THESIS

EXTREMAL VALUES OF THE OCCUPANCY FRACTION FOR THE ANTIFERROMAGNETIC ISING MODEL

Submitted by

Olivia LeBlanc

Department of Computer Science

In partial fulfillment of the requirements

For the Degree of Master of Science

Colorado State University

Fort Collins, Colorado

Summer 2025

Master's Committee:

Advisor: Ewan Davies

Sanjay Rajopadhye
Vinayak Prabhu
Maria Gillespie

ABSTRACT

EXTREMAL VALUES OF THE OCCUPANCY FRACTION FOR THE ANTIFERROMAGNETIC ISING MODEL

The Ising model is a mathematical model of magnetism which is frequently studied in statistical physics and computer science. For the antiferromagnetic version of the model, there is known to be a computational threshold in the complexity of sampling from the model at given magnetization on Δ-regular graphs. The value of this threshold can be determined by minimizing the occupancy fraction of the model, but prior to this paper an explicit formula was not known. This work solves the minimization problem for the majority of the relevant parameter space in the case $\Delta = 3$, determining the value of this threshold. Our methods also yield results on the minimization and maximization problems in other areas of the parameter space, painting a more complete picture of the occupancy fraction's behavior in 3-regular graphs.

TABLE OF CONTENTS

Chapter 1: Introduction

The Ising model is a mathematical framework used to study magnetism and phase transitions. It is an example of a two-state spin system, meaning it consists of a series of "spins" which can take one of two values (+ or -) laid out in a graph structure. The model defines a probability distribution over different spin configurations, taking into account interactions between spins and the presence of an external magnetic field. Only interactions between adjacent spins are captured, allowing for the study of how local interactions affect macroscopic properties. While this was originally designed to model magnetism, the generality of the model can be applied to a variety of other phenomena. It has since been used in areas such as genetics [1], social sciences [2], and neuroscience [3].

We study the antiferromagnetic Ising model, whose local interactions favor opposing spins. A major algorithmic challenge associated with this model is computation of the *partition function*, which is a weighted sum over all possible configurations under the model. It takes the role of a normalization factor, ensuring the probabilities given by the model sum to one. Exact computation of the partition function is #P-hard in general [4], and naive computation takes exponential time. Due to this, there has been a focus on finding an approximate answer in polynomial time. However, there are known to be parameter regions for which even approximating the partition function (to within an arbitrary constant factor) is an NP-hard problem [5, 6].

There is a nearly complete characterization of the complexity of this problem for regular graphs, such that at any point in the parameter space (except on the phase transition boundary), either an efficient approximation algorithm has been created [7, 8] or it has been proven that one cannot exist [5, 6] (assuming RP$\neq$NP). This phase transition boundary acts as a threshold, such that all parameter values for which there exists an efficient approximation algorithm lie on one side, and all parameter values for which there is no efficient approximation algorithm lie on the other. The behavior of the problem exactly at this threshold remains unknown.

Our focus is the related problem of approximating the partition function at a fixed magnetization on regular graphs, for which a similar threshold has been shown to exist. This threshold was established by Davies and Perkins [9], who again prove the existence of a boundary such that an algorithm for polynomial-time approximation exists on one side, but not on the other (unless RP=NP). Interestingly, the authors do not include a formula for the location of the threshold.

The threshold for fixed-magnetization approximation was formulated for the case of Δ-regular graphs in terms of the *occupancy fraction* α. For each $\Delta \geq 3$ there is a critical value $\alpha_c(\Delta)$ such that approximation can be done in polynomial time when $\alpha < \alpha_c$ and unless RP=NP, it cannot be done when $\alpha > \alpha_c$ [9]. Davies and Perkins establish the critical value $\alpha_c(\Delta)$ to be the result of minimizing the occupancy fraction over Δ-regular graphs, but finding an explicit value for this was left as a conjecture.

1.1 Contributions

In this thesis, we address a fundamental extremal problem within the framework of the anti-ferromagnetic Ising model: determining extrema of the occupancy fraction, in terms of both minimization and maximization. The minimization problem is of key interest, as its solution determines a computational threshold in approximately sampling from the model.

The primary contribution of this thesis is the solution of the minimization problem across almost the entire relevant parameter space for the case of 3-regular graphs. We achieve this through tackling a linear programming relaxation of the graph-based optimization problem. This allows us to establish the value of the computational threshold for most of the relevant parameter values and offer insights into the behavior of this problem in other areas of the parameter space.

1.2 Acknowledgment

This thesis includes material co-authored with Ewan Davies [10]. This work was supported in part by NSF grant CCF-2309707.

Chapter 2: Background

This section will formally define the Ising model and its properties that are relevant to our problem.

2.1 The Ising Model

The Ising model defines a probability distribution over spin configurations in a graph. We use σ to represent an assignment of a spin in $\{+,-\}$ to each vertex of a graph.

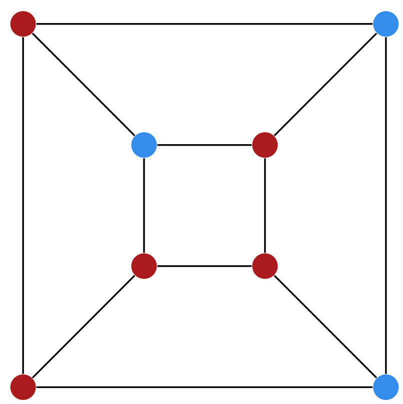

Figure 2.1: An example of an Ising spin assignment on a 3-regular graph. Vertices assigned $+$ or $-$ are colored red or blue respectively.

We will first define two key properties of a spin assignment, as they will arise throughout our discussion of the model.

- $m_G(\sigma)$ is the number of edges of G whose endpoints get the same spin under σ (we call such edges *monochromatic*)

- $n_G(+,\sigma)$ is the number of vertices which receive the spin $+$ under σ

For an edge activity parameter B and external field λ, the partition function $Z_G(B, \lambda)$ of the Ising model on a graph G is given by

$$Z_G(B, \lambda) := \sum_{\sigma: V(G) \to \{+,-\}} B^{m_G(\sigma)} \lambda^{n_G(+,\sigma)}$$

We define the Ising measure $\mu_{G,B,\lambda}$ on spin assignments $\sigma : V(G) \to \{+, -\}$ via

$$\mu_{G,B,\lambda}(\sigma) = \frac{B^{m_G(\sigma)} \lambda^{n_G(+,\sigma)}}{Z_G(B, \lambda)}.$$

The model is antiferromagnetic if $B \in (0, 1)$, and as can be seen in the definition of the Ising measure, this regime assigns higher weights to assignments with fewer monochromatic edges (that is, it "prefers" edges between opposing spins). Without loss of generality we consider $0 < \lambda \leq 1$ because the model is symmetric under swapping the spins and taking $\lambda \mapsto 1/\lambda$.

The magnetization of a spin assignment σ is $M(\sigma) = \sum_{v \in V} \sigma_v$, and given a fixed graph, fixing the magnetization is equivalent to fixing $n_G(+, \sigma)$. If we view the partition function as a polynomial in λ we have

$$Z_G(B, \lambda) = \sum_{k=0}^{|V|} c_k \lambda^k$$

for some coefficients $c_k = c_k(G, B)$ which depend on the graph and the edge activity parameter B. We write $\mathcal{B}_k(G)$ for the subset of spin assignments σ such that $n_G(+, \sigma)$ is exactly k, giving

$$c_k = \sum_{\sigma \in \mathcal{B}_k(G)} B^{m_G(\sigma)}.$$

For each given B, we are interested in establishing a threshold in k for the computational problem of approximating $c_k(G, B)$ in the class of Δ-regular graphs. To discuss the threshold, we introduce the *occupancy fraction* α, which is defined by:

$$\alpha_G(B, \lambda) = \frac{\lambda}{|V(G)|} \frac{\partial}{\partial \lambda} \ln Z_G(B, \lambda).$$

Lemma 1. *The occupancy fraction $\alpha_G(B,\lambda)$ is the expected fraction of vertices which receive spin $+$ in a sample from $\mu_{G,B,\lambda}$.*

Proof.

$$\alpha_G(B,\lambda) = \frac{\lambda}{|V(G)|}\frac{\partial}{\partial\lambda}\ln Z_G(B,\lambda) = \frac{1}{|V(G)|}\frac{\lambda\frac{\partial}{\partial\lambda}Z_G(B,\lambda)}{Z_G(B,\lambda)}$$

$$= \frac{1}{|V(G)|}\frac{\sum_\sigma n_G(+,\sigma)B^{m_G(\sigma)}\lambda^{n_G(+,\sigma)}}{Z_G(B,\lambda)} = \frac{1}{|V(G)|}\sum_\sigma n_G(+,\sigma)\mu_{G,B,\lambda}(\sigma)$$

$$= \frac{1}{|V(G)|}\mathbb{E}\big[|\{\sigma(v) = + \mid \sigma \sim \mu_{G,B,\lambda}\}|\big]$$

□

It has been established in some generality that for two-spin antiferromagnetic systems, minimizing α over a class of graphs of interest determines the computational threshold [9]. Let $\mathcal{G}_\Delta$ be the class of Δ-regular graphs and let

$$\alpha_{\inf}(\Delta,B,\lambda) = \inf_{G\in\mathcal{G}_\Delta}\alpha_G(B,\lambda).$$

In order to define the computational threshold, it is also necessary to define the critical parameter value λ_c. Given $B_c(\Delta) = \frac{\Delta-2}{\Delta}$, the explicit function giving λ_c is:

$$\lambda_c(\Delta,B) = \frac{1-\sqrt{r/s}}{1+\sqrt{r/s}}\left(\frac{1+\sqrt{rs}}{1-\sqrt{rs}}\right)^{\frac{B_c-1}{B_c+1}}, \tag{2.1}$$

where $B_c = B_c(\Delta)$ and

$$r = \frac{B_c - B}{B_c + B}, \qquad s = \frac{1-B}{1+B}.$$

See chapter 4 of the book [11] for more details.

The result of Davies and Perkins relevant to the antiferromagnetic Ising model in regular graphs establishes a computational threshold α_c as follows. Note that a *fully polynomial-time randomized approximation scheme* or FPRAS for a quantity $Q(x)$ is a family of randomized algorithms for each $\varepsilon > 0$ that compute with probability at least 3/4 a value $\hat{Q}(x)$ such that $e^{-\varepsilon} \leq \hat{Q}(x)/Q(x) \leq e^{\varepsilon}$, in time polynomial in the input size (e.g. the number of vertices of the graph represented by x) and $1/\varepsilon$.

Theorem 1 (Davies and Perkins [9, Thm. 3])**.** *For $\Delta \geq 3$ and $0 < B < B_c(\Delta)$, we write $\alpha_c = \alpha_{\inf}(\Delta, B, \lambda_c(\Delta, B))$. Then α_c is a computational threshold in the following sense.*

(a) *For every $\alpha < \alpha_c$ there is an FPRAS for $c_{\lfloor \alpha n \rfloor}(G, B)$ for all n-vertex Δ-regular graphs G.*

(b) *Unless NP=RP, for every $\alpha \in (\alpha_c, 1/2]$ there is no FPRAS for $c_{\lfloor \alpha n \rfloor}(G, B)$ for n-vertex Δ-regular graphs G.*

Given this theorem, determining the computational threshold requires finding an explicit formula for α_c. Davies and Perkins conjectured that for $\Delta \geq 3$ and $B \in (0, B_c(\Delta))$, α_c is the occupancy fraction of the complete graph $K_{\Delta+1}$ evaluated at edge activity B and external field $\lambda_c(\Delta, B)$ [9]. We confirm this conjecture in the case $\Delta = 3$ and $B \in (0, 0.3128)$, noting that in the case of $\Delta = 3$ the conjecture covers values of B up to $B_c(3) = 1/3$.

Additionally, our results are closely related to the *free energy density* given by

$$F_G(B, \lambda) = \frac{1}{|V(G)|} \ln Z_G(\lambda).$$

The free energy density includes a convenient normalization factor that facilitates the comparison of graphs on different numbers of vertices. It is easy to see that

$$F_G(B, \lambda) = \int_0^{\lambda} \frac{1}{\ell} \alpha_G(B, \ell) \, \mathrm{d}\ell,$$

and hence an extremal result for the occupancy fraction in a suitably "downward-closed" region $\{(B, \ell) : 0 \leq \ell \leq \lambda\}$ immediately implies an extremal result for the free energy density.

Chapter 3: Related Work

Many other works tackle similar extremal problems to the one we study. A large area of focus is the study of the hard-core model, which is a probability distribution over independent sets in a graph. It is often used as a simple model of a gas, where a set being independent corresponds to gas particles not overlapping. This model has many similarities with the antiferromagnetic Ising model, in that it also contains two states (a vertex can either be in or out of a set) and "prefers" interactions between opposite spins (adjacent vertices cannot both be in an independent set). Due to this, both of these models are classified as two-state antiferromagnetic spin systems.

The equivalent minimization problem for the hard-core model was solved, with $K_{\Delta+1}$ being confirmed as the minimizer of both the free energy and the occupancy fraction over Δ-regular graphs [12]. These calculations were done implicitly by Cutler and Radcliffe [12] as part of a different proof related to counting independent sets. See this paper [13] for interpretation of this result in the language of occupancy fraction. The proof is short, and Cutler and Radcliffe obtain their result by establishing a bound on the number of ways to add a vertex to an independent set of size $t-1$ to yield an independent set of size t. They also make use of the fact that counting the number of independent sets in a clique is simple, since no independent set can contain more than one vertex.

The maximization problem for the hard-core model on regular graphs has also been solved in terms of both the free energy [14, 15] and more recently the occupancy fraction [16], resulting in $K_{\Delta,\Delta}$ being the maximizer over the entire parameter range for Δ-regular graphs. The free energy result was proven by Zhao in 2010 [14], verifying a conjecture of both Alon [17] and Kahn [15]. This result had already been proven in the special case that the graphs being considered were both Δ-regular and bipartite, due to an entropy approach done by Kahn [15] and later extended by Galvin and Tetali [18]. Zhao's proof introduced the *bipartite swapping trick*, which allowed him to reduce the problem to the bipartite case and use these known results.

An entropy-free proof of the bipartite case of the free energy upper bound was published in 2015 by Lubetzky and Zhao [19]. The key method in their proof is a generalized form of Hölder's inequality, which places an upper bound on the integral of a product of norms using the product of integrals. They first formulate the problem in terms of graph limits, allowing them to write a very short proof which follows immediately from the chosen form of Hölder's inequality.

Similar problems have been tackled for the Ising model itself, although the ferromagnetic version appears to be more well-studied. The ferromagnetic model differs in a key way from the antiferromagnetic one: there exists an FPRAS for the partition function at all parameter values [4]. However, Carlson et al. [20] show that the ferromagnetic model still contains algorithmic challenges of its own. The authors focus on the problem of sampling from the Ising model at fixed magnetization, proving the existence of a threshold much like the one we study in the antiferromagnetic model. Their techniques make use of correlation inequalities which do not apply to the antiferromagnetic model.

For the antiferromagnetic model, there are known results on extremal values of the free energy density, which is closely related to the occupancy fraction. The relevant free energy maximization result was proven in 2020, revealing that $K_{\Delta,\Delta}$ maximizes the free energy density in the entire parameter range [21, Corollary 1.15]. The authors make use of the bipartite swapping trick, which was introduced by Zhao to prove the maximization result for the hard-core model [14].

Given a graph G, the method carefully constructs two sets based on G and uses the idea of swapping vertex labels to construct an injection between them. The sets are set up in such a way that the existence of this injection directly implies an upper bound on the number of graph homomorphisms. Sah et al. use a modification of this technique along with induction to achieve their result on the free energy density [21].

Chapter 4: Results

Theorem 2. *The complete graph K_4 minimizes the occupancy fraction $\alpha_G(B, \lambda_c(3, B))$ over 3-regular graphs when $0 \leq B \leq 0.3128$.*

Theorem 2 confirms a conjecture of Davies and Perkins [9] that the complete graph minimizes the occupancy fraction in the case $\Delta = 3$ for the majority of the relevant parameter range. The conjecture for the case $\Delta = 3$ covers values of $B \in (0, \frac{1}{3})$, while we show this is true for the slightly smaller range $B \in (0, 0.3128)$. In this smaller range, Theorem 2 determines the value of the threshold for approximating the partition function at a fixed-magnetization. Inspired by a wider interest in graph theory, we also tackle the minimization question in 3-regular graphs over the entire parameter range $(B, \lambda) \in [0, 1]^2$. In this wider parameter range, our results on the occupancy fraction directly imply results on the free energy as well.

Theorem 3. *The complete graph K_4 minimizes the occupancy fraction $\alpha_G(B, \lambda)$ over 3-regular graphs in the union of the regions below:*

1. $\mathcal{R}_1^{\min} := \{(B, \lambda) : \frac{59}{100} \leq B \leq 1 \wedge 0 \leq \lambda \leq \frac{177}{200}B - \frac{267}{1000}B^2\}$
2. $\mathcal{R}_2^{\min} := \{(B, \lambda) : 0 \leq B \leq \frac{59}{100} \wedge 0 \leq \lambda \leq \frac{9}{10}B\}$
3. $\mathcal{R}_3^{\min} := \{(B, \lambda) : \frac{1}{10} \leq B \leq \frac{1}{2} \wedge \frac{1}{2}B \leq \lambda \leq \frac{99}{100}B\}$
4. $\mathcal{R}_4^{\min} := \{(B, \lambda) : \frac{1}{2} \leq B \leq \frac{9}{10} \wedge \frac{3}{20} \leq \lambda \leq \min\{B, \frac{53}{100}\}\}$
5. $\mathcal{R}_5^{\min} := \{(B, \lambda) : \frac{3}{5} \leq B \leq \frac{49}{50} \wedge \frac{7}{20} + \frac{11}{50}B \leq \lambda \leq \frac{3}{8} + \frac{6}{25}B\}$
6. $\mathcal{R}_6^{\min} := \{(B, \lambda) : \frac{1}{4} \leq B \leq \frac{9}{20} \wedge \frac{49}{50}B \leq \lambda \leq -\frac{17}{200} + \frac{7}{5}B\}$
7. $\mathcal{R}_7^{\min} := \{(B, \lambda) : \frac{56}{125} \leq B \leq \frac{3}{5} \wedge \frac{99}{100}B \leq \lambda \leq \frac{133}{200} - \frac{13}{50}B\}$
8. $\mathcal{R}_8^{\min} := \{(B, \lambda) : \frac{13}{50} \leq B \leq \frac{9}{25} \wedge -\frac{3}{40} + \frac{34}{25}B \leq \lambda \leq -\frac{9}{100} + \frac{36}{25}B\}$

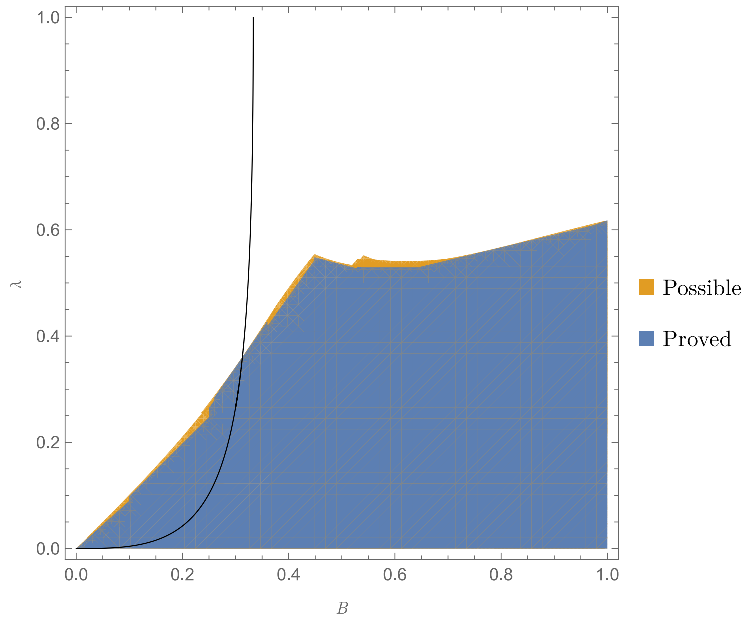

Figure 4.1: Minimizing the occupancy fraction over 3-regular graphs. The blue region is union of the regions in Theorem 3 where we proved that K_4 minimizes the occupancy fraction. The yellow region is where we numerically evaluated that it is theoretically possible for our method to show that K_4 minimizes the occupancy fraction, but we did not push the symbolic proofs of dual feasibility to the limit. The black line is $\lambda_c(3, B)$.

The complete graph K_4 also minimizes the free energy density $F_G(B, \lambda)$ over 3-regular graphs in the union of these regions.

The complete graph K_4 does not minimize the occupancy fraction $\alpha_G(B, \lambda)$ in the entire region $[0, 1]^2$. There is a region of the parameter space $\mathcal{R}^{\text{Pet}}$ where the Petersen graph has smaller occupancy fraction. It is also false that the minimizer is always either K_4 or the Petersen graph, as a third graph which we call the Goose graph[1] *has smaller occupancy fraction than either of these two in a region $\mathcal{R}^{\text{Goose}}$ of $[0, 1]^2$.*

We show the union of the regions in Theorem 3 in Figure 4.1; the K_4, Petersen, and Goose graphs in Figure 4.2; and a plot of where the Petersen and Goose graphs have smaller occupancy fraction in Figure 4.3.

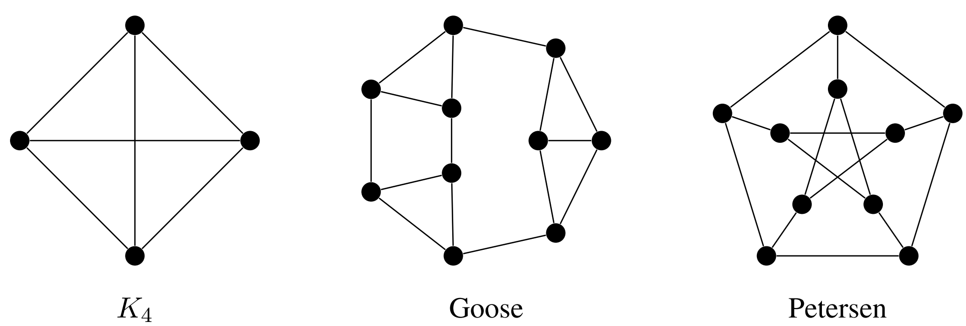

Figure 4.2: K_4, the Goose graph, and the Petersen graph. As stated in Theorem 3, the Goose graph and Petersen graph act as counterexamples to the idea that K_4 may minimize the occupancy fraction over all of $[0, 1]^2$. See Figure 4.3.

As seen in Figure 4.3, it is not possible K_4 is the minimizer in the entirety of $[0, 1]^2$, as there are graphs with lower occupancy fraction in some regions of the square. However, this neither proves nor disproves the original conjecture: that K_4 is the minimizing graph when $\lambda = \lambda_c(3, B)$.

Our methods yield insights into the corresponding problem of maximizing the occupancy fraction of the antiferromagnetic Ising model too.

[1]The name arises from a graph6 representation of the graph: `I}GOOSE@W`.

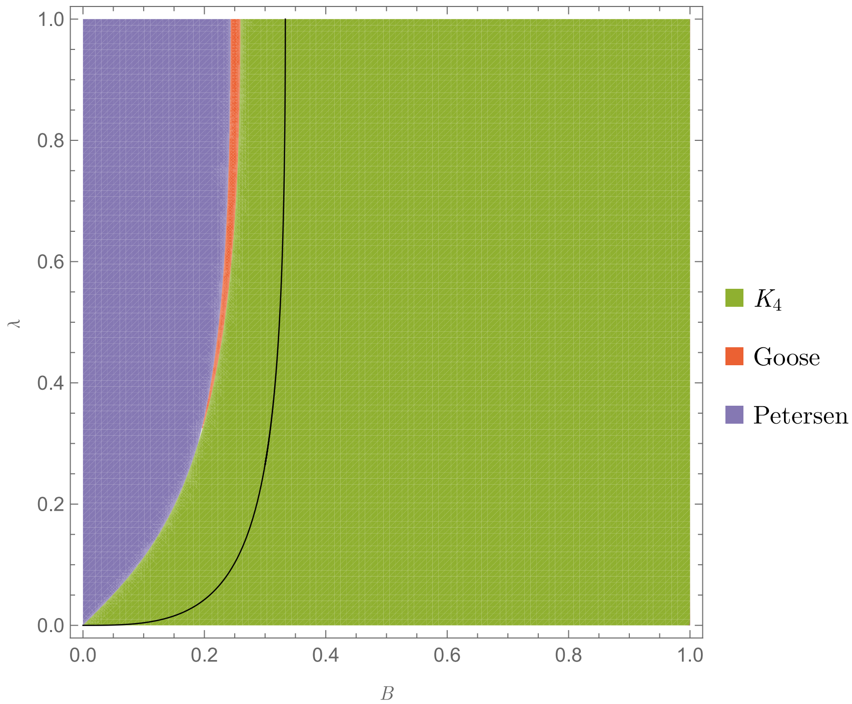

Figure 4.3: Considering K_4, Goose, and Petersen, the plot shows the region where each of them has the smallest occupancy fraction out of the three. The black line is $\lambda_c(3, B)$.

Theorem 4. *The complete bipartite graph $K_{3,3}$ maximizes the occupancy fraction $\alpha_G(B, \lambda)$ over 3-regular graphs in the union of the regions below:*

1. $\mathcal{R}_1^{\max} := \{(B, \lambda) : 0 \leq B \leq 1/5 \wedge 0 \leq \lambda \leq 3B/10\}$,
2. $\mathcal{R}_2^{\max} := \{(B, \lambda) : 1/5 \leq B \leq 2/5 \wedge 0 \leq \lambda \leq 4B/10\}$,
3. $\mathcal{R}_3^{\max} := \{(B, \lambda) : 2/5 \leq B \leq 1 \wedge 0 \leq \lambda \leq 5B/10\}$,

The relevant free energy maximization result for the antiferromagnetic Ising model is known, with $K_{\Delta,\Delta}$ being confirmed as the maximizing graph in the entire parameter range [21, Corollary 1.15]. Our results for the occupancy fraction in Theorem 4 can be used to recover this result on the free energy, but only in the subset of the parameter range for which we establish $K_{3,3}$ as the maximizer.

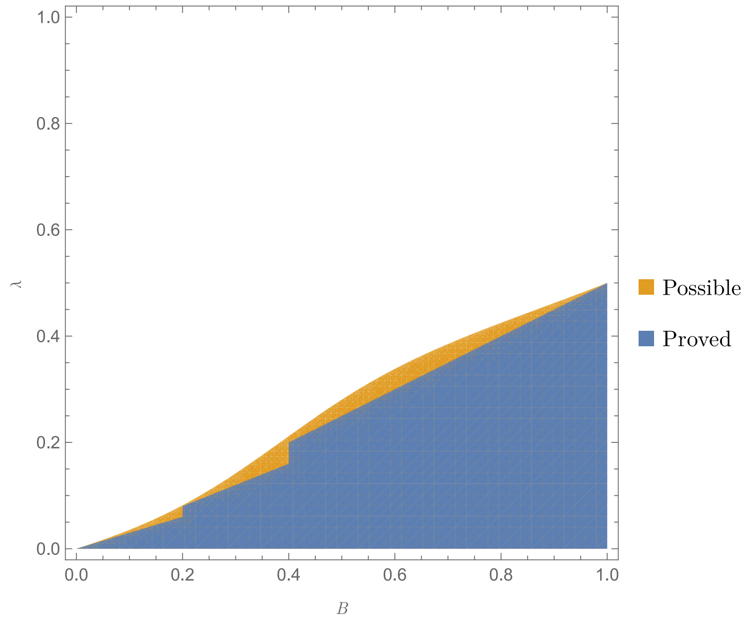

Figure 4.4: Maximizing the occupancy fraction over 3-regular graphs. The blue region is union of the regions in Theorem 4 where we proved that $K_{3,3}$ maximizes the occupancy fraction. The yellow region is where we numerically evaluated that it is theoretically possible for our method to show that $K_{3,3}$ maximizes the occupancy fraction, but we did not push the symbolic proofs of dual feasibility to the limit.

Chapter 5: Method

We apply the *occupancy method*, which has been developed fairly recently for this type of problem and takes a probabilistic approach. This method requires three key ingredients: a class of graphs, a spin model, and a physical observable. Note that an observable is a quantity related to the spin system that can be observed by looking at samples from the model. The occupancy fraction is an example of this because one can compute the density of the set of vertices with spin $+$ in samples of the model. The partition function (and free energy) are *not* observables; it is not clear how to compute the partition function from some samples from the model. Given the three ingredients, the occupancy method can then be used to attempt to find which graph in the class minimizes (or maximizes) the observable under the spin model in question.

Applications of the occupancy method appear in works which study a variety of spin systems, including the hard-core, Potts, and Widom–Rowlinson models [16, 22–26]. The method focuses on the idea of solving a linear programming relaxation of the relevant extremal problem. The variables of this linear program are typically created by sampling a small portion of the graph along with a spin assignment from the model. In this case, one chooses a spin assignment σ distributed according to the model, randomly chooses a small subgraph F of the graph, and then reveals the σ restricted to some prescribed subset of the vertices of F.

The observable is then written in terms of these variables, and set as the objective of the linear program. Constraints are developed by finding relationships that hold for any graph in the class. Formulating the problem as a linear program in this way often expands the feasible set from the original graph based problem, but if a solution to the linear programming relaxation can be shown to arise from the spin model on a graph in the class, it must also be a solution to the original extremal problem.

In our case, we consider the class of 3-regular graphs, the Ising spin model, and the occupancy fraction (our observable). Using this setup, minimizing the occupancy fraction over a class of graphs corresponds to finding the minimum occupancy fraction over the feasible set of variables which arise when one considers the spin model on graphs in the class.

5.1 Local Views

As with previous applications of the occupancy method, we first define random variables based on sampling a small subgraph and spin assignment. These variables will then be used to formulate linear constraints.

We define a *local view* through the following experiment:

1. Fix Δ, B, λ and let G be a Δ-regular graph.

2. Select a vertex $\mathbf{u}$ of G uniformly at random.

3. Select a spin assignment $\boldsymbol{\sigma}$ from the Ising model $\mu_{G,B,\lambda}$.

4. Let the random variable $\mathbf{L}$ be the induced subgraph $G[N^2[\mathbf{u}]]$ together with the spin assignment $\boldsymbol{\tau} = \boldsymbol{\sigma}|_{N^2(\mathbf{u})}$ restricted to the vertices at distance 2 from $\mathbf{u}$.

We call these variables local views because they capture local spin information and graph structure specific to the class of graphs being studied. It is important to retain the information that $\mathbf{u}$ in $\mathbf{L}$ was the vertex sampled from the graph, so we may consider $\mathbf{L}$ a rooted graph with $\mathbf{u}$ identified as the root, as well as having a spin assignment for $N^2(\mathbf{u})$. Formally, a local view is then a tuple $\mathbf{L} = (H, u, \tau)$ where H is a graph containing the specified vertex u, and τ is a spin assignment to a (possibly empty) subset of the vertices of H.

In all graphs H of interest, the specified u will have Δ neighbors, and each of these neighbors also has Δ neighbors, but beyond that we must enumerate all possible rooted graphs (H, u) and spin assignments τ to $N^2(u)$ in H. In fact, for the Ising model on 3-regular graphs there are exactly 23 local views up to isomorphism (where we consider only graph isomorphisms that fix the root). See Appendix B for images of these 23 local views.

The variables of the optimization problem are then the probabilities of the local views in the distribution given by the Ising model.

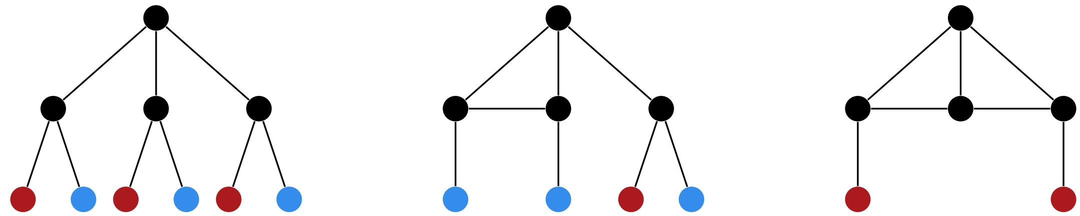

Figure 5.1: Examples of local views that are possible in a 3-regular graph. The top-most vertex is the root u of the local view, and vertices assigned $+$ or $-$ are colored red or blue respectively.

5.2 Optimization

The occupancy method consists of using these local views to reframe our graph-based optimization problem in terms of a linear program relaxation. In basic applications of the occupancy method, the relaxed optimization problems can be analyzed without serious computation. However, there are more complex examples (e.g. [27, 28]) for which the number of variables becomes unmanageable and one turns to a computer analysis. For our case, enumerating all the local views, which give us our variables for the optimization problem, is a nontrivial combinatorial problem in its own right. Davies encounters this challenge while applying the occupancy method to a similar problem [27], and turns to *canonical isomorph-free generation* techniques [29]. We similarly apply these techniques to our local view generation, making the process more efficient by preventing generation of isomorphic local views (see Appendix A).

Aside from the practical issues of scaling one's approach to large numbers of variables, there are two main problems that can arise. Firstly, while any given finite linear program can be solved in finite time, we have an infinite family of linear programs parameterized by the relevant variables Δ, B and λ. It is not necessarily straightforward to establish for the entire family that some candidate solution (as a function of Δ, B, λ) is indeed optimal.

In order to address this, we follow the standard approach to exploit linear programming duality. First, we construct a solution for the dual whose objective value is the occupancy fraction of the conjectured graph, which involves solving a linear system defined by the active primal constraints to find dual variables. We must then show this potential solution is dual feasible, which consists of proving that a large number of inequalities involving polynomials in the parameters B and λ all hold. In general, these inequalities could also include the variable Δ (see these works [14, 16] for examples), but this is not the case for us as we fix Δ in the creation of our linear program.

A second, more fundamental problem is that there can be a gap between the true graph-defined optimal value and the optimal value of the linear relaxation. We confront each of these problems in our study of the antiferromagnetic Ising model and we are only partially successful in solving these issues, which is why our main result applies only to a subset of the desired parameter space.

A novel difficulty in our case is that exact computation requires handling algebraic reals. In previous works (e.g. [16, 27, 28]) one can choose rational values of the parameters and employ exact rational linear program solvers to investigate the problem, such as to find active primal constraints. However, as seen in Equation (2.1), $\lambda_c(\Delta, B)$ contains square roots and is not necessarily rational when $\Delta \geq 3$ is an integer and $B \in [0, 1] \cap \mathbb{Q}$. For our results related to the fixed-magnetization computational threshold, we must investigate our linear programming relaxation exactly at this $\lambda_c(\Delta, B)$ curve, requiring us to work with algebraic reals. This necessitates the use of more flexible combinatorial optimization software and costs significantly more computation time.

5.3 Proof

In this section we prove Theorems 2, 3 and 4 with the assistance of a computer. We must first enumerate the possible local views and collect them in a set we denote $\mathcal{L}$. As mentioned in the Optimization section, this process is done using canonical isomorph-free generation (see Appendix A).

Each Δ-regular graph G gives us a probability distribution $x = x(G, B, \lambda)$ over $\mathcal{L}$ which we represent as a vector in $[0, 1]^{\mathcal{L}}$. We establish a set of linear constraints on the feasible vectors x using the fact that in any regular graph, a uniform random neighbor $\mathbf{v}$ of a uniform random vertex $\mathbf{u}$ is itself distributed uniformly over the vertices. Thus, we have two ways of computing statistics of a uniform random vertex given the distribution x of the random local view. Our constraints arise from equating the two ways of writing these statistics.

For convenience, we define for a local view $\mathbf{L} = (H, u, \tau)$ the distribution μ_L on spin assignments of H as the Ising model $\mu_{H,B,\lambda}$ on H conditioned on the spins of $N^2(u)$ agreeing with τ. We deliberately suppress the dependence of μ_L on B and λ as they remain fixed throughout.

One can express the occupancy fraction $\alpha_G(B, \lambda)$ as the probability that a uniform random vertex $\mathbf{u}$ receives spin $+$ under $\mu_{G,B,\lambda}$. By the spatial Markov property of the Ising model, this is equal to the probability that in the random local view $\mathbf{L} = (H, u, \tau)$ distributed according to $x(G, B, \lambda)$, the vertex u gets spin $+$ under the measure μ_L on spin assignments to H. That is, we define α_L to be the probability that under μ_L, the vertex u in the local view $\mathbf{L}$ gets spin $+$ and we have

$$\alpha_G(B, \lambda) = \sum_{L \in \mathcal{L}} \alpha_L x_L,$$

where x_L is the distribution of the random local view $\mathbf{L}$ in G. Importantly, this gives us a linear objective function which we wish to minimize over feasible x.

For linear constraints on the feasible x, note that one can also compute the probability that a uniform random neighbor of u receives spin $+$ under μ_L, which is a different computation over the local views. This provides one linear equality constraint on the distribution x that we wish to use, though we do not explicitly add this constraint to our linear program because we turn to richer statistics to obtain more constraints. For a local view $\mathbf{L} = (H, u, \tau)$ and $j \in \{0, 1, \ldots, \Delta\}$, let $\gamma^u(j)_L$ be the probability that under μ_L, exactly j neighbors of u get spin $+$. Similarly, let $\gamma^{N(u)}(j)_L$ be the probability that under μ_L, a uniform random neighbor of u has exactly j neighbors with spin $+$. Then for any x which arises from running the local view experiment on a graph and any $0 \leq j \leq \Delta$,

$$\sum_{L \in \mathcal{L}} \left(\gamma^u(j)_L - \gamma^{N(u)}(j)_L \right) x_L = 0,$$

giving $\Delta + 1$ linear constraints on x. Constraints of this type first appear in an occupancy fraction based proof on the number of matchings in a graph [16]. Together with the facts that the entries of x are nonnegative and sum to one, we can express a linear programming relaxation of the extremal problem of minimizing the occupancy fraction.

Definition 1. For a given set $\mathcal{L}$ of local views and set J of constraint indices, the *primal occmin program* is the linear program with variable $x \in \mathbb{R}^{\mathcal{L}}$ given by

$$\begin{aligned} \min \sum_{L\in\mathcal{L}} \alpha_L x_L \text{ s.t.} \\ \sum_{L\in\mathcal{L}} x_L = 1 \\ \sum_{L\in\mathcal{L}} \left(\gamma^u(j)_L - \gamma^{N(u)}(j)_L\right) x_L = 0 \qquad \forall j \in J \\ x_L \geq 0 \qquad \forall L \in \mathcal{L}. \end{aligned}$$

The *primal occmax program* is identical except for the fact that the objective is to be maximized.

Given $\mathcal{L}$ and J as above, the *dual occmin program* is the linear program with variable $y \in \mathbb{R}^{\{p\}\cup J}$ given by

$$\begin{aligned} \max y_p \text{ s.t.} \\ y_p + \sum_{j\in J} \left(\gamma^u(j)_L - \gamma^{N(u)}(j)_L\right) y_j \leq \alpha_L \qquad \forall L \in \mathcal{L}. \end{aligned}$$

Note that we use p to index the dual variable corresponding to the primal constraint $\sum_{L\in\mathcal{L}} x_L = 1$ and the set J to index the other dual variables. We define the corresponding *dual occmax program* analogously.

The dual occmin and dual occmax programs are the duals (in a standard linear programming sense) of the primal occmin and primal occmax programs, respectively.

5.3.1 Minimization: Theorem 2 and Theorem 3

Since K_4 has diameter 1, one of our local views is a rooted copy of K_4 together with an empty spin assignment (as there are no vertices at distance two from the root), and the vector x corresponding to K_4 has a 1 in the entry indexed by K_4 and zeros elsewhere. Thus, to argue that this x from K_4 is optimal in the primal occmin program with local views $\mathcal{L}$ and constraint indices J, it suffices to show that the optimum value achieved in the program with local views $\mathcal{L}' := \mathcal{L} \setminus \{K_4\}$ (and the same J) is at least the occupancy fraction of K_4. Given this, in any 3-regular graph without a component isomorphic to K_4 we know that the occupancy fraction is at least that of K_4. If the graph does have components isomorphic to K_4 then we can use the fact that for a disjoint union $G = G_1 \sqcup G_2$ we have

$$|V(G)|\alpha_G(B, \lambda) = |V(G_1)|\alpha_{G_1}(B, \lambda) + |V(G_2)|\alpha_{G_1}(B, \lambda).$$

That is, to prove Theorems 2 and the statement about K_4 being the minimizer in Theorem 3, it suffices to fix J and perform the following steps with the primal and dual occupancy programs with local view set $\mathcal{L}'$.

1. Identify the local views indexing three tight dual constraints, say $\{L_1, L_2, L_3\} \subset \mathcal{L}'$.
2. Solve for $y = (y_p, y_0, y_1)$ in the linear system obtained by setting the dual constraints indexed by (L_1, L_2, L_3) to equality.
3. Check that every dual constraint holds for $y = (\alpha_{K_4}(B, \lambda), y_0, y_1)$.

Note that for this to have any chance of working, we require that the y_p obtained is at least $\alpha_{K_4}(B, \lambda)$. This will always be the case when (B, λ) are such that indeed K_4 minimizes the occupancy fraction over 3-regular graphs, but even when this does hold, our linear program can be a loose relaxation of the graph-theoretic problem and the method fails.

Since we focus on the $\Delta = 3$ case, we could in theory fix J to be $\{0, 1, 2, 3\}$. However, the constraint corresponding to $j = 3$ follows from the others as probabilities must sum to one, and empirical evaluation indicated $j = 2$ is also redundant. Because of this, we fix the constraint index set to be $J = \{0, 1\}$.

Having done the above steps, we have shown that there is some y feasible in the dual which achieves objective value equal to $\alpha_{K_4}(B, \lambda)$, and hence by LP duality the optimum of the primal is at least $\alpha_{K_4}(B, \lambda)$. The first step looks a bit mysterious, but can be done empirically by solving the primal program and looking at the nonnegative values in optimum x. The reason that we solve for $y = (y_p, y_0, y_1)$ in the three chosen tight dual constraints and then set y_p to $\alpha_{K_4}(B, \lambda)$ is that having taken out the local view K_4 the optimum value in the LPs with local view set $\mathcal{L}'$ could be strictly greater than $\alpha_{K_4}(B, \lambda)$. Thus, setting y_p to the slightly smaller $\alpha_{K_4}(B, \lambda)$ gives us slightly more room to establish the remaining dual constraints. While one could imagine other methods for finding dual variables inspired by complementary slackness, this method proved reasonably successful for us.

When following the above steps, the dual constraints (each of which is indexed by some $L \in \mathcal{L}'$) become inequalities equivalent to

$$0 \leq \alpha_L - \alpha_{K_4(B,\lambda)} - \sum_{j \in \{0,1\}} \left(\gamma^u(j)_L - \gamma^{N(u)}(j)_L \right) y_j,$$

where each term is a rational function of B and λ. Thus, provided one can establish that a set of 22 inequalities (19 if you discount the three that we know must hold by construction) of rational functions hold, one can prove Theorem 2 and the minimization statement in Theorem 3 for regions of the parameter space.

We use the general-purpose mathematical computing programs SageMath [30] and Wolfram Engine [31] to assist with the following.

1. Generate $\mathcal{L}'$.

2. Manually identify a collection $\mathcal{T}$ of sets of three tight dual constraints.

3. For each set $T \in \mathcal{T}$, solve for the relevant dual variables and manually identify a region of the parameter space $(B, \lambda) \in [0, 1]^2$ in which we wish to prove symbolically that the dual constraints hold with these variables.

4. Prove symbolically that in the specified region the dual constraints indeed hold.

This constitutes a computer-assisted proof of the statement about regions where K_4 is the minimizer in Theorem 3, and Theorem 2 follows also because the curve given by $\lambda = \lambda_c(\Delta, B)$ lies in the union of the regions in which we show dual feasibility for $B \in (0, 0.3128)$.

While we could in principle continue for larger B, at some point shortly after $B = 0.3128$ the optimum value in the LP is smaller than that of K_4. We verified this for $B = 0.32$ and $\lambda_c(3, B)$ over the algebraic reals symbolically with SageMath.

The statement about K_4 minimizing the free energy density follows immediately from the occupancy fraction result and the fact that the union of our regions is downward-closed in the sense that, for any (B, λ) in the union of the regions, it is true that (B, ℓ) is in the union for $0 \leq \ell \leq \lambda$.

To complete the proof of Theorem 3, we use a computer to search through all 3-regular graphs on at most 14 vertices and plot the obtained minimizers. While 14 vertices is not necessarily representative of the behavior one could observe in general, it is already enough to prove that the minimizer is not always K_4, and that it is not always either K_4 or the Petersen graph.

5.3.2 Maximization: Theorem 4

The proof of Theorem 4 is analogous. This time, it is not the case that a single local view represents $K_{3,3}$, but it is true that disjoint unions of $K_{3,3}$ are the only 3-regular graphs in which the three neighbors of the root u of the local view always see the same spins on the vertices at distance two from u. That is, the only three local views one can obtain in $K_{3,3}$ are the ones shown in Figure 5.2, and disjoint unions of $K_{3,3}$s are the only graphs with this property. One can check that the only feasible solution to the LP supported on these three local views is indeed the exact distribution on local views one gets in $K_{3,3}$. That is, our constraints are enough to capture this key property.

Then we consider only the set of tight constraints corresponding to those three local views, and perform the same kind of dual feasibility analysis over the entire local view set $\mathcal{L}$. The only differences are that the y_p obtained is $\alpha_{K_{3,3}(B,\lambda)}$ and that the dual constraint inequalities are in the

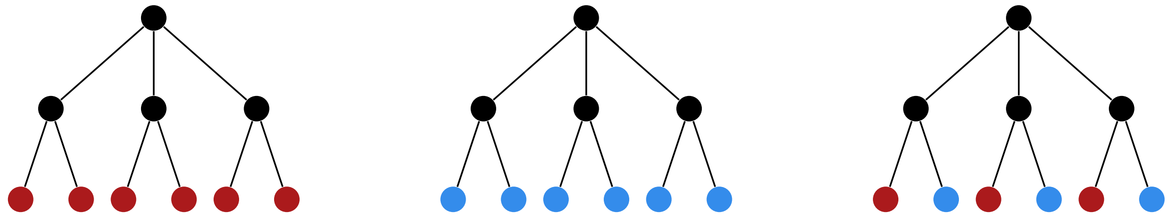

Figure 5.2: Local views one can obtain in $K_{3,3}$. The topmost vertex is the root u of the local view, and vertices assigned $+$ or $-$ are colored red or blue respectively.

opposite direction. That is, we must check inequalities of the form

$$0 \geq \alpha_L - \alpha_{K_{3,3}(B,\lambda)} - \sum_{j \in \{0,1\}} \left(\gamma^u(j)_L - \gamma^{N(u)}(j)_L \right) y_j$$

to verify dual feasibility.

Chapter 6: Discussion

We obtained only partial results for the minimization problem, as our methods did not extend to $B = \frac{1}{3}$. Throughout this process, we gathered evidence that the problem includes unique challenges that are not present in similar problems involving other well-known spin systems. In this section we identify a few ways our method could be improved, and problems we encountered.

6.1 Proving Dual Feasibility for Larger Region

In the process of solving for dual variables, we chose sets of three constraints to set to equality. These were chosen based on an empirical evaluation which resulted in four constraint sets for the minimization problem, and one for the maximization problem. In principle, one could consider every subset $T \in \mathcal{L}'$ of size 3, solve those dual constraints for equality, and find the maximal region of $[0, 1]^2$ in which the corresponding dual feasibility inequalities hold. We decided not to perform this rather extensive computation as we verified that this would not cover the entire parameter space, leaving out for example $B = 0.32$ and $\lambda = \lambda_c(3, 0.32)$.

Additionally, we solved numerically for the four maximal regions the sets of tight constraints we considered in order to plot Figure 4.3, but we did not prove symbolically that dual feasibility holds in these exact regions, settling for a slightly smaller region bounded by simpler functions of B and λ. This is because we found that Wolfram Engine could show dual feasibility rather quickly in regions bounded by linear or quadratic functions, and so manually fit such regions inside the plotted ones. It appears that more specialized tools and significant computing power would be required to prove symbolically that the minimizer is K_4 for the entire region in which the linear program is a tight relaxation of the graph-defined optimization problem.

6.2 Deeper Local Views

The setup for occupancy fraction optimization that we use is flexible, and one can define "deeper" local views that capture more graph-theoretic structure in the problem. We attempted to extend our local views to depth three, which consisted of changing the vertex set from $N^2[u]$ to $N^3[u]$ and storing spin information on vertices of distance three from u. We applied the proof technique to these depth three local views in 3-regular graphs and found (after significant engineering effort and computation time) that there are 2637 local views in this case. This is just small enough that one can investigate the corresponding primal occmin program numerically, but the overhead of exact rational or algebraic real computation proved too much to handle. While we believe that this program could enlarge the known region in $[0, 1]^2$ where K_4 is optimal, it seems computationally infeasible to extend our dual feasibility analysis to this many local views. Not only would there be roughly 100 times more inequalities to check for dual feasibility, each such inequality is now a much larger polynomial with significantly higher degree than in the case of local views of depth 2.

Initial numerical experiments suggest that the deeper program can show K_4 to be the minimizer in a larger region, but that this program still fails to be a tight relaxation in the entire parameter range.

6.3 Additional Constraints

For many parameter values for which we did not prove $K_{\Delta+1}$ was the minimizing graph, the output of the primal occmin program was a collection of probabilities that could not correspond to a real Δ-regular graph. To combat this, we tried adding a different type of constraints to see if the program could be pushed further. The type of constraints we explored hinged on the idea of 'flipping' the spin of a vertex from $+$ to $-$ or $-$ to $+$. Constraints of this nature were used to tackle a similar problem for the hard-core model, with much more success [26]. A hurdle to adapting this type of constraint to our problem is that the change in the probability of a configuration due to flipping a spin is not uniquely defined by the local view as it is for the hard-core model. In the hard-core model, a vertex can always be removed from an independent set, resulting in a known, quantifiable change in the probability of the given set being selected from the model. However, the effect of flipping a spin under the Ising model depends on the neighborhood of the vertex in question, including vertices whose spins are unknown (since our local views only have spins on the second neighbors of u). We tried bounding the impact of flipping a spin, which yielded legal constraints, but the bound was extremely loose. Because of this, the constraints were not able to significantly widen the parameter range in which our program was a tight relaxation.

Chapter 7: Conclusions

We succeed in determining the value of the computational threshold in sampling from the antiferromagnetic Ising model for the parameter range $B \in (0, 0.3128)$, which covers the majority of the relevant space $B \in (0, \frac{1}{3})$. This confirms a conjecture of Davies and Perkins [9] that the complete graph $K_{\Delta+1}$ minimizes the occupancy fraction over Δ-regular graphs for the case $\Delta = 3$, but only in the region $B \in (0, 0.3128)$. We also obtain results on the minimization and maximization problems in other areas of the parameter space, allowing us to confirm that K_4 does not minimize the occupancy fraction in the entire range of possible B, λ values.

With additional time and computing power, we expect our methods could be pushed further. However, since our linear program was not a tight relaxation of the original optimization problem, it would take more fundamental changes in order to solve the problem over the entire parameter space relevant to the computational threshold.

It would be interesting to develop new methods to understand occupancy fractions, as for the closely related free energy density there are more techniques available, such as induction [21] and Hölder inequality [19] based approaches.

Bibliography

[1] J. Majewski, H. Li, and J. Ott, "The Ising model in physics and statistical genetics," en, *Am. J. Hum. Genet.*, vol. 69, no. 4, pp. 853–862, 2001.

[2] D. Stauffer, "Social applications of two-dimensional Ising models," *American Journal of Physics*, vol. 76, no. 4, pp. 470–473, 2008, ISSN: 0002-9505, DOI: 10.1119/1.2779882.

[3] E. Schneidman, M. J. Berry 2nd, R. Segev, and W. Bialek, "Weak pairwise correlations imply strongly correlated network states in a neural population," en, *Nature*, vol. 440, no. 7087, pp. 1007–1012, 2006.

[4] M. Jerrum and A. Sinclair, "Polynomial-time approximation algorithms for the Ising model," *SIAM Journal on Computing*, vol. 22, no. 5, pp. 1087–1116, 1993, ISSN: 0097-5397, DOI: 10.1137/0222066.

[5] A. Galanis, D. Štefankovič, and E. Vigoda, "Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models," *Combinatorics, Probability and Computing*, vol. 25, no. 4, pp. 500–559, 2016, DOI: 10.1017/S0963548315000401.

[6] A. Sly and N. Sun, "The computational hardness of counting in two-spin models on d-regular graphs," in *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, Los Alamitos, CA, USA: IEEE Computer Society, 2012, pp. 361–369, DOI: 10.1109/FOCS.2012.56.

[7] A. Sinclair, P. Srivastava, and M. Thurley, "Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs," *Journal of Statistical Physics*, vol. 155, no. 4, pp. 666–686, 2014, ISSN: 1572-9613, DOI: 10.1007/s10955-014-0947-5.

[8] L. Li, P. Lu, and Y. Yin, "Correlation decay up to uniqueness in spin systems," in *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '13, New Orleans, Louisiana: Society for Industrial and Applied Mathematics, 2013, pp. 67–84, ISBN: 9781611972511.

[9] E. Davies and W. Perkins, "Approximately counting independent sets of a given size in bounded-degree graphs," *SIAM Journal on Computing*, vol. 52, no. 2, pp. 618–640, 2023, ISSN: 0097-5397, DOI: 10.1137/21M1466220.

[10] E. Davies and O. LeBlanc, "On the occupancy fraction of the antiferromagnetic Ising model," 2024, preprint, arXiv: 2412.18070 [math.CO].

[11] R. J. Baxter, *Exactly Solved Models in Statistical Mechanics*. London ; New York: Academic Press, 1982, ISBN: 978-0-12-083180-7.

[12] J. Cutler and A. Radcliffe, "The maximum number of complete subgraphs in a graph with given maximum degree," *Journal of Combinatorial Theory, Series B*, vol. 104, pp. 60–71, 2014, ISSN: 00958956, DOI: 10.1016/j.jctb.2013.10.003.

[13] E. Davies, M. Jenssen, W. Perkins, and B. Roberts, "Tight bounds on the coefficients of partition functions via stability," *Journal of Combinatorial Theory, Series A*, vol. 160, pp. 1–30, 2018, ISSN: 00973165, DOI: 10.1016/j.jcta.2018.06.005.

[14] Y. Zhao, "The number of independent sets in a regular graph," *Combinatorics, Probability and Computing*, vol. 19, no. 2, pp. 315–320, 2010, ISSN: 0963-5483, 1469-2163, DOI: 10.1017/S0963548309990538.

[15] J. Kahn, "An entropy approach to the hard-core model on bipartite graphs," *Combinatorics, Probability and Computing*, vol. 10, no. 3, pp. 219–237, 2001, ISSN: 0963-5483, 1469-2163, DOI: 10.1017/S0963548301004631.

[16] E. Davies, M. Jenssen, W. Perkins, and B. Roberts, "Independent sets, matchings, and occupancy fractions," *Journal of the London Mathematical Society*, vol. 96, no. 1, pp. 47–66, 2017, ISSN: 00246107, DOI: 10.1112/jlms.12056.

[17] N. Alon, "Independent sets in regular graphs and sum-free subsets of finite groups," *Israel Journal of Mathematics*, vol. 73, no. 2, pp. 247–256, 1991, ISSN: 1565-8511, DOI: 10.1007/BF02772952.

[18] D. Galvin and P. Tetali, "On weighted graph homomorphisms," in *Graphs, Morphisms and Statistical Physics*, ser. DIMACS Ser. Discrete Math. Theoret. Comput. Sci. Vol. 63, Amer. Math. Soc., Providence, RI, 2004, pp. 97–104, DOI: 10.1090/dimacs/063/07.

[19] E. Lubetzky and Y. Zhao, "On replica symmetry of large deviations in random graphs," *Random Struct. Algorithms*, vol. 47, no. 1, pp. 109–146, 2015, ISSN: 1042-9832, DOI: 10.1002/rsa.20536.

[20] C. Carlson, E. Davies, A. Kolla, and W. Perkins, "Computational thresholds for the fixed-magnetization Ising model," in *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, Rome Italy: ACM, 2022, pp. 1459–1472, ISBN: 978-1-4503-9264-8, DOI: 10.1145/3519935.3520003.

[21] A. Sah, M. Sawhney, D. Stoner, and Y. Zhao, "A reverse Sidorenko inequality," *Inventiones mathematicae*, 2020, ISSN: 0020-9910, 1432-1297, DOI: 10.1007/s00222-020-00956-9.

[22] E. Cohen, W. Perkins, and P. Tetali, "On the Widom–Rowlinson occupancy fraction in regular graphs," *Combinatorics, Probability and Computing*, vol. 26, no. 2, pp. 183–194, 2017, ISSN: 0963-5483, 1469-2163, DOI: 10.1017/S0963548316000249.

[23] E. Davies, M. Jenssen, W. Perkins, and B. Roberts, "On the average size of independent sets in triangle-free graphs," *Proceedings of the American Mathematical Society*, vol. 146, no. 1, pp. 111–124, 2017, ISSN: 0002-9939, 1088-6826, DOI: 10.1090/proc/13728.

[24] E. Davies, M. Jenssen, W. Perkins, and B. Roberts, "Extremes of the internal energy of the Potts model on cubic graphs," *Random Structures & Algorithms*, vol. 53, no. 1, pp. 59–75, 2018, ISSN: 10429832, DOI: 10.1002/rsa.20767.

[25] E. Davies, R. J. Kang, F. Pirot, and J.-S. Sereni, "Graph structure via local occupancy," 2020, arXiv: 2003.14361 [math].

[26] Y. Zhao, "Extremal regular graphs: Independent sets and graph homomorphisms," *The American Mathematical Monthly*, vol. 124, no. 9, p. 827, 2017, ISSN: 00029890, DOI: 10.4169/amer.math.monthly.124.9.827.

[27] E. Davies, "Counting proper colourings in 4-regular graphs via the Potts model," *The Electronic Journal of Combinatorics*, vol. 25, no. 4, P4.7, 2018, ISSN: 1077-8926, DOI: 10.37236/7743.

[28] G. Perarnau and W. Perkins, “Counting independent sets in cubic graphs of given girth,” *Journal of Combinatorial Theory, Series B*, vol. 133, pp. 211–242, 2018, ISSN: 00958956, DOI: 10.1016/j.jctb.2018.04.009.

[29] B. D. McKay, “Isomorph-free exhaustive generation,” *Journal of Algorithms*, vol. 26, no. 2, pp. 306–324, 1998, ISSN: 0196-6774, DOI: 10.1006/jagm.1997.0898.

[30] W. A. Stein et al, *Sage Mathematics Software (Version 10.4)*, The Sage Development Team, 2024, [Online]. Available: www.sagemath.org.

[31] Wolfram Research, Inc., *Wolfram Engine (Version 14.1)*, 2024, [Online]. Available: www.wolfram.com/engine.

[32] B. McKay, *Structure generation and approximate counting*, Presented at KU Leuven as part of the International Francqui Professor Lecture Series, 2019, [Online]. Available: https://www.ugent.be/we/winst/nl/slides-brendan-mckay-lecture-kuleuven.pdf.

[33] D. Stolee, *The canonical augmentation method*, Presented at Cumberland Conference on Combinatorics, Graph Theory & Computing, Louisville, KY., 2011, [Online]. Available: https://lidicky.name/dstolee/presentations/Sto11-CanonicalAugmentations.pdf.

Appendix A: Canonical Isomorph-Free Generation

In order to generate our local views efficiently, we use canonical isomorph-free generation, which can be used to generate a list of combinatorial objects such that no two are isomorphic. This essentially prevents the creation of "redundant" elements, since the labeling of the vertices does not matter for our application.

We start with the idea that graphs can be incrementally constructed by starting with a single vertex, and progressively adding vertices. This construction process can be viewed as a tree, such that the root is a single vertex, and the children of any node X are all possible results of extending X by adding a vertex and any associated edges. If we consider the root node to be the 1^{st} layer, the i^{th} layer of the tree would then contain all graphs on i vertices. The problem is that this tree would contain many isomorphic copies of the same graphs.

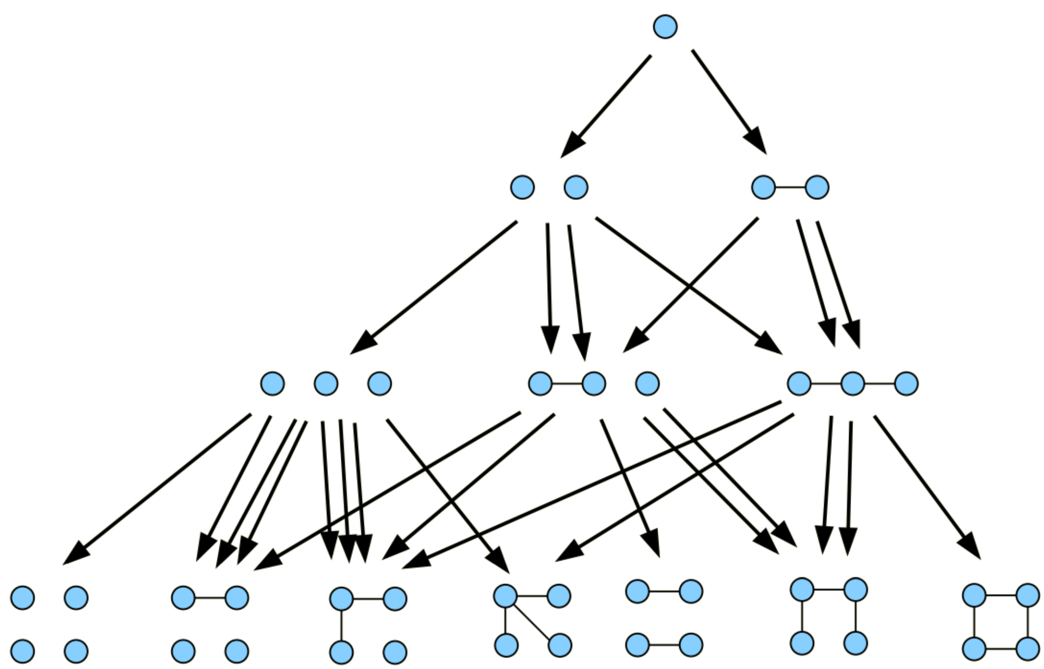

Figure A.1: An example of this construction tree for the case of generating triangle-free planar graphs. To conserve space, nodes with isomorphic graphs have been collapsed. Taken from a presentation by Brendan McKay [32].

The method we use to address this was introduced by Brendan McKay in 1998 [29]. McKay's algorithm prunes branches of this construction tree which will create isomorphic copies. This appendix provides an overview of the algorithm for the case of graph generation, see McKay's paper [29] for a more thorough description and proof that his method ensures no isomorphs are generated.

In order to decide which one of the isomorphic copies of a given graph to keep, the idea of a *canonical deletion* is introduced. A canonical deletion is designed to choose a part of the graph to delete such that the selected part is the same for all graphs in an isomorphism class. This concept is key to McKay's algorithm, and can be used to keep only graphs which were generated in a canonical way. In the algorithm, it will essentially play the role of identifying which extension "should have" been done to get to a given graph [33].

There are two key changes made from a naive generation to ensure no isomorphs are created. The first is that only one expansion is tried from each set of isomorphic expansions. This prevents a single graph from having two isomorphic children in the construction tree. However, it is still possible to generate isomorphs, as two different graphs can be extended to create isomorphic children. To prevent these, the second change is to check whether each expansion is the inverse of a canonical deletion. If it is not, the expansion is not done, as the resulting graph would not be created in a canonical way.

We break down our process of local view generation, assuming we start with one root vertex, and add one layer at a time. It consists of the following steps:

1. Add vertices for new layer, connecting them to the previous layer
2. 'Fill' the layer by adding any edges within it
3. Repeat steps 1 and 2 until the local view has the desired depth
4. Assign spins to the last layer of vertices

Steps 1, 2, and 4 can all introduce isomorphs, so we use canonical isomorph-free generation in slightly different ways at each step.

Appendix B: Local Views for the Ising Model (3-regular)

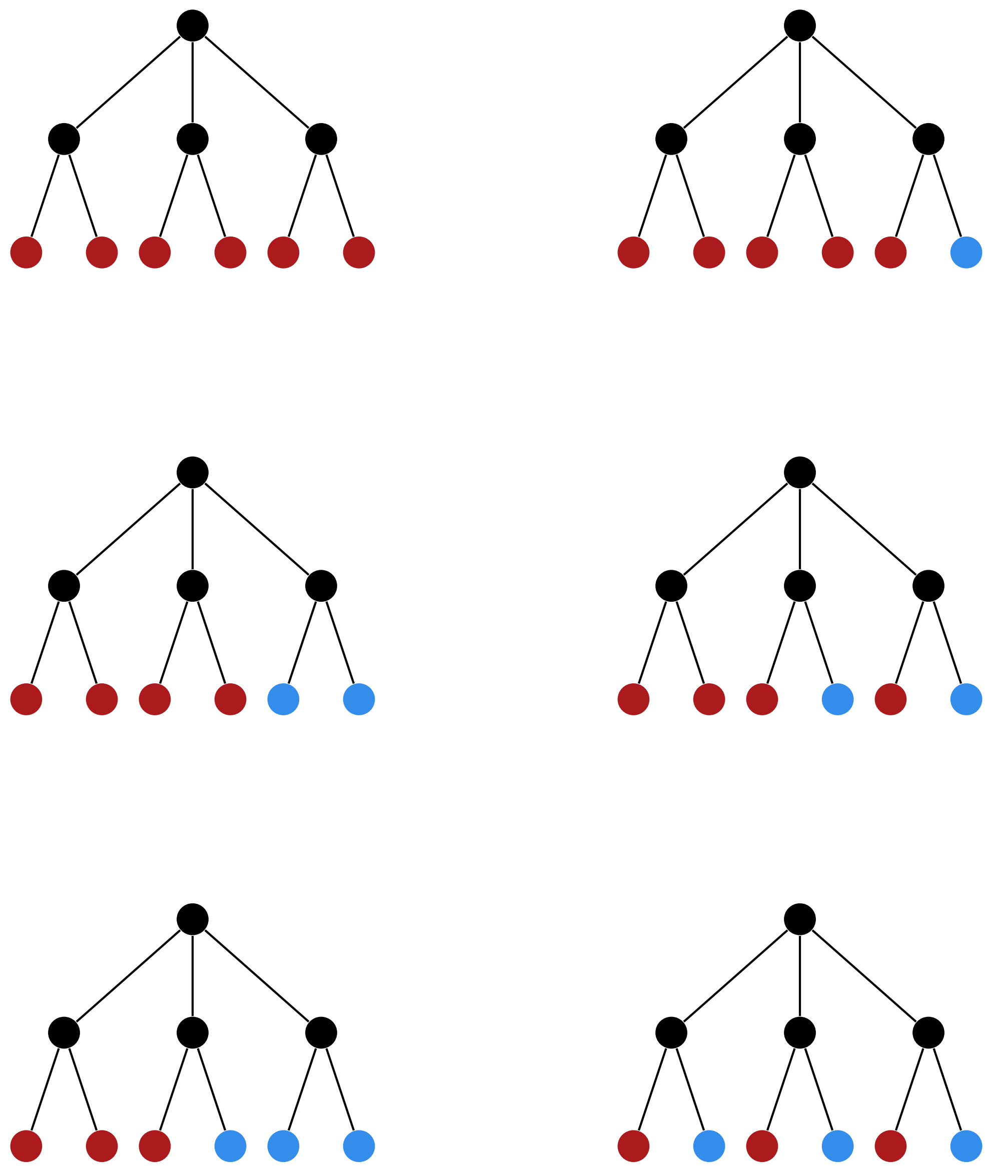

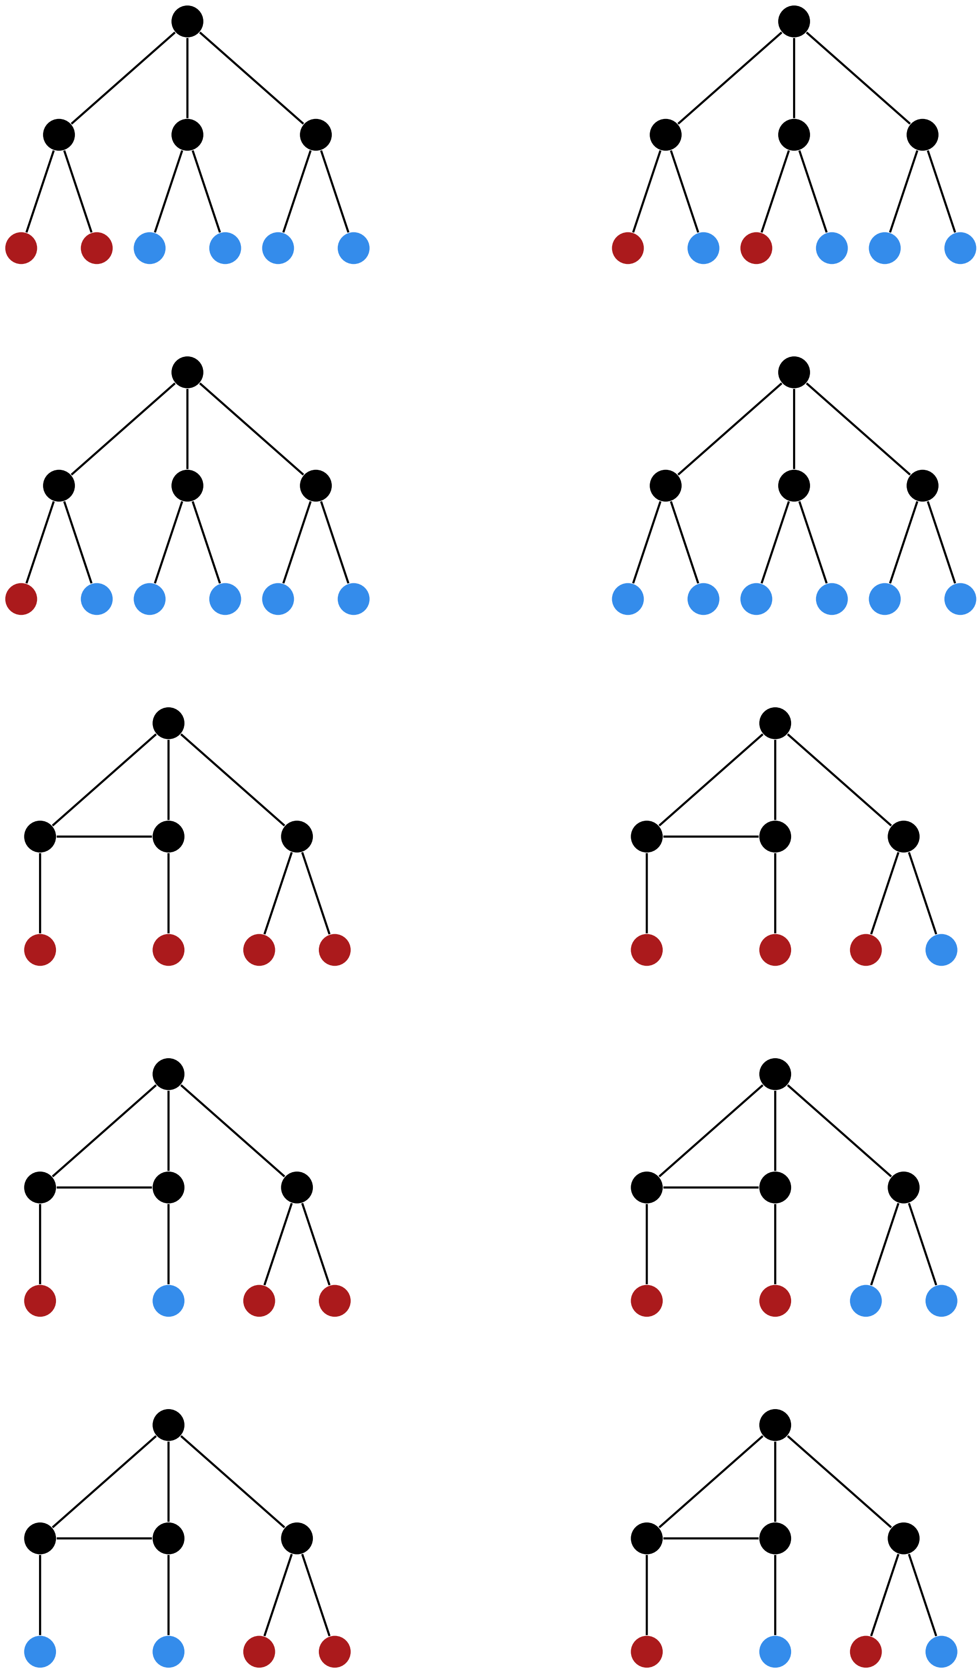

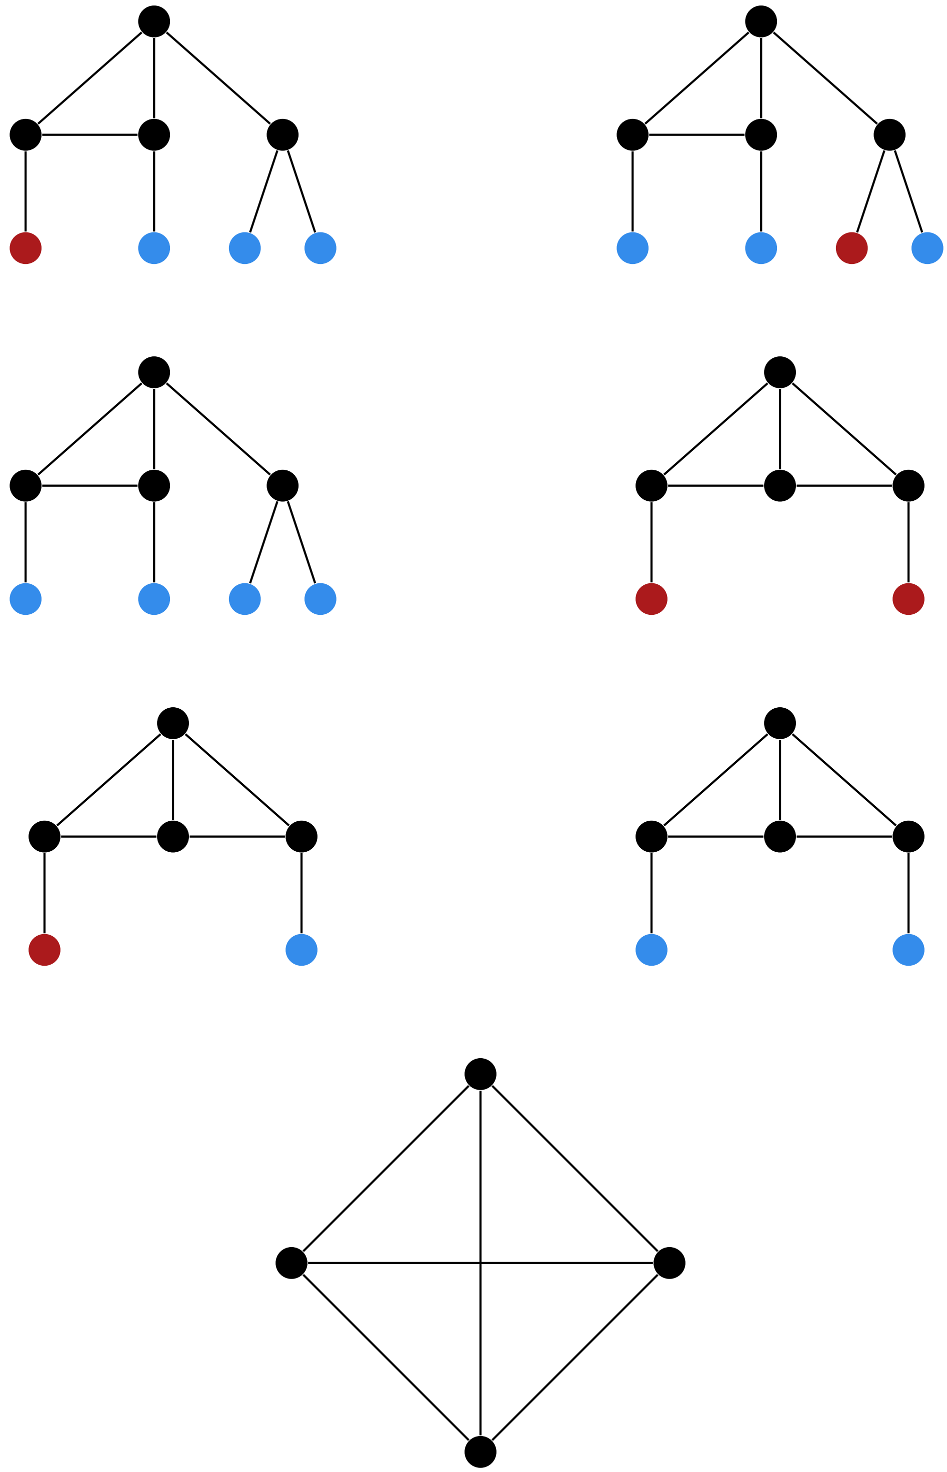

Appendix C: Supplementary Material

The code we used as part of our proofs can be found at github.com/ed359/IsingOccupancy/tree/d3paper, along with instructions. You will require both SageMath (sagemath.org) and Wolfram Engine (www.wolfram.com/engine).