x) = cp(x)$, for some non-zero rational number c , and c itself can be determined in polynomial time by comparing the one known coefficient of $p(x)$ with the corresponding coefficient of $b(x)$.

We now consider the general case. Thinking of $p(x)$ as a polynomial in $\mathbb{Q}[x]$, we consider its unique factorization

$$p(x) = c \prod_{i=1}^k p_i(x)^{d_i},$$

where $c \in \mathbb{Q}$, and $p_i(x) \in \mathbb{Z}[x]$ are distinct irreducible primitive polynomials with positive leading terms. We then have

$$p'(x) = c \sum_{i=1}^k d_i p_i'(x) p_i(x)^{d_i-1} \prod_{\substack{1 \leq j \leq k \\ j \neq i}} p_j(x)^{d_j}.$$

Using these expansions (and the fact that the p_i are distinct irreducible polynomials), we get that

$$\gcd(p(x), p'(x)) = \prod_{i=1}^k p_i(x)^{d_i-1} \quad \text{and} \quad \frac{p'(x)}{p(x)} = \frac{g(x)}{h(x)},$$

where

$$g(x) := \sum_{i=1}^k d_i p_i'(x) \prod_{\substack{1 \leq j \leq k \\ j \neq i}} p_j(x) \quad \text{and} \quad h(x) := \prod_{i=1}^k p_i(x), \quad (5)$$

with $\gcd(g(x), h(x)) = 1$. (To see this, note that since the $p_i(x)$ are distinct irreducible polynomials, none of them can divide $g(x)$.)

Using the hypothesis of the lemma, we therefore get

$$\frac{g(x)}{h(x)} = \frac{a(x)}{b(x)}.$$

Now, since $g(x), h(x)$ and $a(x), b(x)$ are co-prime in pairs, Fact II.3 implies that there is a rational number $c_1 \neq 0$ such that $a(x) = c_1 g(x)$ and $b(x) = c_1 h(x)$.

Using the fact that $h(x)$ is a primitive polynomial (by Lemma II.1, since it is a product of primitive polynomials), and the assumption that $b(x)$ is primitive, we then get that $c_1 = 1$. We can then factor $b(x) = h(x)$ in deterministic polynomial time, using the algorithm of Lenstra, Lenstra and Lovász (Theorem II.2), to obtain the factors $p_i(x)$ of $h(x)$. Thus, from $a(x)$ and $b(x)$, we can obtain the polynomials $g(x) = a(x)$ and $h(x) = b(x)$ as well as the distinct irreducible factors $p_i(x)$ of $p(x)$ in time polynomial in the description length of the input. It only

¹For readers familiar with [1], we point out that the main results in that paper sought to isolate instances where this case applies.

remains to determine the exponents d_i , which we now do using the uniqueness of the partial fraction expansion of $g(x)/h(x)$ with $p_i(x)$ as the denominators.

We note first that the representation in eq. (5) implies the existence of such an expansion:

$$\frac{a(x)}{b(x)} = \frac{g(x)}{h(x)} = \sum_{i=1}^k \frac{d_i p_i'(x)}{p_i(x)}. \quad (6)$$

Let $n \geq k - 1$ be the degree of $g(x)$.² Since we already know $g(x)$ and the $p_i(x)$, we can match coefficients of powers of x in the definition of $g(x)$ in eq. (5) to set up $n + 1 \geq k$ equations to determine the d_i . Note that all the coefficients in this system of equations have description lengths polynomial in the description length of the $p_i(x)$, which themselves have length polynomial in the description length of $a(x)$ and $b(x)$. Since Gaussian elimination can be carried out in time polynomial in the length of the input (see, e.g., [15]), this system of equations can be solved in polynomial time. Further, the definition of $g(x)$ implies the existence of at least one solution to this system, while the expansion in eq. (6) and Theorem II.4 on the uniqueness of partial fraction expansions implies that this solution is unique. Thus, the unique solution of the system yields the d_i , and we can then determine the product $\prod_{i=1}^k p_i(x)^{d_i} = p(x)/c$. The constant c can then be determined by comparing the known coefficient of $p(x)$ with the corresponding coefficient of the product. ■

Remark III.1. As observed by an anonymous referee, it is in fact possible to avoid the use of the Lenstra, Lenstra and Lovász algorithm in the procedure for determining $p(x)$ from $a(x)$ and $b(x)$, as we now explain. Let $t(x)$ be the unique polynomial such that

$$p'(x) = a(x)t(x) \quad \text{and} \quad p(x) = b(x)t(x). \quad (7)$$

Note that the conditions in eq. (7) yield a system of linear equations for the coefficients of $t(x)$ (namely, that obtained by equating the coefficients of the polynomials $a(x)t(x)$ and $(b(x)t(x))'$). Now, the proof of the lemma, in particular the discussion in the two paragraphs following eq. (5), shows that such a $t(x)$ must satisfy $t(x) := c \prod_{i=1}^k p_i(x)^{d_i-1}$, where the p_i are the distinct primitive irreducible factors of $p(x)$ with exponents d_i , so that $p(x) = c \prod_{i=1}^k p_i(x)^{d_i}$. Note also that the rest of the proof establishes that the $p_i(x)$ and the d_i are uniquely determined given $a(x)$ and $b(x)$. Together with the one known coefficient of $p(x)$, this implies that $t(x)$ itself is uniquely determined given $a(x)$ and $b(x)$, and hence is the unique solution of the system of linear equations given by eq. (7). Thus solving this system allows us to efficiently determine $t(x)$, and hence also $p(x)$.

Combining the above lemma with Theorem II.5, we get the following corollary.

Corollary III.2. *Let $p(x) \in \mathbb{Q}[x]$ be an unknown polynomial of degree n . Then given $2n + 1$ distinct evaluation points $y_1, y_2, \dots, y_{2n+1} \in \mathbb{Q}$, the corresponding values $\frac{p'(y_i)}{p(y_i)}$, and one non-zero coefficient of $p(x)$, we can uniquely determine $p(x)$ in time polynomial in the description length of the input.*

Proof: We first use Theorem II.5 (and the algorithm outlined in the Remark following it) to obtain polynomials $a(x)$ and $b(x)$ in $\mathbb{Q}[x]$ such that $\frac{p'(x)}{p(x)} = \frac{a(x)}{b(x)}$, $a(x)$ and $b(x)$ are co-prime, and $b(x)$ is primitive. Note then that these $a(x)$ and $b(x)$ satisfy the hypotheses of Lemma III.1, so we can now apply the algorithm in the lemma to obtain $p(x)$. ■

IV. PROOFS OF HARDNESS RESULTS: THE HARD CORE MODEL

We now show how Corollary III.2 can be used to prove the hardness results listed in Section I-B. All of these proofs follow a similar template, in which the first step is to use a direct application of the Corollary to prove a somewhat weaker hardness result for the problem in which the corresponding model parameters are allowed to be part of the input. Proving the results when the parameters are fixed then requires some more combinatorial

²This follows from the facts that $\deg(h(x)) \geq k$, and $\deg(h(x)) - \deg(g(x)) = \deg(p(x)) - \deg(p'(x)) = 1$.

work which is mostly independent of the first step. This template is similar to the strategy used by Sinclair and Srivastava [1], and the main difference between that paper and this appears in the implementation of the first step. As in [1], the first step starts with an attempt to use a rational interpolation procedure to compute the corresponding partition function using a hypothetical algorithm for computing the mean observable in question. The hardness of computing the mean observable would then follow if the partition function itself is known to be #P-hard (as it is in all the cases we consider). In [1], Lee-Yang type theorems were used to guarantee that the rational interpolation indeed succeeds in recovering the partition function. In contrast, in this paper we accept the possibility that cancellations may not allow the rational interpolation procedure to recover the partition function. However, we then invoke Corollary III.2 to ensure that the partition function can nevertheless be obtained after some further processing of the results of the rational interpolation.

In this section, we demonstrate the template by instantiating it for the hard core model (Theorem I.3). The details for the other models are similar and are deferred to the next section. For the hard core model, our proof involves a reduction from the well-known #P-hard problem of counting independent sets in bounded degree graphs. Concretely, we use the following hardness result due to Greenhill [16].

Theorem IV.1 ([16, Theorem 3.1]). *The problem of counting independent sets in graphs of degree at most 3 is #P-hard.*

Proof of Theorem I.3: Let G be any graph with maximum degree at most 3. We assume that we have an algorithm \mathcal{A} that, given a graph K on n vertices of maximum degree at most 4, computes in time polynomial in n the average size $M_H(K, \lambda)$ of an independent set in the hard core model with vertex activity λ on K (recall that λ is a fixed parameter, and not an input to \mathcal{A}). We then use this algorithm to compute the number of independent sets in G , thus completing the proof via Theorem IV.1.

Our starting point is the earlier observation that the average size of an independent set $M_H(G, x)$ can be written in terms of the derivative of the logarithm of the partition function:

$$M_H(G, x) = x \frac{d}{dx} \log Z_H(G, x) = x \frac{Z'_H(G, x)}{Z_H(G, x)}. \quad (8)$$

Let n be the number of vertices in G . Suppose for the moment that we have the values $M_H(G, \lambda_i)$ for $2n + 1$ distinct, positive rational λ_i , $1 \leq i \leq 2n + 1$. Note that eq. (8) implies that if the λ_i have bit length polynomial in n , then so do the values $M_H(G, \lambda_i)/\lambda_i = \frac{Z'_H(G, \lambda_i)}{Z_H(G, \lambda_i)}$. Since all these λ_i are positive, the latter function is well defined (i.e., has non-zero denominator) at all these points, and hence we can use Corollary III.2 to compute the polynomial $Z_H(G, x)$, and hence also the value $Z_G(x, 1)$. This shows that if we could somehow use the hypothetical algorithm \mathcal{A} to evaluate $M_H(G, \cdot)$ at multiple efficiently computable values of λ , we would be done.

To do this, we modify the input graph G in a manner almost identical to that used in [1], whose notation we closely follow. Specifically, let $P(k)$ denote the path on k nodes. Let p_k^+ denote the partition function $Z_H(P(k), \lambda)$ restricted to those independent sets which occupy the “leftmost” node of $P(k)$. Similarly, let p_k^- denote the partition function $Z_H(P(k), \lambda)$ restricted to those independent sets which do *not* occupy the “leftmost” node of $P(k)$. Define $r_k = \frac{p_k^+}{p_k^-}$. As we shall see later, the p_k are polynomials in λ which can be computed efficiently as long as $k = \text{poly}(n)$.

Now, consider the graph $G(k)$ obtained from G by attaching to each node of G a separate copy of $P(k)$ so that any node v in G is connected by an edge to the left-most vertex of its corresponding copy of $P(k)$. Consider now an independent set I in $G(k)$, and denote by I' its projection on the vertices of G . It is easy to see that the total contribution to the partition function $Z_H(G(k), \lambda)$ of all independent sets in $G(k)$ with a *fixed* projection I' on G is given by $(p_{k+1}^+)^{|I'|} (p_{k+1}^-)^{n-|I'|}$. This in turn implies that

$$Z_H(G(k), \lambda) = (p_{k+1}^-)^n Z_H(G, r_{k+1}).$$

Viewing r_k, p_k^+ and p_k^- as formal functions of λ and denoting their formal derivatives with respect to λ by the

fluxion notation \dot{r}_k , \dot{p}_k^+ and \dot{p}_k^- , we then have

$$\begin{aligned}
M_H(G(k), \lambda) &= \frac{n\lambda\dot{p}_{k+1}^-}{\dot{p}_{k+1}^-} + \lambda \frac{d}{d\lambda} \log Z_H(G, r_{k+1}) \\
&= \frac{n\lambda\dot{p}_{k+1}^-}{\dot{p}_{k+1}^-} + \lambda r_{k+1} \frac{\partial}{\partial r_{k+1}} \log Z_H(G, r_{k+1}) \\
&= \frac{n\lambda\dot{p}_{k+1}^-}{\dot{p}_{k+1}^-} + \frac{\lambda r_{k+1}}{r_{k+1}} M_H(G, r_{k+1}), \tag{9}
\end{aligned}$$

where in the last line we use the definition of M_H . Thus, all we need to show is that for $2 \leq i \leq 2n + 2$, the r_i are distinct, $r_i \neq 0$, and all the values \dot{p}_i^+ , \dot{p}_i^- and \dot{r}_i are efficiently computable. This is sufficient because we can then consider the graphs $G(k)$ for $k = 1, 2, \dots, 2n + 1$ and then use eq. (9) and the numbers $M_H(G(k), \lambda)$ output by the algorithm \mathcal{A} to compute $M_H(G, x)$ at $2n + 1$ distinct positive rational values of x given by the r_{k+1} . As discussed above, this is sufficient to complete the reduction.

To establish the requisite properties of the p_i and the r_i , we now consider their computation in detail. In particular, we have $p_1^+ = r_1 = \lambda$, $p_1^- = 1$, while $\dot{p}_1^+ = \dot{r}_1 = 1$, and $\dot{p}_1^- = 0$. We then have the following recurrences:

$$\dot{p}_{k+1}^+ = \lambda \dot{p}_k^-, \tag{10}$$

$$\dot{p}_{k+1}^- = \dot{p}_k^- + \dot{p}_k^+, \tag{11}$$

$$r_{k+1} = \frac{\lambda}{1 + r_k}, \tag{12}$$

$$\dot{p}_{k+1}^+ = \lambda \dot{p}_k^- + \dot{p}_k^-, \tag{13}$$

$$\dot{p}_{k+1}^- = \dot{p}_k^- + \dot{p}_k^+, \text{ and} \tag{14}$$

$$r_{k+1} = \frac{\dot{p}_{k+1}^+ \dot{p}_{k+1}^- - \dot{p}_{k+1}^+ \dot{p}_{k+1}^-}{(\dot{p}_{k+1}^-)^2}.$$

Note that eqs. (10), (11), (13) and (14) show that the p_k , r_k and their derivatives can be computed in time polynomial in k , so that if $k = \text{poly}(n)$, they are all efficiently computable. The fact that $\dot{r}_k \neq 0$ follows from eq. (12) and a simple induction. It remains to prove that the r_k are distinct. To see this, note that we have $r_{k+1} = f(r_k) := \frac{\lambda}{1+r_k}$. When $\lambda > 0$, f has the following two properties, which can be verified by a direct computation:

- 1) f is strictly decreasing on the positive real line, and has a unique positive fixed point x^* which is smaller than λ .
- 2) $f \circ f$ is strictly increasing and has the unique positive fixed point x^* . Further $f(f(x)) > x$ when $x < x^*$ and $f(f(x)) < x$ when $x > x^*$.

The above two properties imply that r_1, r_3, r_5, \dots form a strictly decreasing sequence all of whose elements are strictly larger than x^* , while r_2, r_4, r_6 form a strictly increasing sequence all of whose elements are strictly smaller than x^* . This implies that all the r_k are distinct, and completes the proof. (Note that since G was of maximum degree at most 3, the maximum degree of the $G(k)$ is at most 4.) ■

V. PROOFS OF HARDNESS RESULTS: THE ISING AND MONOMER-DIMER MODELS

In this section, we instantiate the template described in Section IV to prove Theorems I.1, I.2 and I.4. The combinatorial simulations required in the proofs of Theorems I.1 and I.4 are essentially identical to those in [1], so we only point out the minor differences. However, the proof of Theorem I.2 on the hardness of mean energy of the Ising model requires a different construction due to Dyer and Greenhill [17]. We proceed to describe that reduction first.

A. Proof of Theorem I.2

We reduce from the problem of computing the Ising partition function $Z_I(G, \beta, 1)$, which is known to be #P-hard via the following theorem.

Theorem V.1 ([17, Theorem 5.1]). *Fix any positive rational $\beta \neq 1$. Then the problem of computing the partition function $Z_I(G, \beta, 1)$ of the Ising model on graphs of maximum degree at least 3 is #P-hard.*

We are now ready to prove Theorem I.2.

Proof of Theorem I.2: We follow the same template as that used in the proof of Theorem I.3. Let G be any graph with maximum degree at most 3. We assume that we have an algorithm \mathcal{A} that, given a graph K with n vertices, m edges and maximum degree at most 4, computes in time polynomial in n and m the mean energy $M_{IE}(K, \beta, 1) =: M_{IE}(K, \beta)$ of the Ising model with edge activity β and vertex activity 1 on K (recall that β is a fixed parameter, and not an input to \mathcal{A}). We then use this algorithm to compute the partition function $Z_I(G, \beta, 1)$, thus completing the proof via Theorem V.1.

Our starting point again is the observation that the mean energy $M_{IE}(G, \beta)$ can be written in terms of the derivative of the logarithm of the partition function:

$$M_{IE}(G, x) = x \frac{d}{dx} \log Z_I(G, x, 1) = x \frac{Z'_I(G, x, 1)}{Z_I(G, x, 1)}. \quad (15)$$

Let n be the number of vertices in G . Suppose for the moment that we have the values $M_{IE}(G, \beta_i)$ for $2m + 1$ distinct, positive rational β_i , $1 \leq i \leq 2m + 1$. Note that eq. (15) implies that if the β_i have bit length polynomial in n , then so do the values $M_{IE}(G, \beta_i)/\beta_i = \frac{Z'_I(G, \beta_i, 1)}{Z_I(G, \beta_i, 1)}$. Since all these β_i are positive, the latter function is well defined (i.e., has non-zero denominator) at all these points, and hence we can use Corollary III.2 to compute the polynomial $Z_I(G, x, 1)$, and thus also the value $Z_I(G, \beta, 1)$ (since the degree of the latter as a polynomial in β is at most m). This shows that if we could somehow use the hypothetical algorithm \mathcal{A} to evaluate $M_{IE}(G, \cdot)$ at multiple efficiently computable values of β , we would be done.

To do this, we use the stretching operation defined by Dyer and Greenhill [17]. Formally, given a graph G and a non-negative integer k , the k -stretch $S_k G$ of G is defined as the graph obtained by replacing each edge of G with a path on $k + 1$ nodes with the end points of the path at the endpoints of the original edge. Thus, for example, $S_1 G = G$, and $S_2 G$ is the graph where each edge of G is replaced by a path of length two. Specializing Corollary 2.1 of Dyer and Greenhill [17] to our case, we get that

$$Z_I(S_k G, \beta, 1) = \alpha_k^m Z_I(G, r_k, 1), \quad (16)$$

where m is the number of edges in G , $\alpha_1 = 1$, $\gamma_1 = \beta$, $r_k := \frac{\gamma_k}{\alpha_k}$ and for larger k these quantities satisfy the recurrences

$$\begin{aligned} \alpha_{k+1} &= \alpha_k + \beta \gamma_k, \\ \gamma_{k+1} &= \beta \alpha_k + \gamma_k, \text{ and} \\ r_{k+1} &= \frac{\beta + r_k}{1 + \beta r_k}. \end{aligned} \quad (17)$$

Viewing the α_k , γ_k and r_k as functions of β , we can also write recurrences for their derivatives (denoted using the fluxion notation) with respect to β :

$$\begin{aligned} \dot{\alpha}_{k+1} &= \dot{\alpha}_k + \beta \dot{\gamma}_k + \gamma_k, \\ \dot{\gamma}_{k+1} &= \beta \dot{\alpha}_k + \alpha_k + \dot{\gamma}_k, \text{ and} \\ \dot{r}_{k+1} &= \frac{\gamma_{k+1} \dot{\alpha}_{k+1} - \alpha_{k+1} \dot{\gamma}_{k+1}}{\alpha_{k+1}^2}. \end{aligned}$$

These recurrences show that as long as k is polynomial in m , the description length of all these quantities is also polynomially bounded in the input size.

As before, we can now use eq. (16) to obtain an expression for $M_{IE}(S_k G, \beta)$ in terms of the above quantities:

$$M_{IE}(S_k G, \beta) = \frac{m\beta\dot{\alpha}_k}{\alpha_k} + \frac{\beta\dot{r}_k}{r_k} M_{IE}(G, r_k).$$

Thus, again as before, if the r_k are distinct for $k = 1, 2, \dots, 2m + 1$, and the \dot{r}_k are non-zero, we would be done (since we could use the values of $M_{IE}(S_k G, \beta)$ obtained from the algorithm to obtain the values of $M_{IE}(G, \beta_k)$).

The fact that $\dot{r}_k \neq 0$ again follows by an easy induction using eq. (17). To show that r_k are distinct, we use somewhat different strategies for the ferromagnetic case ($\beta < 1$) and the anti-ferromagnetic case ($\beta > 1$). In the former case, an induction proves that the r_k form a strictly increasing sequence, which guarantees distinctness. In the anti-ferromagnetic case, the argument is very similar to that made in the case of the hard core model: it again turns out that r_1, r_3, r_5, \dots form a strictly decreasing sequence that always remains larger than 1, while r_2, r_4, r_6, \dots form a strictly increasing sequence that always remains smaller than 1. As before, this guarantees distinctness of the r_k and completes the proof. (Note that since G was of maximum degree at most 3, so are the $S_k G$ for all k .) ■

B. Proofs of Theorems I.1 and I.4

We now proceed to sketch the proofs of Theorems I.1 and I.4. Since the combinatorial parts of these proofs are virtually identical to some calculations that appeared in [1], we merely describe the places where the arguments differ.

Proof of Theorem I.1: We again reduce from the problem of computing $Z_I(G, \beta, 1)$ described in Theorem V.1. Let G be any graph with maximum degree at most 3 and n vertices. We assume that we have an algorithm \mathcal{A} that, given a graph K with n vertices and maximum degree at most 4, computes in time polynomial in n the mean magnetization $M_{IM}(K, \beta, \lambda)$ of the anti-ferromagnetic Ising model with edge activity $\beta > 1$ and vertex activity λ on K (recall that β and λ are fixed parameters, and not inputs to \mathcal{A}). We then use this algorithm to compute the partition function $Z_I(G, \beta, 1)$, thus completing the proof via Theorem V.1.

Proceeding exactly as in the proof of Theorem I.3 for the hard core model (except for the use of eq. (2) in place of eq. (1)), we then conclude using Corollary III.2 that we only need to show how to evaluate $M_{IM}(G, \beta, x)$ for some $2n + 1$ distinct, efficiently computable values of x . To do this we again consider the graphs $G(k)$, and refer to the analysis of the Ising model on these graphs in Appendix C of [1], in particular eqs. (21)–(23), (25)–(27) and (30) of that paper, which show that for $k > 1$ (using our notation)

$$M_{IM}(G(k), \beta, \lambda) = \frac{n\lambda p_{k+1}}{p_{k+1}} + \frac{\lambda r_{k+1}}{r_{k+1}} M_{IM}(G(k), \beta, r_{k+1}),$$

where the quantities p_{k+1}, r_{k+1} and their derivatives with respect to λ (denoted in the fluxion notation) are efficiently computable as long as k is polynomially bounded. Using the same notation as in that paper, it only remains to show that the \dot{r}_k remain non-zero and the r_k remain distinct even in the case $\beta > 1$ (the rest of the analysis in that paper is for $\beta < 1$, i.e., the ferromagnetic case). The fact that the \dot{r}_k are non-zero follows from an induction using eq. (23) of that paper, while the distinctness follows by an argument exactly analogous to that used in the case of the hard core model, except that the cases $\lambda > 1$ and $\lambda < 1$ need to be considered separately. Note also that, as expected, the distinctness argument does not work when $\lambda = 1$ (in this case all the r_k are 1), since the problem is trivially solvable in polynomial time in this special case. Note also that since the maximum degree of G is at most 3, the maximum degree of the $G(k)$ is at most 4, which completes the proof. ■

We now prove the hardness result for the monomer-dimer model. For this model, no changes from the treatment in [1] are required in the combinatorial part of the argument, so we only describe the steps leading to it. We use the following hardness result for counting matchings to do the reduction.

Theorem V.2 ([18, Table 1.1]). *The problem of counting matchings (equivalently, computing $Z_D(G, 1)$) in graphs of maximum degree at most 4 is #P-hard.*

Proof of Theorem I.4: Let G be any graph with maximum degree at most 4 and n vertices. We assume that we have an algorithm \mathcal{A} that, given a graph K on n vertices of maximum degree at most 5, computes in time polynomial in n the average number of monomers $M_D(K, \lambda)$ of the monomer-dimer model with vertex activity λ on K (recall that λ is fixed parameter, and not an input to \mathcal{A}). We then use this algorithm to compute the partition function $Z_D(G, 1)$, thus completing the proof via Theorem V.2.

Again, we proceed in a manner analogous to that in the case of the hard core model, and conclude using Corollary III.2 that it is sufficient to be able to evaluate $M_D(G, x)$ for $2n + 1$ distinct, efficiently computable values of x using the algorithm \mathcal{A} . However, the proof of Theorem 4 in Appendix C of [1] gives a method for doing exactly this, by running the algorithm \mathcal{A} on the graphs $G(k)$ constructed above. This allows us to complete the reduction, once we observe that since G was of maximum degree at most 4, the $G(k)$ are of maximum degree at most 5. ■

ACKNOWLEDGMENTS

We thank an anonymous referee for observing that an appeal to the polynomial factoring algorithm of Lenstra, Lenstra and Lovász is not necessary in the proof of our main technical lemma.

LJS and PS were supported in part by NSF grant CCF-1319745. AS was supported by in part by NSF grant CCF-1420934 and the Simons Institute for the Theory of Computing.

REFERENCES

- [1] A. Sinclair and P. Srivastava, “Lee–Yang theorems and the complexity of computing averages,” *Comm. Math. Phys.*, vol. 329, no. 3, pp. 827–858, 2014. A preliminary version appeared in *Proc. 45th ACM Symp. Theory Comput. (STOC) 2013*. [Online]. Available: <https://dx.doi.org/10.1007/s00220-014-2036-7>
- [2] E. Ising, “Beitrag zur Theorie des Ferromagnetismus,” *Z. Phys.*, vol. 31, pp. 253–258, Feb. 1925.
- [3] O. J. Heilmann and E. H. Lieb, “Theory of monomer-dimer systems,” *Comm. Math. Phys.*, vol. 25, no. 3, pp. 190–232, 1972. [Online]. Available: <https://dx.doi.org/10.1007/BF01877590>
- [4] M. Jerrum and A. Sinclair, “Polynomial-time approximation algorithms for the Ising model,” *SIAM J. Comput.*, vol. 22, no. 5, pp. 1087–1116, 1993.
- [5] D. Randall and D. Wilson, “Sampling spin configurations of an Ising system,” in *Proc. 10th ACM-SIAM Symp. Discret. Algorithms (SODA)*, ser. SODA ’99. Philadelphia, PA, USA: SIAM, 1999, pp. 959–960. [Online]. Available: <http://dl.acm.org/citation.cfm?id=314500.314945>
- [6] M. Jerrum and A. Sinclair, “Approximating the permanent,” *SIAM J. Comput.*, vol. 18, no. 6, pp. 1149–1178, Dec. 1989. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/0218077>
- [7] T. D. Lee and C. N. Yang, “Statistical theory of equations of state and phase transitions. II. Lattice gas and Ising model,” *Phys. Rev.*, vol. 87, no. 3, pp. 410–419, Aug. 1952.
- [8] E. Horowitz, “Algorithms for partial fraction decomposition and rational function integration,” in *Proc. 2nd ACM Symp. Symbolic and Algebraic Manipulation*, ser. SYMSAC ’71. New York, NY, USA: ACM, 1971, pp. 441–457. [Online]. Available: <http://doi.acm.org/10.1145/800204.806314>
- [9] R. G. Tobey, “Algorithms for anti-differentiation of rational functions,” Ph.D. dissertation, Harvard University, 1967.
- [10] J. Borcea and P. Brändén, “Pólya-Schur master theorems for circular domains and their boundaries,” *Ann. Math.*, vol. 170, no. 1, pp. 465–492, Jul. 2009. [Online]. Available: <https://dx.doi.org/10.4007/annals.2009.170.465>

- [11] —, “The Lee-Yang and Pólya-Schur programs. I. Linear operators preserving stability,” *Invent. Math.*, vol. 177, no. 3, pp. 541–569, 2009. [Online]. Available: <https://dx.doi.org/10.1007/s00222-009-0189-3>
- [12] I. N. Herstein, *Topics in Algebra*, 2nd ed. John Wiley and Sons, 1999.
- [13] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász, “Factoring polynomials with rational coefficients,” *Math. Ann.*, vol. 261, no. 4, pp. 515–534, Dec. 1982.
- [14] N. Macon and D. E. Dupree, “Existence and uniqueness of interpolating rational functions,” *Am. Math. Mon.*, vol. 69, no. 8, pp. 751–759, Oct. 1962. [Online]. Available: <http://www.jstor.org/stable/2310771>
- [15] J. Edmonds, “Systems of distinct representatives and linear algebra,” *J. Res. Natl. Bur. Stand. B Math. & Math. Phys.*, vol. 71B, no. 4, pp. 241–245, 1967.
- [16] C. Greenhill, “The complexity of counting colourings and independent sets in sparse graphs and hypergraphs,” *Comput. Complexity*, vol. 9, no. 1, pp. 52–72, 2000. [Online]. Available: <https://dx.doi.org/10.1007/PL00001601>
- [17] M. E. Dyer and C. S. Greenhill, “The complexity of counting graph homomorphisms,” *Random Struct. Algorithms*, vol. 17, no. 3-4, pp. 260–289, 2000.
- [18] S. P. Vadhan, “The complexity of counting in sparse, regular, and planar graphs,” *SIAM J. Comput.*, vol. 31, no. 2, pp. 398–427, Jan. 2001. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/S0097539797321602>

APPENDIX

A. Uniqueness of partial fraction expansions

Proof of Theorem II.4: Suppose for contradiction that there are two sequences of polynomials $(q_i(x))_{i=1}^k$ and $(r_i(x))_{i=1}^k$ satisfying the conclusions of the theorem. Then, by permuting indices if necessary, we can assume that $q_1(x) \neq r_1(x)$. We now show that this leads to a contradiction.

By using the expansion in eq. (3) for $r_i(x)$ and $q_i(x)$ and clearing fractions, we get that

$$a(x) = q_1(x)S_1(x) + \sum_{i=2}^k q_i(x)S_i(x), \text{ and}$$

$$a(x) = r_1(x)S_1(x) + \sum_{i=2}^k r_i(x)S_i(x),$$

where $S_i(x) := \prod_{j \neq i} p_j(x)$. Note that p_1 divides S_j for $j \neq 1$, but does not divide S_1 since the p_i are distinct irreducible polynomials.

Subtracting one equation from the other and reducing the resulting equation modulo $p_1(x)$, we therefore get that $p_1(x)$ must divide $r_1(x) - q_1(x)$ (since it cannot divide $S_1(x)$). However, this is impossible since $p_1(x)$ is irreducible and $r_1(x) - q_1(x)$ is a non-zero polynomial of degree less than that of $p_1(x)$. ■