PHASE TRANSITIONS IN THE COMPLEXITY OF COUNTING

A Thesis
Presented to
The Academic Faculty

by

Andreas Galanis

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in
Algorithms, Combinatorics, and Optimization

School of Computer Science
Georgia Institute of Technology
August 2014

Copyright © 2014 by Andreas Galanis

PHASE TRANSITIONS IN THE COMPLEXITY OF COUNTING

Approved by:

Eric Vigoda, Advisor
School of Computer Science
Georgia Institute of Technology

Will Perkins
School of Mathematics
Georgia Institute of Technology

Dana Randall
School of Computer Science
Georgia Institute of Technology

Daniel Štefankovič
Department of Computer Science
University of Rochester

Prasad Tetali
School of Mathematics
Georgia Institute of Technology

Santosh Vempala
School of Computer Science
Georgia Institute of Technology

Date Approved: 2 May 2014

DEDICATION

To my parents, Georgios and Mary,

And my brother, Dimitris

ACKNOWLEDGEMENTS

First and foremost, I want to express my gratitude to my advisor, Eric Vigoda, whose encouragement and attentiveness have been vital in my endeavors. His insightful research pursuits and guidance have made my studies a highly engaging experience, from which I have learned and developed. For these and lot more that are left out, I will always be indebted to him.

I am also deeply grateful to Daniel Štefankovič, who has been to me a co-advisor rather than simply a co-author. His systematic and inquisitive approach to problems has shaped my thinking from year one and I feel very fortunate for this. Pondering over problems with him has invariably been a creative and enjoyful process, which I will always safeguard as a source of inspiration.

I would also like to thank the members of my committee Will Perkins, Dana Randall, Prasad Tetali, and Santosh Vempala. Not only have they caringly monitored my progress throughout my studies, but they have also provided a warm environment to pursue my research interests. I have greatly benefited from discussions, courses and talks, where they shared their expertise on a variety of topics.

I would also like to acknowledge the ACO program for the wide program of studies and for giving me the opportunity to be among such a great group of people. I would also like to thank all my instructors who have helped expand my knowledge in courses I have taken.

I also want to acknowledge my co-authors Qi Ge and Linji Yang for their excellent collaboration; especially my old office mate Linji. I have also greatly benefited from visits to Dimitris Achlioptas, Amin Coja-Oghlan and Juan Vera and I would like to thank them; especially Dimitris, who stimulated my interest in the area during my Master's degree back in Athens.

Special thanks go to my good friends Ioannis Panageas and Georgios Piliouras, with whom I have spent endless hours working together. I would also like to thank all of my

ACO fellow students with whom we have hung around, and especially Abhishek Banerjee, Arindam Khan and Ying Xiao. All those that I unintentionally forget.

Last but not least, my family for supporting me over all these years.

TABLE OF CONTENTS

SUMMARY

A recent line of works established a remarkable connection for antiferromagnetic 2-spin systems, including the Ising and hard-core models, showing that the computational complexity of approximating the partition function for graphs with maximum degree Δ undergoes a computational transition that coincides with the statistical physics uniqueness/non-uniqueness phase transition on the infinite Δ-regular tree. Despite this clear picture for 2-spin systems, there is little known for multi-spin systems. We present the first analog of the above inapproximability results for multi-spin systems.

The main difficulty in previous inapproximability results was analyzing the behavior of the model on random Δ-regular bipartite graphs, which served as the gadget in the reduction. To this end one needs to understand the moments of the partition function. Our key contribution is connecting: (i) induced matrix norms, (ii) maxima of the expectation of the partition function, and (iii) attractive fixed points of the associated tree recursions (belief propagation). We thus obtain a generic analysis of the Gibbs distribution of any multi-spin system on random regular bipartite graphs. We also treat in depth the k-colorings and the q-state antiferromagnetic Potts models.

Based on these findings, we prove that for Δ constant and even $k < \Delta$, it is NP-hard to approximate within an exponential factor the number of k-colorings on triangle-free Δ-regular graphs. We also prove an analogous statement for the antiferromagnetic Potts model. We systematize the approach to obtain a general theorem for the computational hardness of counting in antiferromagnetic spin systems, which we ultimately use to obtain the inapproximability results for the k-colorings and q-state antiferromagnetic Potts models, as well as (the previously known results for) antiferromagnetic 2-spin systems. The criterion captures in an appropriate way the statistical physics uniqueness phase transition on the tree.

CHAPTER I

INTRODUCTION

Over the last two decades, the need to efficiently simulate large-scale random systems has led to an extensive development of approximate sampling techniques. This line of research is concretely captured in the framework of probabilistic graphical models, whose applications can be found in a wide range of areas, including statistical machine learning, bioinformatics and computer vision, to name a few. As a byproduct of the intensive study of graphical models, the classical Markov Chain Monte Carlo approach for approximate sampling has been complemented by elaborate variational methods (see for example [67]). Rapid as the development of approximate sampling techniques might be, the computational limits of these methods and the reasons behind their practical success (or failure) remain for the large part elusive.

Recently, a remarkable connection has been established for *spin systems*, a natural class of graphical models, relating directly the complexity of approximate sampling/counting in computer science and the theory of phase transitions in statistical physics. A crucial part in this connection is in establishing phase transitions for spin systems on random graphs. This thesis is devoted to developing further these ideas and understanding the boundaries of efficient computation.

1.1 Spin systems & Computation: overview

Spin systems are a general framework that captures well-studied models from statistical physics, such as the Ising and Potts models, and also models of combinatorial interest, such as k-colorings and the hard-core lattice gas model defined on independent sets. We next give a general definition of spin systems and later review the most interesting models.

1.1.1 General definition

For an integer $q \geq 2$, a q-spin system is specified by a *symmetric* $q \times q$ *interaction matrix* $\mathbf{B} = (B_{ij})_{i,j \in [q]}$ with non-negative entries, which specify the strength of the interaction between the spins.

For a finite undirected graph $G = (V, E)$, a q-spin system is a probability distribution μ_G over the space Ω_G of all *configurations*, i.e., spin assignments $\sigma : V \to [q]$. The weight of a configuration $\sigma \in \Omega_G$ is the product of neighboring spin interactions, that is,

$$w_G(\sigma) = \prod_{(u,v) \in E} B_{\sigma(u)\sigma(v)}. \tag{1}$$

The spin system has *hard* constraints if there exist $i, j \in [q]$ such that $B_{ij} = 0$; note that every configuration with positive weight does *not* have the spins i, j assigned to adjacent vertices. A *permissive* spin system is a system for which there exists a spin i such that $\min_j B_{ij} > 0$. For example, a system with no hard constraints is permissive. In permissive spin systems, there is always a configuration with positive weight and, more strongly, every partial configuration $\sigma : S \to [q]$ with $S \subseteq V$ has an extension $\bar{\sigma} : V \to [q]$ such that $w_G(\bar{\sigma}) \neq 0$ (provided that for all $u, v \in S$ such that $(u, v) \in E$ it holds that $B_{\sigma(u)\sigma(v)} \neq 0$).

The *partition function* of the spin system is the aggregate weight of the configurations, i.e.,

$$Z_G = \sum_{\sigma \in \Omega_G} w_G(\sigma). \tag{2}$$

The *Gibbs distribution* μ_G is defined as $\mu_G(\sigma) = w_G(\sigma)/Z_G$. In statistical physics, the Gibbs distribution specifies an equilibrium state of the system.

Often, the edge interaction is supplemented by a set of *external fields* which bias the frequency of one or more spins. More precisely, a set of external fields is specified by a q-dimensional vector $\boldsymbol{\lambda} = (\lambda_i)_{i \in [q]}$ with non-negative entries. The weight function (1) is given then by

$$w_G(\sigma) = \prod_{u \in V} \lambda_{\sigma(u)} \prod_{(u,v) \in E} B_{\sigma(u)\sigma(v)}.$$

Note that setting all the λ_i's equal to 1 yields the weight function (1), the so called *no external field* case.

Remark 1. *In the largest portion of this thesis, we will work with Δ-regular graphs. For this class of graphs, there is a generic way to restrict our attention to spin systems with no external fields. In particular, the spin system specified by $\mathbf{B}, \boldsymbol{\lambda}$ is equivalent to a spin system with interaction matrix $\mathbf{B}_{\boldsymbol{\lambda}}$ where the ij-entry of $\mathbf{B}_{\boldsymbol{\lambda}}$ is given by $\lambda_i^{1/\Delta}\lambda_j^{1/\Delta}B_{ij}$.*

For completeness, we next review some well-known models of significant combinatorial or statistical physics interest.

1.1.2 2-spin systems

A 2-spin system is specified by parameters $B_1, B_2 \geq 0$ and $\lambda > 0$. The parameters B_1, B_2 specify the edge interaction, while λ is the external field on the spin 1. To avoid trivial models we will invariably assume that at least one of B_1, B_2 is bigger than 0. The edge interaction matrix for a 2-spin system with parameters B_1, B_2 is given by

$$\mathbf{B} = \begin{bmatrix} B_1 & 1 \\ 1 & B_2 \end{bmatrix}.$$

The system is *antiferromagnetic* if $B_1B_2 < 1$ and *ferromagnetic* if $B_1B_2 > 1$. The classification captures whether neighboring spins are favored to be the same or different.

The two most interesting models of 2-spin systems are the hard-core and Ising models. The hard-core model is an idealized lattice gas model in statistical physics and has only one parameter, the external field λ, which is known as the fugacity or activity of the gas. It corresponds to the special case where $B_1 = 1, B_2 = 0$ (in particular, it is a permissive antiferromagnetic 2-spin system). In the standard formulation of the hard-core model, the spins take 0,1 values and vertices which are assigned spin 1 are called *occupied*, and *unoccupied* otherwise. The external field λ acts on the vertices assigned spin 1. It is straightforward to verify that the Gibbs distribution is supported on configurations which induce an independent set (by looking at the vertices assigned the spin 1). Recall, an independent set is a set of vertices which are mutually non-adjacent. The partition function may be alternatively written as

$$Z_G = \sum_I \lambda^{|I|},$$

where the sum ranges over all independent sets of G. Note that for $\lambda = 1$, the partition function is the total number of independent sets and the Gibbs distribution is the uniform distribution over the set of independent sets.

The Ising model was introduced as a model of ferromagnetism by Lenz, which was studied in Ising's PhD thesis [34]. In the 2-spin setting given above, it corresponds to the particular case $B_1 = B_2 = B$. In the no external field case, the Gibbs distribution of the Ising model can be viewed as a weighted probability distribution on the set of cuts of the graph G. The model is antiferromagnetic when $B < 1$ and ferromagnetic otherwise. Note that when $B < 1$, cuts with larger size are favored, in contrast to the case $B > 1$ where the largest weight configurations are trivially the assignments in which every vertex has the same spin. In the standard formulation of the ferromagnetic Ising model with no external field, the spins take ± 1 values and the Gibbs distribution is given by

$$\mu_G(\sigma) \propto \exp\left(\beta \sum_{(u,v)\in E} \sigma(u)\sigma(v)\right), \tag{3}$$

where $\beta \geq 0$ corresponds to what is known as the inverse temperature. In our setting, $B = \exp(\beta)$. The antiferromagnetic regime is obtained by replacing β with $-\beta$ (and keeping $\beta \geq 0$).

To conclude our review of 2-spin systems, we note that there is no generality to be gained by having extra parameters in the formulation of a 2-spin system, provided that the presence of an external field is allowed.

1.1.3 The k-colorings and the q-state Potts models

We next look at models with more than 2-spins. The most important such example is the k-colorings model, where $k \geq 3$ is the number of colors/spins. Later in this thesis, we will use q instead of k to have a uniform account of spin systems. Recall that a proper k-coloring of a graph G is an assignment of k colors to the vertices of G such that no two adjacent vertices have the same color.

The Gibbs distribution of a graph G in the k-colorings model is the uniform distribution over the set of proper k-colorings of the graph G with k-colors. In particular, the partition function is the total number of proper k-colorings. In contrast to 2-spin systems, note that the Gibbs distribution is not always well-defined if the graph G does not admit a proper k-coloring. This is clearly not an issue when k is greater or equal than the chromatic number $\chi(G)$, however computing the latter is in general an NP-hard problem. For graphs of maximum degree Δ, the well-known Brooks' theorem asserts that $\chi(G) \leq \Delta$ unless the graph is complete or an odd cycle. Our interest will be in the class of graphs G whose maximum degree Δ is a constant (independent of the size of the graph), for which better algorithmic bounds on the chromatic number are known. We postpone a review of the literature regarding the colorability of graphs with maximum degree Δ till Section 1.4.2.

The q-state Potts model is a generalization of the k-colorings model. Precisely, the model has q spins and a parameter $B > 0$ (which roughly corresponds to the temperature). The weight of a configuration σ is $B^{m(\sigma)}$, where $m(\sigma)$ is the number of monochromatic edges under the configuration σ, that is, edges (u, v) with $\sigma(u) = \sigma(v)$. The model is antiferromagnetic when $B < 1$ and antiferromagnetic otherwise. Note that the case $B = 0$ corresponds to the k-colorings model. The edge interaction matrix of the q-state Potts model has off-diagonal entries equal to 1, and on the diagonal equal to B. Finally, observe that the Ising model is the particular case of the Potts model with $q = 2$.

1.1.4 Counting in spin systems

From a computational perspective, the notion of the partition function naturally gives rise to the problem of computing it. The partition function captures the macroscopic properties of a spin system and is important to run simulations. For example, it is well known that an FPRAS for the partition function is equivalent to approximately sampling from the Gibbs distribution (see [39] for more details).

Formally, we will be interested in the following counting problem.

Parameters. integer $q \geq 2$, $q \times q$ symmetric matrix $\mathbf{B}$ with non-negative entries.

Name. #PartitionFunction$(q, \mathbf{B})$.

Input. A graph G.

Output. The partition function Z_G for the q-spin system specified by the interaction matrix $\mathbf{B}$.

The input of #PartitionFunction$(q, \mathbf{B})$ will often be restricted to Δ-regular graphs or graphs with maximum degree Δ, where $\Delta \geq 3$ is a fixed constant. This will be clear from context.

From a complexity viewpoint, #PartitionFunction$(q, \mathbf{B})$ belongs to the complexity class #P (introduced in [66]), which is the class of function problems associated to counting versions of decision problems in the class NP [1]. The problem #PartitionFunction$(q, \mathbf{B})$ is #P-complete unless the matrix $\mathbf{B}$ has rank 1 [12]. When $\mathbf{B}$ has $0, 1$ entries, this is known to hold even for graphs of bounded degree [24], and it is reasonable to expect that #PartitionFunction$(q, \mathbf{B})$ is #P-complete even when restricted to graphs of maximum degree Δ, for all $\Delta \geq 3$. We will thus be interested in efficient approximations of #PartitionFunction$(q, \mathbf{B})$; below we give a formal description of the class of algorithms that will be of interest to us.

Let $f : \Sigma^* \to \mathbb{R}$, where f is a function we are interested in computing and Σ^* is an encoding of the instances. A *polynomial-time randomized approximation scheme* (PRAS) for f is a randomized algorithm that takes as input a pair $(x, \varepsilon) \in \Sigma^* \times (0, 1)$ and outputs in polynomial time (in the length $|x|$ of the input) the value of a random variable Y supported on rational numbers which satisfies $\Pr[|Y - f(x)| \leq \varepsilon |f(x)|] \geq 3/4$. A *fully polynomial-time randomized approximation scheme* (FPRAS) for f has the same guarantees as a PRAS, but the algorithm is required to run in polynomial time both in the length $|x|$ and $1/\varepsilon$. A PTAS and an FPTAS are the deterministic analogues of PRAS and FPRAS, respectively.

Later, we will also need to compare the relative complexity of approximating the partition functions in two different spin systems. This notion was introduced in [23]. Let $f, g : \Sigma^* \to \mathbb{R}$. An *approximation-preserving reduction* ($\mathsf{AP-reduction}$) from f to g is an FPRAS for f that has oracle access to an FPRAS for g; moreover, on input $(x, \varepsilon) \in \Sigma^* \times (0, 1)$, all oracle calls to g should be pairs $(y, \delta) \in \Sigma^* \times (0, 1)$, where $|y|$ is polynomial in $|x|, \varepsilon^{-1}$ and δ is polynomial in $|x|, \varepsilon$.

[1] We assume familiarity with standard complexity classes, such as P, NP, RP. The reader is referred to [58] in the literature.

1.1.5 A sharp computational transition for the hard-core model

A striking computational transition was established for the problem of approximating the partition function in the hard-core model. To describe this computational transition, for $\Delta \geq 3$, let $\lambda_c(\mathbb{T}_\Delta) := (\Delta - 1)^{\Delta-1}/(\Delta - 2)^\Delta$ (the notation $\lambda_c(\mathbb{T}_\Delta)$ will become apparent shortly).

Weitz [68] gave an FPTAS for the partition function of graphs with maximum degree Δ when Δ is constant and $\lambda < \lambda_c(\mathbb{T}_\Delta)$. On the other hand, Sly [63] proved (extended in [26, 27, 64], see also Theorem 11 in this thesis) that, unless $\mathsf{NP} = \mathsf{RP}$, for every $\Delta \geq 3$ and $\lambda > \lambda_c(\mathbb{T}_\Delta)$, there does not exist an FPRAS for the partition function on graphs with maximum degree Δ. In fact, it was proved in [64] that the intractability result remains true even within an exponential factor of approximation. As an interesting special case of these results, observe that $\lambda_c(\mathbb{T}_\Delta) > 1$ iff $\Delta \leq 5$, and thus approximately counting independent sets on graphs with maximum degree Δ admits a polynomial-time algorithm iff $\Delta \leq 5$ (assuming $\mathsf{NP} \neq \mathsf{RP}$).

This sharp computational transition is even more interesting in light of the fact that it coincides with the *uniqueness/non-uniqueness phase transition* on the infinite Δ-regular tree $\mathbb{T}_\Delta$ in statistical physics. Roughly, the uniqueness/non-uniqueness phase transition captures the existence of long-range correlations; we will formulate this precisely later in the introduction. For now, we note that the techniques and results of [68] and [63] established for the first time a strong connection between computational complexity and phase transitions in statistical physics.

This connection was later extended to antiferromagnetic 2-spin systems [44, 64] (for the hardness side, see also Theorem 11). Despite this beautiful picture for 2-spin systems, much less is known for multi-spin systems such as the k-colorings or the q-state Potts models. Our major goal in this thesis is to obtain an analog of the above inapproximability results for multi-spin systems, by exploring the connection with the phase transition on the infinite Δ-regular tree.

More precisely, the main difficulty in the previous inapproximability results for 2-spin systems was to analyze the Gibbs distribution on random bipartite Δ-regular graphs, which are used as gadgets in the reduction. This reduction scheme was first introduced by Dyer, Frieze, and Jerrum [22], who obtained that approximately counting independent sets on graphs of maximum degree Δ is hard whenever $\Delta \geq 25$. The connection of the Gibbs distribution on random bipartite Δ-regular graphs with the phase transition on the infinite Δ-regular tree was formulated and established by Mossel, Weitz, and Wormald [56], who conjectured that approximating the partition function in the hard-core model should be hard whenever $\lambda > \lambda_c(\mathbb{T}_\Delta)$. In a seminal work, Sly [63] proved the conjecture when $\lambda_c(\mathbb{T}_\Delta) < \lambda < \lambda_c(\mathbb{T}_\Delta) + \varepsilon(\Delta)$ for some small $\varepsilon(\Delta) > 0$ and the interesting case $\lambda = 1, \Delta = 6$, introducing a reduction scheme which allowed to go all the way to the threshold.

To give a flavor of our later pursuits, it will be useful to briefly go over the approach in [62]. Sly utilized the results in [56] to construct a gadget based on a bipartite random graph (the gadget itself is a bipartite graph). Roughly, when $\lambda > \lambda_c(\mathbb{T}_\Delta)$, the largest contribution to the partition function of the gadget comes from configurations which have a linear (in the size of the gadget) surplus of occupied vertices on one side of the bipartition (the configuration space splits evenly with respect to this property between the two sides). The gadget thus exhibits a boolean behavior; a configuration σ in the Gibbs distribution is with high probability in one of two "phases", depending on the side of the bipartition which has the largest number of occupied vertices. By making clever connections between the sides of two copies of the gadget yields a graph whose partition function is dominated by the contribution of configurations whose phases on the two gadgets are different. Using a sufficiently large gadget, the construction can be carried out for an arbitrary graph H (rather than a single edge) by replacing each vertex of H by a copy of the gadget. The partition function of the resulting graph is dominated by configurations whose phases on neighboring gadgets are as different as possible. Specifically, the phases of the gadgets correspond to a maximum cut partition of H, yielding a reduction of MaxCut to the problem of approximating the partition function in the hard-core model.

Sly's reduction required substantial technical work, at the heart of which lies the analysis of the hard-core model on random bipartite Δ-regular graphs. The latter analysis is based heavily on insights from phase transitions on the infinite Δ-regular tree. To get inapproximability results for multi-spin systems, we will follow a similar approach, seeking to analyze general spin systems on random bipartite Δ-regular graphs. En route, we will describe the uniqueness phase transition (Section 1.2) and how it manifests itself on random regular graphs (Section 1.3). These pieces will be combined in Section 1.4 to obtain our inapproximability results.

1.2 Phase transitions on lattices

From the viewpoint of statistical physics, a phase transition is an abrupt change in the macroscopic properties of a system due to small changes in the microscopic parameters of the system. A well-known example of a phase transition is the transformation of the states of matter (solid, liquid, gas). To better align however with the mathematical development of phase transitions, it will be useful to describe the phase transition of the Ising model on the two dimensional integer lattice $\mathbb{Z}^2$ (which we will abbreviate in short as 2D Ising). We will work in the standard notation of the Ising model, given in (3).

To motivate the discussion, let us first observe that the Gibbs distribution is well defined only for finite graphs. In statistical physics however, systems have typically a large number of components, effectively infinite. This raises the question of how to describe the equilibrium states in an infinite system as an extension of the (finite) Gibbs distribution. Ideally, the description of such an equilibrium would capture that finite subsystems are in equilibrium with the rest of the system. Let us formalize the last sentence.

Consider the 2D Ising model on $\mathbb{Z}^2$. Let Λ_n be a square box around the origin of $\mathbb{Z}^2$ with length $\lfloor\sqrt{n}\rfloor$. The set of vertices in $\mathbb{Z}^2\backslash\Lambda_n$ which have neighbors in Λ_n will be called the (external) boundary of Λ_n and denoted as $\partial\Lambda_n$. We will consider configurations which assign the same spin to the boundary vertices. Specifically, the (+)-boundary condition will refer to the set of configurations $\sigma : \Lambda_n \cup \partial\Lambda_n \rightarrow \{\pm 1\}$ where all vertices in $\partial\Lambda_n$ are assigned spin $+1$. Define analogously the $(-)$-boundary condition. Next, we will examine whether the boundary condition affects the marginal probability that the origin is assigned the spin $+1$. To do this, denote the origin by ρ and the Gibbs distribution on $\Lambda_n \cup \partial\Lambda_n$ by

$\mu_n := \mu_{\Lambda_n \cup \partial \Lambda_n}$. Let

$$p_n^+ := \mu_n(\sigma(\rho) = +1 \,|\, \forall u \in \partial\Lambda_n, \sigma(u) = +1),$$
$$p_n^- := \mu_n(\sigma(\rho) = +1 \,|\, \forall u \in \partial\Lambda_n, \sigma(u) = -1).$$

Given that the model is ferromagnetic, it is reasonable to expect that $p_n^+ > p_n^-$. The phase transition in $\mathbb{Z}^2$ captures whether the bias persists in the limit $n \to \infty$. More precisely, if

$$p_n^+ - p_n^- \to 0, \tag{4}$$

the model is in the *uniqueness regime* of $\mathbb{Z}^2$, and otherwise in the *non-uniqueness regime*. To justify this terminology, (4) captures whether there are multiple equilibrium states of the infinite system. In particular, whenever (4) holds, it can be proved that finite subregions which are far away are asymptotically independent. It is reasonable to expect that this would hold when the edge interaction is weak, i.e., when β is small. On the other hand, if β is large, then (4) fails and the system is characterized by long range correlations, signifying the existence of multiple equilibrium states. In fact, the transition from the uniqueness regime to the non-uniqueness regime is sharp: Onsager [57] solved the model and proved the existence of a critical temperature $\beta_c(\mathbb{Z}^2)$ such that whenever $\beta > \beta_c(\mathbb{Z}^2)$ the model is in the non-uniqueness regime, and in uniqueness otherwise.

To summarize, the phase transition on $\mathbb{Z}^2$ captures the existence of long range correlations in the large scale limit; whenever the correlations persist, the system has multiple equilibrium states. So far, we have been intentionally obscure on the definition of the equilibrium states in infinite systems, since the mathematical formulation is more strenuous. An equilibrium state in $\mathbb{Z}^2$ is a probability distribution on the space of all configurations, known as an *infinite volume Gibbs measure*. Roughly, an infinite volume Gibbs measure is consistent with all Gibbs distributions on finite subgraphs and all boundary conditions on them. This should bear some resemblance with the argument above; a formal definition and further discussion is given in Section 1.2.1.

We conclude this part by noting that there are other ways to view the phase transition on $\mathbb{Z}^2$: the classical statistical physics view concentrates on singularities of the limit of the logarithm of the partition function or discontinuities of macroscopic properties of the system.

1.2.1 Infinite-volume Gibbs measures

Dobrushin [21] and Lanford and Ruelle [43] formulated the concept of an infinite-volume Gibbs measure. This is a natural extension of the Gibbs distribution in the setting of infinite graphs and can be therefore interpreted as an equilibrium state of an infinite system. In the previous section, we described the motivation behind the definition of an infinite-volume Gibbs measure. In this section, we shall give a formal definition for a q-spin system with interaction matrix $\mathbf{B}$. For a more elaborate treatment, the reader should refer to [29].

Let $G = (V, E)$ be a graph, where V is a countable set of vertices. For convenience, we will assume that the graph is locally finite, i.e., every vertex has bounded degree. A configuration $\sigma : V \to [q]$ will be called feasible if for every edge $(u, v) \in E$, it holds that $B_{\sigma(u)\sigma(v)} > 0$. For a configuration σ and a subset $S \subseteq V$, the restriction of σ to S will be denoted by σ_S.

Let Λ be a finite subgraph of G and denote by $\partial\Lambda$ its external boundary, i.e., the set of vertices which do not belong in Λ but are adjacent to a vertex in Λ. With a minor abuse of notation, for a configuration σ, we will denote by σ_Λ the restriction of σ to the set of vertices in Λ. We will denote by $\mu_{\Lambda\cup\partial\Lambda}$ the finite volume Gibbs distribution on $\Lambda \cup \partial\Lambda$.

A Gibbs measure ν is defined by requiring ν to agree with the finite Gibbs distribution inside every finite region Λ, given that the configuration outside Λ is kept fixed. More precisely, we have the following definition.

Definition 1. *Let $\mathbf{B}$ be the interaction matrix of a q-spin system and $G = (V, E)$ be a locally finite graph. A probability measure ν over the space of all feasible configurations on G is a Gibbs measure for the spin model defined by $\mathbf{B}$, if for every finite region $\Lambda \subseteq G$, and ν-almost every feasible configuration η, it holds that*

$$\nu(\sigma_\Lambda = \tau \,|\, \sigma_{G\setminus\Lambda} = \eta_{G\setminus\Lambda}) = \mu_{\Lambda\cup\partial\Lambda}(\sigma_\Lambda = \tau \,|\, \sigma_{\partial\Lambda} = \eta_{\partial\Lambda}),$$

for every $\tau : \Lambda \to [q]$.

With the definition of an infinite-volume Gibbs measure in place, we are now ready to give a formal definition of uniqueness/non-uniqueness. The definition captures the existence of multiple equilibrium states for the infinite system.

Definition 2. *Let* $\mathbf{B}$ *be the interaction matrix of a* q*-spin system and* $G = (V, E)$ *be a locally finite graph. The* q*-spin system is in the uniqueness regime of* G *if there exists a unique infinite-volume Gibbs measure. Otherwise, the* q*-spin system is in the non-uniqueness regime of* G*.*

To relate with the definition of the uniqueness regime in the 2D Ising model, we give the following equivalent definition of uniqueness, whose proof can be found in e.g. [29].

Lemma 1. *Let* $\mathbf{B}$ *be the interaction matrix of a* q*-spin system and* $G = (V, E)$ *be a locally finite graph. The* q*-spin system is in the uniqueness regime of* G *iff for every finite subgraph* $\Lambda \subseteq G$*, there exists a sequence of finite regions* $\Lambda \subseteq \Lambda_1 \subseteq \Lambda_2 \subseteq \cdots \subseteq \Lambda_n \subseteq \dots$ *with* $\cup_n \Lambda_n = V$ *such that for every pair of feasible configurations* η, η' *on* G*, it holds that*

$$\left|\mu_{\Lambda_n \cup \partial\Lambda_n}(\sigma_\Lambda = \tau \,|\, \sigma_{G\setminus\Lambda_n} = \eta_{G\setminus\Lambda_n}) - \mu_{\Lambda_n \cup \partial\Lambda_n}(\sigma_\Lambda = \tau \,|\, \sigma_{G\setminus\Lambda_n} = \eta'_{G\setminus\Lambda_n})\right| \to 0$$

for all $\tau : \Lambda \to [q]$*.*

Equation (4) for the 2D Ising model captures the condition in Lemma 1, in the case that Λ is the origin of $\mathbb{Z}^2$ and η, η' are the configurations on $\mathbb{Z}^2$ where all vertices are assigned $+1, -1$, respectively. That for the ferromagnetic Ising model it suffices to look only at these two extreme configurations follows from a correlation inequality known as Griffith's inequality.

1.2.2 A primer on the uniqueness threshold on the Bethe lattice

Let $\Delta \geq 3$. The infinite Δ-regular tree, denoted by $\mathbb{T}_\Delta$, is known in statistical physics as the Bethe lattice with coordination number Δ.

The uniqueness threshold on $\mathbb{T}_\Delta$ will be crucial for our arguments. $\mathbb{T}_\Delta$ has been conjectured to be the worst case graph of maximum degree Δ for long range correlations. The reason being that the influence from boundary conditions should be maximized when the number of vertices in the boundary is maximized. In this sense, $\mathbb{T}_\Delta$ has the most vertices at distance ℓ from the root. While the conjecture has been proved to be wrong in general [62], for natural spin systems such as k-colorings or the antiferromagnetic Potts model it is believed to be true. In fact, for 2-spin systems the conjecture has been shown to be true in [68] in the context of the hard-core model and was later extended in [61, 44] for antiferromagnetic 2-spin systems. Later, the reader will find that our inapproximability results hinge upon the long-range correlation on the tree.

Determining the uniqueness threshold on $\mathbb{T}_\Delta$ is generally simpler than on other graphs, due to its acyclic structure. However, technical difficulties are not absent; for example, it is surprising to some extent that the uniqueness threshold for the antiferromagnetic Potts model is not known, despite that it is conjectured to have the particularly simple form $B_c(\Delta) = \frac{\Delta - q}{\Delta}$.

We next give a flavor of the arguments that come into play when determining the uniqueness threshold on the tree, to help motivating the upcoming sections. We do this in the simplest possible case, the hard-core model with activity λ. Note, the uniqueness threshold for the hard-core model has been found by Kelly [42] to be

$$\lambda_c(\mathbb{T}_\Delta) = \frac{(\Delta - 1)^{(\Delta-1)}}{(\Delta - 2)^\Delta},$$

so we will focus on those parts that will be most important for us.

Following the same approach as in the 2D Ising model, we examine the Gibbs distribution for the hard-core model on finite trees of depth n. It will be more convenient to look at complete $(\Delta - 1)$-ary trees of depth n, where the root has degree $\Delta - 1$ (this is to avoid accounting for the root separately, otherwise nothing really changes). We denote the depth n tree with $\mathbb{T}_n$ and its set of leaves by S_n. We will examine two extreme boundary configurations on the leaves which, due to certain monotonicities of the hard-core model on the tree $\mathbb{T}_\Delta$ (or more generally a bipartite graph), turn out to have the largest influence on the root: the $(+)$-boundary, where all the leaves are occupied, and the $(-)$-boundary, where all the leaves are unoccupied. Let p_n^+, p_n^- be the marginal probabilities that the root of $\mathbb{T}_n$ is occupied conditioned on the $(+)$, $(-)$ boundary conditions respectively. Formally, if ρ denotes the root of the tree and $\mu_n := \mu_{\mathbb{T}_n}$ the Gibbs distribution on $\mathbb{T}_n$, let

$$p_n^+ := \mu_n(\sigma(\rho) = 1 \,|\, \forall u \in S_n, \sigma(u) = 1), \quad p_n^- := \mu_n(\sigma(\rho) = 1 \,|\, \forall u \in S_n, \sigma(u) = 0).$$

The hard-core model with activity λ is then in the uniqueness regime of $\mathbb{T}_\Delta$ if $p_n^+ - p_n^- \to 0$ as $n \to \infty$. We note here that, in contrast to the Ising model, it does *not* hold that $p_n^+ > p_n^-$ for all n; instead, it holds that

$$p_{2n}^+ > p_{2n}^- \text{ and } p_{2n+1}^+ < p_{2n+1}^-. \tag{5}$$

This is a consequence of the fact that the hard-core model is an antiferromagnetic 2-spin system. Unlike to the case of $\mathbb{Z}^2$, we can get easily a handle on p_n^+, p_n^- by writing tree recursions. One thus obtains the following recursions

$$p_{n+1}^+ = f(p_n^+), \quad p_{n+1}^- = f(p_n^-), \text{ where } f(x) := 1 - \frac{1}{1 + \lambda(1-x)^{\Delta-1}}. \tag{6}$$

The initial conditions for the recursions are given by $p_0^+ = 1$ and $p_0^- = 0$. It can easily be checked that the fixed point equation $x = f(x)$ has a unique solution x_o for all values of λ. When $\lambda \leq \lambda_c(\mathbb{T}_\Delta)$, the fixed point (or in short fixpoint) x_o is stable and both p_n^+, p_n^- converge to x_0. However, when $\lambda > \lambda_c(\mathbb{T}_\Delta)$, the fixpoint x_o is unstable, i.e., it holds that $f'(x_o) < -1$. This implies that the sequences p_n^+, p_n^- do not converge: otherwise, their limit would have to be equal to x_o, but close to x_o the recursion oscillates.

To better understand what is happening in the regime $\lambda > \lambda_c(\mathbb{T}_\Delta)$, (5) suggests looking at the recursion (6) for two steps, i.e.,

$$p_{2n+2}^+ = g(p_{2n}^+), \quad p_{2n+2}^- = g(p_{2n}^-), \text{ where } g(x) := f(f(x)),$$

and $p_0^+ = 1$ and $p_0^- = 0$. The fixpoint equation $x = g(x)$ has now three fixpoints $x^+ > x_o > x^-$ in the regime $\lambda > \lambda_c(\mathbb{T}_\Delta)$. The unstable fixpoint x_o of f remains unstable for g as well, i.e., $g'(x_o) > 1$. The remaining two fixpoints x^+, x^- satisfy $|g'(x)| < 1$ and are hence stable. Combining the above observations, one can easily see that $p_{2n}^+ \downarrow x^+$ and $p_{2n}^- \uparrow x^-$. This suggests the existence of (at least) two different infinite-volume Gibbs measures in the regime $\lambda > \lambda_c(\mathbb{T}_\Delta)$. The first is obtained by fixing the leaves on T_{2n} to be occupied and taking the weak limit as $n \to \infty$; the marginal probability of the root being occupied is x^+

in the limiting measure. The second is obtained by fixing the leaves on T_{2n} to be unoccupied and taking again the weak limit as $n \to \infty$; in this case, the marginal probability of the root being occupied is x^- in the limiting measure. A more explicit construction of these two infinite-volume Gibbs measures is given in Section 1.2.3.1.

We now summarize our observations. First, we saw that one can get a handle on the uniqueness threshold on the tree by looking at appropriate tree recursions. In the case of the hard-core model, due to the (anti)monotonicity properties that the model exhibits, we had to go up to depth two only and look at the two extreme boundary conditions (in general this need not be the case). Second, the asymptotic stability of the fixpoints provided us with useful information on the location of the phase transition and helped us to identify infinite-volume Gibbs measures.

In the next section, we will see how to generalize in an appropriate way these observations to a general q-spin system.

1.2.3 Translation invariant Gibbs measures on the Bethe lattice

In general, infinite-volume Gibbs measures can be extremely complex and particularly hard to define. However, on lattice graphs it is usually the case that a natural class of Gibbs measures can be identified which are invariant under a subgroup of the automorphism group of the graph. In this section, we will view two such general constructions on the infinite Δ-regular tree. These constructions will naturally lead us to the notion of semi-translation invariant uniqueness, which is a weaker notion of uniqueness, albeit easier to formulate. Most importantly, this auxiliary uniqueness threshold will manifest itself in our later investigation of random Δ-regular graphs, allowing us to connect the phase transition on the infinite Δ-regular tree with properties of the Gibbs distribution in random graphs.

Recall from Section 1.2.2 that the Bethe lattice is an (infinite) regular tree. We will denote the infinite Δ-regular tree by $\mathbb{T}_\Delta$. To ease exposition, it will also be helpful to specify a root ρ on $\mathbb{T}_\Delta$. A translation invariant Gibbs measure on $\mathbb{T}_\Delta$ is a measure ν on the space of configurations which is invariant under all automorphisms of $\mathbb{T}_\Delta$[2]. Among other things, this

[2] An automorphism g of $\mathbb{T}_\Delta$ is either a rotation (there exists a vertex v such that $gv = v$), an inversion (there exists an edge (u, v) such that $gu = v$ and $gv = u$) or a translation. For more details, see e.g., [65, 16].

implies that under ν, the marginal spin distribution of each vertex of the tree is identical, i.e., the distribution $\nu(\sigma(v) = \cdot)$ does not depend on the vertex v. As we shall display at the end of this section, translation invariant Gibbs measures correspond to fixpoints of the *depth-one tree recursions*:

$$\widehat{R}_i \propto \Big(\sum_{j=1}^{q} B_{ij} R_j\Big)^{\Delta-1}. \tag{7}$$

The fixpoints of the tree recursions (7) are those $\mathbf{r} = (R_1, \ldots, R_q)$ such that:

$$\widehat{R}_i \propto R_i \text{ for all } i \in [q].$$

Remark 2. *If $\mathbf{r} = (R_1, \ldots, R_q)$ is a fixpoint of (7), so is $c\mathbf{r} = (cR_1, \ldots, cR_q)$ for any $c > 0$. In other words, the recursion (7) is scale-free.*

Before proceeding, for the sake of a concrete example, we align the recursions (7) with the recursion (6) for the hard-core model. Recall that the interaction matrix of the hard-core model is given by $\left[\begin{smallmatrix} 1 & 1 \\ 1 & 0 \end{smallmatrix}\right]$. Hence, (7) takes the form:

$$\widehat{R}_1 \propto (R_1 + R_2)^{\Delta-1}, \quad \widehat{R}_2 \propto \lambda R_1^{\Delta-1}.$$

Note here the minor change to account for the presence of the external field, one can alternatively start from Remark 1 to derive the same recursion but this requires slightly more work. If we set $\hat{p} = \widehat{R}_2/(\widehat{R}_1 + \widehat{R}_2)$ and $p = R_2/(R_1 + R_2)$, we obtain $\hat{p} = f(p)$, exactly as in (6).

As we saw in Section 1.2.2, for the hard-core model we needed to look at two steps of the tree recursions to get the correct limiting behavior, as a consequence of the fact that it is an antiferromagnetic spin system. It is natural thus to also examine the Gibbs measures arising from fixpoints after iterating two steps of the tree recursions.

We may define a semi-translation invariant Gibbs measure on $\mathbb{T}_\Delta$ as a measure ν which is invariant under parity-preserving automorphisms of $\mathbb{T}_\Delta$; the marginal spin distribution of a vertex of the tree depends now on the parity of its distance from the root of the tree. Semi-translation invariant Gibbs measures correspond to fixpoints of the *depth-two tree recursions*:

$$\widehat{R}_i \propto \Big(\sum_{j=1}^{q} B_{ij} C_j\Big)^{\Delta-1} \quad \text{and} \quad \widehat{C}_j \propto \Big(\sum_{i=1}^{q} B_{ij} R_j\Big)^{\Delta-1}. \tag{8}$$

The fixpoints of the tree recursions (8) are those pairs $(\mathbf{r}, \mathbf{c})$ with $\mathbf{r} = (R_1, \ldots, R_q)$ and $\mathbf{c} = (C_1, \ldots, C_q)$ such that:

$$\widehat{R}_i \propto R_i \text{ for all } i \in [q] \ , \quad \widehat{C}_j \propto C_j \text{ for all } j \in [q].$$

Remark 3. *Analogously to Remark 2, observe that if* $(\mathbf{r}, \mathbf{c})$ *is a fixpoint of* (8)*, so is* $(c_1\mathbf{r}, c_2\mathbf{c})$ *for any* $c_1, c_2 > 0$*. In other words, the recursion* (8) *is scale-free.*

Definition 3. *A fixpoint of* (8) *will be called translation invariant iff* $R_i \propto C_i$ *for all* $i \in [q]$*.*

Following the observations in Section 1.2.2, it will be important to have the notion of stability for a fixpoint of the tree recursions. The following definition captures whether the tree recursions (7), (8) are stable under small perturbations around a fixpoint. The formal definition uses the Jacobian matrix of a vector valued function $f : \mathbb{R}^n \to \mathbb{R}^n$. The Jacobian of f at a point $\mathbf{x} \in \mathbb{R}^n$ is the $n \times n$ matrix whose i, j entry is the partial derivative of the i-th component of f with respect to x_j (the j-th component of $\mathbf{x}$).

In this languange, our observations in Section 1.2.2 translate as follows in the case of the hard-core model: in the non-uniqueness regime, the Jacobian of the depth-one tree recursions at the unique fixpoint has spectral radius greater than one, while the Jacobian of the depth-two tree recursions at each of the other two fixpoints has spectral radius less than one. We capture this in a more general setting as follows.

Definition 4. *A fixpoint* $\mathbf{r} = (R_1, \ldots, R_q)$ *of the depth-one tree recursions* (7)*, viewed as a function* $f : (R_1, \ldots, R_q) \mapsto (\widehat{R}_1, \ldots, \widehat{R}_q)$*, is* Jacobian attractive *if the Jacobian of* f *at* $\mathbf{r}$ *has spectral radius less than* 1*.*

Similarly, a fixpoint $(\mathbf{r}, \mathbf{c})$ *with* $\mathbf{r} = (R_1, \ldots, R_q), \mathbf{c} = (C_1, \ldots, C_q)$ *of the depth-two tree recursions* (8)*, viewed as a function* $f : (R_1, \ldots, R_q, C_1, \ldots, C_q) \mapsto (\widehat{R}_1, \ldots, \widehat{R}_q, \widehat{C}_1, \ldots, \widehat{C}_q)$*, is* Jacobian attractive *if the Jacobian of* f *at* $(\mathbf{r}, \mathbf{c})$ *has spectral radius less than* 1*.*

We next define the notion of semi-translation invariant uniqueness on $\mathbb{T}_\Delta$. This is a simpler notion of uniqueness, which captures whether there are multiple semi-translation invariant Gibbs measures. In turn, the latter corresponds to checking whether there are multiple fixpoints to the tree recursions (8).

Definition 5. *A spin system with interaction matrix* $\mathbf{B}$ *is in the* semi-translation invariant uniqueness *regime of* $\mathbb{T}_\Delta$*, if there is a unique fixpoint to the tree recursions* (8)*. Otherwise, the model is in the* semi-translation invariant non-uniqueness *regime of* $\mathbb{T}_\Delta$*.*

We now discuss the interplay between Definitions 4 and 5. It can be proved by a variational argument (described in Section 1.3) that a translation invariant fixpoint, i.e., a fixpoint of the recursions (7), always exists. In the uniqueness regime of $\mathbb{T}_\Delta$, the fixpoints of both the recursions (7) and (8) should correspond to the same infinite-volume Gibbs measure. On the other hand, if there exists a fixpoint of either (7) or (8) which is not stable, then we will prove (see Theorem 2) that there must exist another fixpoint of (7) or (8), signifying that the spin system is in the non-uniqueness regime of $\mathbb{T}_\Delta$.

It should be noted at this point that *a spin system may be in the semi-translation invariant uniqueness regime of* $\mathbb{T}_\Delta$*, but not in the uniqueness regime of* $\mathbb{T}_\Delta$, a concrete example will be given shortly. However, for the most interesting models, the relevant thresholds are typically close if not identical. In particular, for 2-spin systems, the semi-translation invariant non-uniqueness threshold coincides with the uniqueness threshold. For the antiferromagnetic Potts model, the uniqueness threshold is conjectured to be $B_c(\mathbb{T}_\Delta) = \frac{\Delta-q}{\Delta}$ and we will prove that this is the semi-translation invariant non-uniqueness threshold. For k-colorings there is a slight discrepancy: uniqueness holds when $k \geq \Delta + 1$, while semi-translation invariant uniqueness holds for $k \geq \Delta$.

1.2.3.1 Gibbs measures from fixpoints of the tree recursions

Here, we will explicitly show that the tree recursions (7) and (8) correspond to translation and semi-translation invariant invariant Gibbs measures on $\mathbb{T}_\Delta$. Note that this argument is well-known; we just give the construction for completeness.

Given a fixpoint of the depth-one tree recursions (7), a Gibbs measure on $\mathbb{T}_\Delta$ for the spin system specified by $\mathbf{B}$, can be defined by a *broadcasting process* (see [55]); the spin of the root is chosen according to a probability distribution (specified by the fixpoint) and is then propagated along the edges of the tree, where each edge acts as a q-ary channel.

Specifically, for $\mathbf{r} = (R_1, \dots, R_q)$, define the probability vector $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_q)$ given by

$$\alpha_i = \frac{R_i \sum_{j=1}^q B_{ij} R_j}{\sum_{i,j} B_{ij} R_i R_j},$$

and consider the $q \times q$ matrix $\mathbf{M} = (M_{ij})_{i,j \in [q]}$ where $M_{ij} = B_{ij} R_j / \sum_{j=1}^q B_{ij} R_j$. It is immediate to verify that $\mathbf{M}$ is stochastic, so we can view $\mathbf{M}$ as a transition kernel. Moreover, by construction $\boldsymbol{\alpha}\mathbf{M} = \boldsymbol{\alpha}$, so that $\boldsymbol{\alpha}$ is the stationary distribution of the kernel $\mathbf{M}$. With these observations, it can easily be shown that the following broadcasting process generates a translation invariant Gibbs measure ν on $\mathbb{T}_\Delta$: first choose the spin at the root ρ according to the probability vector $\boldsymbol{\alpha}$. The spin of the root is then propagated along the edges of the tree; if u is the parent of v in the tree, then

$$\nu(\sigma(v) = j \,|\, \sigma(u) = i) = M_{ij}.$$

It is straightforward to check that each vertex has marginal spin distribution $\boldsymbol{\alpha}$.

For a fixpoint of the depth-two tree recursions (8), a slight modification of the above argument is possible. Namely, for $\mathbf{r} = (R_1, \dots, R_q), \mathbf{c} = (C_1, \dots, C_q)$, define the probability vectors $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_q), \boldsymbol{\beta} = (\beta_1, \dots, \beta_q)$ by

$$\alpha_i = \frac{R_i \sum_{j=1}^q B_{ij} C_j}{\sum_{i,j} B_{ij} R_i C_j}, \quad \beta_j = \frac{C_j \sum_{i=1}^q B_{ij} R_i}{\sum_{i,j} B_{ij} R_i C_j}. \tag{9}$$

Consider also the following $q \times q$ matrices: $\mathbf{M}^+ = (M^+_{ij})_{i,j \in [q]}$, whose (i,j)-entry is given by $B_{ij} C_j / \sum_{j=1}^q B_{ij} C_j$, and $\mathbf{M}^- = (M^-_{ij})_{i,j \in [q]}$, whose (i,j)-entry equals $B_{ij} R_j / \sum_{j=1}^q B_{ij} R_j$. As before, it is immediate to verify that $\mathbf{M}^+, \mathbf{M}^-$ are stochastic and by construction $\boldsymbol{\alpha}\mathbf{M}^+ = \boldsymbol{\beta}$ and $\boldsymbol{\beta}\mathbf{M}^- = \boldsymbol{\alpha}$. There are two possible ways now to define the broadcadsting process, which are symmetric up to interchanging $\boldsymbol{\alpha}, \boldsymbol{\beta}$ and $\mathbf{M}^+, \mathbf{M}^-$. We describe explicitly one of the two: first choose the spin at the root ρ according to the probability vector $\boldsymbol{\alpha}$. The spin of the root is then propagated along the edges of the tree; if u is the parent of v in the tree and v has odd distance from the root of the tree, then

$$\nu(\sigma(v) = j \,|\, \sigma(u) = i) = M^+_{ij}.$$

Otherwise,

$$\nu(\sigma(v) = j \,|\, \sigma(u) = i) = M^-_{ij}.$$

As in the translation-invariant case, it is straightforward to check that each vertex at odd level has marginal distribution $\boldsymbol{\beta}$, while each vertex at even level has marginal distribution $\boldsymbol{\alpha}$.

1.3 Phase transitions on random regular graphs

The behavior of spin systems on sparse random graphs has been the subject of intensive study in statistical physics, probability and probabilistic combinatorics, see for example [18]. We will use such graphs as gadgets for our inapproximability results. In this section, we make the crucial connection between the Gibbs distribution of such graphs and Gibbs measures, laying down the foundation for the most technical part of our work.

There is a well-known connection between sparse random graphs and random trees (see for example [17]). This connection is particularly simple for random Δ-regular graphs. Namely, a random Δ-regular graph *converges locally* to the infinite Δ-regular tree. We will not give the general definition of local convergence for graphs (see [6, 2]) but rather what it means in our setting. Let v be a uniformly random vertex of a random Δ-regular graph G with n vertices and let $t \geq 1$ be an arbitrary integer. The t-step neighborhood of v is the set of vertices in G whose graph distance from v is at most t. The local convergence of a random Δ-regular graph to $\mathbb{T}_\Delta$ can be stated as follows: for all n sufficiently large, with probability $1 - o(1)$ over the choice of v and the graph G, the t-step neighborhood of v is isomorphic to the t-step neighborhood of the root of $\mathbb{T}_\Delta$. In other words, for all but a vanishing fraction of vertices of a random Δ-regular graph G, the local neighborhood of a vertex looks like a complete tree. An immediate consequence of the local convergence to the tree is that the number of short cycles (i.e., constant length, not depending on n) in a random Δ-regular graph is $o(n)$.

Given this property, it is reasonable to expect that the analysis of the Gibbs distribution on random Δ-regular graphs would be connected with the infinite-volume Gibbs distribution on $\mathbb{T}_\Delta$. However, whether this connection actually holds turns out to *depend on the spin system*. Let us first explain in more detail how one can obtain an alignment between the two settings and then what can go wrong.

We will restrict our attention to two classes of graph distributions, random Δ-regular graphs and random bipartite Δ-regular graphs, which we will use for our inapproximability results. The approach we will ultimately describe for analyzing the Gibbs distribution on random graphs will be much more general for the class of random bipartite Δ-regular graphs, so we will focus on this distribution first and discuss extensions to the random Δ-regular graph case afterwards.

For a q-spin system with interaction matrix $\mathbf{B}$, our goal is to understand the Gibbs distribution on a random Δ-regular bipartite graph $G = (V, E)$ (with bipartition $V = V_1 \cup V_2$) by looking at the distribution of spin values in V_1 and V_2. Let $n = |V_1| = |V_2|$. For a configuration $\sigma : V \to [q]$, we denote the set of vertices assigned spin i by $\sigma^{-1}(i)$. For q-dimensional probability vectors $\boldsymbol{\alpha}, \boldsymbol{\beta}$, let

$$\Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}} = \left\{\sigma : V \to \{1, \dots, q\} \,\middle|\, |\sigma^{-1}(i) \cap V_1| = \alpha_i n,\, |\sigma^{-1}(i) \cap V_2| = \beta_i n \text{ for } i = 1, \dots, q\right\},$$

that is, configurations in $\Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ assign $\alpha_i n$ and $\beta_i n$ vertices in V_1 and V_2 the spin value i, respectively. For example, in the case of the hard-core model, the set $\Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ would correspond to the independent sets of G which have αn vertices in V_1 and βn vertices in V_2, upon identifying $\boldsymbol{\alpha}, \boldsymbol{\beta}$ with the scalars α, β since for 2-spin systems we only need two variables to capture the spin frequencies. Our interest will be in the total weight $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ of configurations in $\Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}$, namely

$$Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}} = \sum_{\sigma \in \Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}} w_G(\sigma).$$

What would be the pairs $\boldsymbol{\alpha}, \boldsymbol{\beta}$ with the largest $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$? Let us give a high level motivation to this question. We expect that for sufficiently large n, the $\boldsymbol{\alpha}, \boldsymbol{\beta}$ with the largest $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ dominate the partition function of G in the following sense: the partition function of G is just a sum of a polynomial (in n) number of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ which are typically exponential in n. Thus, the only pairs $\boldsymbol{\alpha}, \boldsymbol{\beta}$ with non-negligible contribution to the partition function should maximize $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ and should correspond in some sense to equilibrium states in the Gibbs distribution of the random graph.

To go one step further, let us look again at the example of the hard-core model and use the view with the local tree-like structure of the graph to speculate how the equilibrium might look. As previously, we consider independent sets with a prescribed number of $\alpha n, \beta n$ vertices in V_1, V_2. In a random such configuration, a vertex $v \in V_1$ would typically see a complete tree whose vertices at odd levels are occupied with probability β and at even levels are occupied with probability α. Thus, one would expect to see in the limit $n \to \infty$ a semi-translation invariant Gibbs measure of the infinite Δ-regular tree. The reader may have observed that this view is on a shaky ground; the phrase "in a random such configuration" implicitly assumes that a random configuration in the hard-core model distribution is decorrelated from the random graph. Despite this deficit of the argument, we now have an overview of what we would like to formalize.

A great portion of this thesis is targeted to establish the picture above for *general spin systems on random Δ-regular bipartite graphs.* To do this, and following previous approaches for the hard-core model [56], we will look at the moments of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ with respect to the distribution of the random Δ-regular bipartite graph, from hereon denoted by $\mathcal{G}$. For the next two subsections, we will focus on analyzing the leading terms of the first and second moments of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$, namely

$$\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta}) = \Psi_1^{\mathbf{B}}(\boldsymbol{\alpha},\boldsymbol{\beta}) := \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big],$$
$$\Psi_2(\boldsymbol{\alpha},\boldsymbol{\beta}) = \Psi_2^{\mathbf{B}}(\boldsymbol{\alpha},\boldsymbol{\beta}) := \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\Big[\big(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big)^2\Big].$$

The function Ψ_1 will be used to make connections to the depth-two recursions (8) and hence semi-translation invariant Gibbs measures on $\mathbb{T}_\Delta$. The function Ψ_2 will be used to identify those $(\boldsymbol{\alpha},\boldsymbol{\beta})$ for which $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ is asymptotically independent from the choice of the random graph.

1.3.1 First moment

To start, denote the leading term of the first moment of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ as:

$$\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta}) = \Psi_1^{\mathbf{B}}(\boldsymbol{\alpha},\boldsymbol{\beta}) := \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big],$$

We will be interested in the global maximizers of Ψ_1, which (at least in expectation) capture the configurations which have the largest contribution in the partition function of a random Δ-regular bipartite graph as $n \to \infty$. This motivates the following definition.

Definition 6. *For a q-spin system with interaction matrix* $\mathbf{B}$*, a* dominant phase *is a pair* $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ *of q-dimensional probability vectors which maximize* $\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$.

Now, we have a concrete function whose maxima hopefully capture the equilibrium states of a random bipartite Δ-regular graph G, we will address this point later. Prior to that, we first look at the critical points of Ψ_1 to get a better understanding of the dominant phases. The following definition will help to make the picture more clear.

Definition 7. *For a q-spin system with interaction matrix* $\mathbf{B}$*, a* phase *is a pair* $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ *of q-dimensional probability vectors which is a critical point of* $\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$.

A conceptual observation made in [56] in the context of the hard-core model was that the phases of a random Δ-regular bipartite graph correspond to the semi-translation invariant Gibbs measures on $\mathbb{T}_\Delta$. This extends rather easily to the general q-spin model as well, see Section 3.2.2 for a derivation in our setting. In particular, we restate for the reader's convenience the depth-two tree recursions on $\mathbb{T}_\Delta$.

$$\widehat{R}_i \propto \Big(\sum_{j=1}^{q} B_{ij} C_j\Big)^{\Delta-1} \quad \text{and} \quad \widehat{C}_j \propto \Big(\sum_{i=1}^{q} B_{ij} R_i\Big)^{\Delta-1}. \tag{8}$$

The phase $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ corresponding to a fixpoint $(R_1, \ldots, R_q, C_1, \ldots, C_q)$ of the tree recursions is given by (9), which we restate here for concreteness:

$$\alpha_i = \frac{R_i \sum_{j=1}^{q} B_{ij} C_j}{\sum_{i,j} B_{ij} R_i C_j}, \quad \beta_j = \frac{C_j \sum_{i=1}^{q} B_{ij} R_i}{\sum_{i,j} B_{ij} R_i C_j}. \tag{9}$$

Recall here that the vectors $\boldsymbol{\alpha}, \boldsymbol{\beta}$ are the marginal spin distributions of even and odd vertices in the semi-translation invariant Gibbs measure on $\mathbb{T}_\Delta$ corresponding to the fixpoint $(R_1, \ldots, R_q, C_1, \ldots, C_q)$ of (8), cf. Section 1.2.3.1. This establishes a firm connection between critical points of Ψ_1 and Gibbs measures on $\mathbb{T}_\Delta$.

However, even this level of understanding is not sufficient for our purposes. The reason being, that the critical points of a function f can be local minima, local maxima or saddle points of the function. We need hence a more refined picture to restrict our attention to the phases that can be dominant phases. This investigation will also allow us to connect the semi-translation invariant uniqueness threshold with the Jacobian stability of fixpoints.

To obtain a handle on the local maxima of a function, we use the Hessian criterion: a critical point $\mathbf{x} \in \mathbb{R}^t$ of a function $f : \mathbb{R}^t \to \mathbb{R}$ is a local maximum of f if the Hessian matrix of f at $\mathbf{x}$ is negative definite. When the Hessian criterion is true, we call the critical point a *Hessian local maximum.* In Section 3.4, we show the following connection between the stability of the tree recursions and Hessian local maxima of Ψ_1.

Theorem 2. *Jacobian attractive fixpoints of the depth-two tree recursions* (8) *correspond to Hessian local maxima of* Ψ_1*. Moreover, if* $\mathbf{B}$ *is ergodic (irreducible and aperiodic), Hessian local maxima of* Ψ_1 *correspond to Jacobian attractive fixpoints of the depth-two tree recursions.*

In particular, we have the following straightforward corollary.

Corollary 3. *For every ergodic* $\mathbf{B}$ *(irreducible and aperiodic), a phase of a random* Δ*-regular bipartite graph is a local maximum of* Ψ_1 *iff the corresponding fixpoint of the depth-two tree recursions is Jacobian attractive.*

Before concluding this section, let us comment briefly on the restriction in the second part of Theorem 2, which requires $\mathbf{B}$ to be irreducible and aperiodic. Irreducibility is a natural condition on the spin system; if not, the spin system may be decomposed into sub-spin systems whose interaction matrices are irreducible. For connected graphs, the partition function of the initial spin system is just the sum of the partition functions for each sub-spin system. Thus, the irreducibility condition merely excludes degenerate cases. The aperiodicity condition is also kind of a degenerate one, in the following sense. The interaction matrix $\mathbf{B}$ of a periodic spin system must have period two (since $\mathbf{B}$ is symmetric) and hence such a spin system is only interesting on bipartite graphs (otherwise the partition function is zero). Thus, the condition is rather an artifact of studying the spin system on a bipartite graph (for random regular graphs for example the condition may be immediately dropped). By restricting to appropriate subspaces, one can still show an analog of Theorem 2. We do not pursue this path here since for spin systems studied in statistical physics and computer science, it is usually the case that the system remains interesting on non-bipartite graphs, cf. the spin systems in Section 1.1.

1.3.2 Second moment

Having established a connection between dominant phases and Gibbs measures on $\mathbb{T}_\Delta$, we next come to perhaps the most intriguing aspect of establishing that the dominant phases are indeed "equilibria" of a random Δ-regular graph. We first explain why the first moment argument above is not sufficient. Using Markov's inequality, one can easily show that only dominant phases can have significant contribution to the partition function of G. This information alone however is insufficient, since we would like a lower bound for the contribution of dominant phases to the partition function to argue that the dominant phases correspond to natural equilibrium states of the random graph (with high probability over the choice of the graph). Ideally, the lower bound would be in terms of the expectation, and we are naturally led to consider the second moment of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$.

The classical second-moment method establishes the concentration of a random variable X around its expectation $\mathbf{E}[X]$, provided that the ratio $\mathbf{E}[X^2]/(\mathbf{E}[X])^2$ is close to 1. While this direct approach will not work for us, there is a substitute of the method, known as the small subgraph conditioning method, which works extremely well for probability spaces corresponding to random Δ-regular graphs when the ratio $\mathbf{E}[X^2]/(\mathbf{E}[X])^2$ is a constant (strictly greater than 1). The method still needs to argue however about the asymptotic order of the second moment, in our case $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]$.

As a necessary step, we need to consider first the leading term of the second moment, which we recall that it is given by

$$\Psi_2(\boldsymbol{\alpha},\boldsymbol{\beta}) = \Psi_2^{\mathbf{B}}(\boldsymbol{\alpha},\boldsymbol{\beta}) := \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\Big[\big(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big)^2\Big].$$

The following theorem relates the second moment to the first moment, for any spin system on random bipartite Δ-regular graphs. The proof is given in Section 3.3.

Theorem 4. *For any spin system, for all $\Delta \geq 3$,*

$$\max_{\boldsymbol{\alpha},\boldsymbol{\beta}} \Psi_2(\boldsymbol{\alpha},\boldsymbol{\beta}) = 2 \max_{\boldsymbol{\alpha},\boldsymbol{\beta}} \Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta}).$$

In particular, for every dominant phase $(\boldsymbol{\alpha},\boldsymbol{\beta})$, it holds that $\Psi_2(\boldsymbol{\alpha},\boldsymbol{\beta}) = 2\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta})$.

Theorem 4 suggests that for a dominant phase $(\boldsymbol{\alpha},\boldsymbol{\beta})$, $\mathbf{E}_{\mathcal{G}}\big[\big(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big)^2\big]/\big(\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]\big)^2$ should be asymptotically equal to a constant and hints that $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ should be close to its expectation. Though certain difficulties do emerge (the implicit constant turns out to be a constant greater than 1, see Section 1.3.3 for a more thorough discussion), it all comes down to proving Theorem 4.

Second moment arguments analogous to the one given in Theorem 4 tend to be difficult. To illustrate the difficulty, let us give the development of the proof for Theorem 4 in the case of the hard-core model. This was first done in [56], where they showed the equality $\Psi_2(\boldsymbol{\alpha},\boldsymbol{\beta}) = 2\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta})$ for the dominant phase at $\lambda = \lambda_c(\mathbb{T}_\Delta)$ for all $\Delta \geq 3$. By a convexity/continuity argument, they were able to conclude the existence of $\varepsilon(\Delta) > 0$ such that the same equality holds for $\lambda_c(\mathbb{T}_\Delta) < \lambda < \lambda_c(\mathbb{T}_\Delta) + \varepsilon(\Delta)$. Their argument was used in the seminal inapproximability results of [63], which due to the same technical reason were only obtained under the same condition and the interesting subcase $\lambda = 1$, $\Delta = 6$. The analysis of [56] was extended in [26] for certain regimes of λ, Δ in a way that was sufficient to conclude the inapproximability results for $\Delta = 3$ and $\Delta \geq 6$ for all $\lambda > \lambda_c(\mathbb{T}_\Delta)$. After that, in [27] the remaining cases $\Delta = 4, 5$ were settled by a somewhat easier analysis which applied to a subclass of general antiferromagnetic 2-spin systems. Simultaneously, [64] managed to circumvent the second-moment analysis for the hard-core model (and general 2-spin systems) by a different approach. Needless to say that the reason why the second-moment approach seemed to work for random bipartite Δ-regular graphs, even in the 2-spin case, remained elusive. To add more to this mystery, it was well-known that the second-moment approach on $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ fails for *random regular graphs* after a certain threshold (see, for example, [20] for more details).

The analysis for the proof of Theorem 4, which appeared in [28], is surprisingly simple given the difficulty in the previous approaches even in the rather restricted setting of 2-spin systems. There are two key ideas: the first one is the reformulation of the function $\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta})$ with an auxiliary function $\Phi(\mathbf{r},\mathbf{c})$. The reformulation is such that the maximum of Ψ_1 over $\boldsymbol{\alpha},\boldsymbol{\beta}$ equals the maximum Φ over $\mathbf{r},\mathbf{c}$. At the same time, the maximum of Φ can easily be expressed as a matrix norm of the interaction matrix $\mathbf{B}$. The second key idea is more of an observation: the second moment can be interpreted as a *paired-spin* model, i.e., a q^2-spin system with interaction matrix $\mathbf{B} \otimes \mathbf{B}$, where $\otimes$ denotes the Kronecker product. Then the maximum of Ψ_2 is captured by the norm of the matrix $\mathbf{B} \otimes \mathbf{B}$. It then comes down to connect the two norms, at which point we can utilize a result of Bennett [7], stating that the relevant norms are multiplicative over tensor product. For concreteness, we next give in more detail the connection between Ψ_1 and Φ.

First, we need to define the function Φ. We will need to set up some notation. Recall that the p-norm of a vector $\mathbf{x} \in \mathbb{R}^n$ is given by

$$\|\mathbf{x}\|_p := \left(\sum_{i=1}^{n} |x_i|^p \right)^{1/p}.$$

We will also use the subordinate matrix norm (also known as the induced matrix norm) which will be denoted as $\|\cdot\|_{p\to q'}$ (this is different than the usual notation $\|\cdot\|_{p\to q}$ used in the literature, since we have reserved q for the number of spins). For a matrix $\mathbf{A}$, the subordinate matrix norm $\|\mathbf{A}\|_{p\to q'}$ is defined as:

$$\|\mathbf{A}\|_{p\to q'} = \max_{\|\mathbf{x}\|_p=1} \|\mathbf{A}\,\mathbf{x}\|_{q'}.$$

Note that if $\mathbf{A}$ has non-negative entries then one can restrict the maximization to $\mathbf{x}$ with non-negative entries. A well-known example of an induced norm is the spectral norm $\|\cdot\|_{2\to 2}$.

Let $p = \Delta/(\Delta - 1)$. For non-negative q-dimensional vectors $\mathbf{r} := (R_1, \ldots, R_q), \mathbf{c} := (C_1, \ldots, C_q)$, define $\Phi(\mathbf{r}, \mathbf{c})$ by:

$$\exp\big(\Phi(\mathbf{r}, \mathbf{c})/\Delta\big) = \frac{\mathbf{r}^{\intercal}\mathbf{B}\mathbf{c}}{\|\mathbf{r}\|_p \|\mathbf{c}\|_p}.$$

The maximum of Φ can be compactly expressed in terms of matrix norms. as follows:

$$\max_{\mathbf{r},\mathbf{c}} \exp\big(\Phi(\mathbf{r}, \mathbf{c})/\Delta\big) = \max_{\mathbf{c}} \max_{\mathbf{r}} \frac{\mathbf{r}^{\intercal}\mathbf{B}\mathbf{c}}{\|\mathbf{r}\|_p \|\mathbf{c}\|_p} = \max_{\mathbf{c}} \frac{\|\mathbf{B}\mathbf{c}\|_\Delta}{\|\mathbf{c}\|_p} = \|\mathbf{B}\|_{p\to\Delta}, \tag{10}$$

where the second equality follows from matrix norm duality (see for example [33]).

The following theorem connects tree recursions, the function Φ and the function Ψ_1. The proof is given in Section 3.2.

Theorem 5. *There is a one-to-one correspondence between the fixpoints of the tree recursions and the critical points of* Φ *(both considered for non-negative* $\mathbf{r} = (R_1, \ldots, R_q)$, $\mathbf{c} = (C_1, \ldots, C_q)$ *in the projective space, that is, up to scaling by constants as in Remark 3).*

The following transformation $(\mathbf{r}, \mathbf{c}) \mapsto (\boldsymbol{\alpha}, \boldsymbol{\beta})$ *given by:*

$$\alpha_i = \frac{R_i^{\Delta/(\Delta-1)}}{\sum_i R_i^{\Delta/(\Delta-1)}} \quad \text{and} \quad \beta_j = \frac{C_j^{\Delta/(\Delta-1)}}{\sum_j C_j^{\Delta/(\Delta-1)}} \tag{11}$$

yields a one-to-one correspondence between the critical points of Φ *and the critical points of* Ψ_1 *(in the region defined by* $\alpha_i \geq 0, \beta_j \geq 0$ *and* $\sum_i \alpha_i = 1, \sum_j \beta_j = 1$*).*

Moreover, for the corresponding critical points $(\mathbf{r}, \mathbf{c})$ *and* $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ *one has*

$$\Phi(\mathbf{r}, \mathbf{c}) = \Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta}). \tag{12}$$

Finally, for ergodic $\mathbf{B}$ *(irreducible and ergodic), the local maxima of* Φ *and* Ψ_1 *happen at the critical points (that is, there are no local maxima on the boundary).*

The first item states that the critical points of Φ are given by the tree recursions; this follows from a straightforward differentiation. Equation (11) gives the correspondence between fixpoints of the tree recursions and the critical points of Ψ_1; note that this equation is seemingly different than (9), though the two expressions can easily be seen to be equivalent using that $(\mathbf{r}, \mathbf{c})$ are fixpoints of (8). Equation (12) is the most important piece for the second-moment analysis; it asserts that the values of the functions Ψ_1 and Φ at the corresponding critical points are equal. Since the final item guarantees that the maxima of Ψ_1, Φ happen at their critical points, we obtain that they have the same maximum value yielding the matrix norm formulation of Ψ_1 we wanted. We remark that the last point can also be obtained even when the maximum happens at the boundary; the restriction of ergodicity imposed on $\mathbf{B}$ is to make the connection more transparent.

1.3.3 Gibbs distribution of a random graph and short cycles

Theorem 4 establishes that for a dominant phase $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ the exponential order of the second moment of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ is twice the exponential order of the first moment of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$. The question we would like to address now is whether this can be used to establish concentration for $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$.

In the ideal scenario, we would be able to apply Chebyshev's inequality to get immediately a sharp concentration result. For this method to be succesful, it must be the case that the ratio $\mathbf{E}_{\mathcal{G}}\big[\big(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big)^2\big]/\big(\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]\big)^2$ converges asymptotically to 1. An intensive calculation of the asymptotic ratio, yields however

$$\lim_{n\to\infty}\frac{\mathbf{E}_{\mathcal{G}}\big[\big(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big)^2\big]}{\big(\mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big]\big)^2}=C>1, \tag{13}$$

for a constant C depending on the interaction matrix $\mathbf{B}$ and Δ, see Lemma 43 for an explicit form. Thus, the direct second-moment method only yields the existence of a graph G where $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ is comparable to its expected value. This is for several reasons prohibitive: first, it does not allow to get a statement which holds with high probability over the choice of the random graph; second, if there exist multiple dominant phases, i.e., the q-spin system is in the semi-translation invariant non-uniqueness regime of $\mathbb{T}_\Delta$, we are not guaranteed that the dominant phases coexist in at least one graph, which is a crucial component in utilizing the uniqueness phase transition. Fortunately, there exists a well-established method to deal with this apparent (13) failure of the second-moment method, known as the small subgraph conditioning method.

The small subgraph conditioning method was introduced in [60] to prove that a random Δ-regular contains asymptotically almost surely (a.a.s.) a Hamilton cycle. Roughly speaking, the method provides a way to get a.a.s results when the second-moment method fails, in the particular case (though common in the random regular graph setting) where the ratio of the second moment of a variable to the first moment squared converges to a constant strictly greater than 1. The method was first used to analyze spin models on random regular graphs in [56] and was subsequently used in [63, 27]. The principle behind the method is that the variance of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ comes from the presence of short cycles, which as we discussed in Section 1.3 are few, but still appear with non-zero probability (asymptotically with n, they follow the Poisson distribution).

In our setting, this comes about almost naturally. We have already seen that $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]$ is determined by the Gibbs measure on the infinite Δ-regular tree. On the other hand, we do expect a deviation from the expectation since a graph $G\sim\mathcal{G}$ does have $o(n)$ vertices which are contained in constant sized cycles. Thus, it is reasonable to expect that $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ fluctuates from its expectation. It is equally reasonable to expect the fluctuations to depend on the presence of small cycles which occur with small but non-zero probability. The surprising aspect of (14), a consequence of applying the small subgraph conditioning method, is that there is an explicit handle on these fluctuations.

The following lemma is a special case of the slightly stronger Lemma 36 which will help us to make the above explicit. The proof is given in Section 4.2.1.

Lemma 6. *Suppose that $(\boldsymbol{\alpha},\boldsymbol{\beta})$ is a dominant phase on random bipartite Δ-regular graph, which corresponds to a Jacobian attractive fixpoint of the tree recursions. Let $G \sim \mathcal{G}$ and denote by X_{in}, $i = 1, 2, \dots,$ the number of cycles of length $2i$ in G. There exist random variables W_{mn}, a deterministic function of $X_{1n}, X_{2n}, \dots, X_{mn}$, such that for every $\varepsilon > 0$*

$$\lim_{m\to\infty} \limsup_{n\to\infty} \Pr_{\mathcal{G}}\Big(\Big[\big|\frac{Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}}{\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]} - W_{mn}\big| > \varepsilon\Big]\Big) = 0, \tag{14}$$

There also exists a positive constant $c > 0$ such that $W_{mn} > c$ uniformly in m, n.

The rather unintuitive limit in (14) should be understood as follows: if we let m, n sufficiently large, the random variable $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}/\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]$ is well-approximated by the random variables W_{mn}, with large probability over the choice of the random graph. Equation (14) (or rather the slightly stronger version in Lemma 36) will be crucial for proving the properties of the gadget in our inapproximability results since it will allow us to tie $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ to the underlying graph. Thus, even when there exist multiple dominant phases, we can hope to compare their contribution to the partition function by understanding the corresponding random variables W_{mn}.

1.3.4 Extensions to random regular graphs

In this section, we discuss extensions of our techniques to random Δ-regular graphs. First, a word of caution; we cannot possibly hope to get as general results as in previous sections. For example, in the k-colorings model on random Δ-regular graphs the second-moment approach cannot possibly work when the number of colors is less than $\frac{\Delta}{2\log\Delta}(1+o(1))$, since this is with high probability the chromatic number of the random graph. However we do identify two cases where the approach works well. The most important one is in the context of ferromagnetic spin systems. Let us first the notion of ferromagnetism we use.

Definition 8. *A model is called* **ferromagnetic** *if* $\mathbf{B}$ *is positive definite. Equivalently we have that all of its eigenvalues are positive and also that*

$$\mathbf{B} = \hat{\mathbf{B}}^{\intercal}\hat{\mathbf{B}}, \tag{15}$$

for some $q \times q$ *matrix* $\hat{\mathbf{B}}$.

The most alluring aspect of this definition is that for a ferromagnetic model, neighboring vertices prefer to have the same spin, matching the intuition with the 2-spin setting. More generally, we have the following simple application of the Cauchy-Schwarz inequality. If $\mathbf{B}$ is ferromagnetic, then for probability vectors $\mathbf{z}_1, \mathbf{z}_2$, it holds that

$$(\mathbf{z}_1^{\intercal}\mathbf{B}\mathbf{z}_1)(\mathbf{z}_2^{\intercal}\mathbf{B}\mathbf{z}_2) \geq (\mathbf{z}_1^{\intercal}\mathbf{B}\mathbf{z}_2)^2,$$

In particular, if we plug in the above inequality the vectors with a single 1 in the positions i and j respectively, we obtain that any two spins i, j induce a ferromagnetic two-spin system. The definition is discussed more thoroughly in Section 8.1.2.

To state our results for ferromagnetic models, let us present briefly the slightly different setting. Let $G = (V, E)$ be a random Δ-regular graph with $n = |V|$. For a q-dimensional probability vector $\boldsymbol{\alpha}$, denote by $\Sigma^{\boldsymbol{\alpha}}$ the set of configurations σ which assign to $\alpha_i n$ vertices the spin i for each $i \in [q]$. Let

$$Z_G^{\boldsymbol{\alpha}} = \textstyle\sum_{\sigma \in \Sigma^{\boldsymbol{\alpha}}} w_G(\sigma).$$

We again study $Z_G^{\boldsymbol{\alpha}}$ by looking at the moments $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$ and $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha}})^2]$, where the expectation is over the distribution of the random Δ-regular bipartite graph G, from hereon denoted by $\mathcal{G}$ (no confusion should arise with the notation for bipartite random Δ-regular graphs).

Denote the leading term of the first and second moments as:

$$\Psi_1(\boldsymbol{\alpha}) = \Psi_1^{\mathbf{B}}(\boldsymbol{\alpha}) := \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha}}\big].$$
$$\Psi_2(\boldsymbol{\alpha}) = \Psi_2^{\mathbf{B}}(\boldsymbol{\alpha}) := \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\Big[\big(Z_G^{\boldsymbol{\alpha}}\big)^2\Big].$$

Dominant phases are vectors $\boldsymbol{\alpha}$ which are maximizers of $\Psi_1(\boldsymbol{\alpha})$. Analogously to Theorem 2, we prove the following in Section 8.3.

Theorem 7. *For a ferromagnetic model, Jacobian attractive fixpoints of the (depth-one) tree recursions are in one-to-one correspondence with the Hessian local maxima of* Ψ_1*.*

Also, we obtain the analogue of Theorem 4 for ferromagnetic models on random regular graphs. The reason that this is possible in the case of ferromagnetic Potts model is the Cholesky decomposition in Definition 8, which allows us to formulate the maximum of the first moment as a matrix norm (which in general is not possible).

Theorem 8. *For a ferromagnetic model,*

$$\max_{\boldsymbol{\alpha}} \Psi_2(\boldsymbol{\alpha}) = 2 \max_{\boldsymbol{\alpha}} \Psi_1(\boldsymbol{\alpha}).$$

Specifically, for dominant phases $\boldsymbol{\alpha}$*,* $\Psi_2(\boldsymbol{\alpha}) = 2\Psi_1(\boldsymbol{\alpha})$.

Theorem 8 is proved in Section 8.3.1. Combining Theorem 8 with the small subgraph conditioning method allows us to prove concentration for $Z_G^{\boldsymbol{\alpha}}$ (see Lemma 107). In particular, we verify the so-called *Bethe prediction* for general ferromagnetic models on random Δ-regular graphs, which is captured in our setting by equation (16) in the following theorem.

Theorem 9. *Let* $\mathbf{B}$ *specify a ferromagnetic model. Then, if there exists a Hessian dominant phase, it holds that*

$$\lim_{n\to\infty} \frac{1}{n}\mathbf{E}_{\mathcal{G}}[\log Z_G] = \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}[Z_G]. \tag{16}$$

Theorem 9 is proved in Section 8.6. Note that for a ferromagnetic model the interaction matrix $\mathbf{B}$ is positive definite and hence the entries on the diagonal are all positive. Thus Z_G is always positive for every graph G.

We note here that in contrast to $\lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}[Z_G]$, even proving the existence of $\lim_{n\to\infty} \frac{1}{n}\mathbf{E}_{\mathcal{G}}[\log Z_G]$ is far from trivial. There are however general techniques based on the interpolation method to establish the existence of the limit [4], which do not provide however its explicit value.

Theorem 9 can be extended to general models (not necessarily ferromagnetic) on random Δ-regular graphs under the stronger assumption that there is a unique semi-translation invariant Gibbs measure on $\mathbb{T}_\Delta$. In this setting, one also obtains the analogue of Theorem 8 and as a consequence concentration for $Z_G^{\boldsymbol{\alpha}}$ for the (unique) dominant phase $\boldsymbol{\alpha}$, which can be used to verify in complete analogy the Bethe prediction. Similar results are well-known in the literature (assuming uniqueness), the most general to the best of our knowledge are in [19].

Theorem 10. *Let* $\mathbf{B}$ *be the interaction matrix of a spin system such that* $Z_G > 0$ *for all* Δ-*regular graphs. Assume that there exists a unique semi-translation invariant Gibbs measure on* $\mathbb{T}_\Delta$ *and the corresponding fixpoint of the tree recurrences* (7) *is Jacobian attractive. Provided that the matrix* $\mathbf{B}$ *is regular (non-zero determinant),*

$$\lim_{n\to\infty} \frac{1}{n}\mathbf{E}_{\mathcal{G}}[\log Z_G] = \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}[Z_G]. \tag{17}$$

Theorem 10 is proved in Section 8.6. The assumption $Z_G > 0$ for all Δ-regular graphs is to avoid pathological cases where $\log Z_G \equiv -\infty$ in which case the quantities are not well-defined. It is satisfied by many classes of models, e.g., permissive models (cf. Section 1.1.1), such as the hard-core and antiferromagnetic Potts models, or even non-permissive, such as k-colorings when $k \geq \Delta + 1$. Finally, the restriction that $\mathbf{B}$ is regular is to ensure that we can apply the small subgraph conditioning method (see the relevant Lemma 31).

We conclude this section by noting that for natural models such as the hard-core or the k-colorings model, the second moment has been analyzed even in non-uniqueness regimes (see for example [1], [14]). Needless to say that Theorem 9 does not improve on these results; the scope of Theorem 9 is to provide a simple criterion (uniqueness of semi-translation invariant measures) which applies for general spin systems.

1.4 The complexity lens

1.4.1 Inapproximability for antiferromagnetic 2-spin systems

As we mentioned in Section 1.1.5, the complexity of approximately counting in antiferromagnetic 2-spin systems on graphs of maximum degree Δ is now well understood; recall that these are systems where neighboring vertices are favored to have different spins. Building on the works of Weitz [68] and Sly [63] for the hard-core model, the infinite Δ-regular tree has been identified to be the worst case graph for the persistence of long-range correlations in antiferromagnetic 2-spin systems. For 2-spin antiferromagnetic models, this establishes a beautiful picture connecting the computational complexity of approximating the partition function to statistical physics phase transitions in the infinite tree.

For the hard-core model on graphs of maximum degree Δ, there exists an FPTAS for the partition function for all $\lambda < \lambda_c(\mathbb{T}_\Delta)$, and for $\lambda > \lambda_c(\mathbb{T}_\Delta)$ the problem is intractable. Analogous results hold for general antiferromagnetic 2-spin systems; the uniqueness threshold on $\mathbb{T}_\Delta$ maps the boundary for efficient computation. For the algorithmic side, see [61, 44], and for the intractability side see [64]. As a straightforward application of our upcoming general theorem 14, we can also easily cover the computational hardness regime in the 2-spin setting.

Theorem 11. *Let $\Delta \geq 3$. For an antiferromagnetic 2-spin system in the non-uniqueness regime of $\mathbb{T}_\Delta$, unless $\mathsf{NP} = \mathsf{RP}$, there is no FPRAS for approximating the partition function for triangle-free Δ-regular graphs. Moreover, there exists $\varepsilon = \varepsilon(B_1, B_2, \Delta) > 0$ such that, unless $\mathsf{NP} = \mathsf{RP}$, one cannot approximate the partition function within a factor $2^{\varepsilon n}$ for triangle-free Δ-regular graphs (where n is the number of vertices).*

Theorem 11 is proved in Section 6.1.

Remark 4. *The hardness result remains true even for Δ-regular graphs with girth at least g, for any constant $g \geq 3$.*

Remark 5. *The condition $\mathsf{NP} \neq \mathsf{RP}$ is because we exclude the possibility of an FPRAS, as is typical in problems of the field. If instead we restricted our attention to an FPTAS, the condition may be replaced with $\mathsf{NP} \neq \mathsf{P}$.*

Remark 6. *The above two remarks apply to the upcoming Theorems 12, 13, 14, 15 as well.*

Before proceeding, we remark that for ferromagnetic 2-spin systems, a seminal algorithm by Jerrum and Sinclair [38] approximates the partition function for all graphs (without degree bound) in the ferromagnetic Ising model (allowing the presence of external field), and this extends to arbitrary ferromagnetic 2-spin systems on regular graphs. For ferromagnetic

2-spin systems on non-regular graphs, the picture is more complicated, see [31, 45] for more details.

1.4.2 NP-hardness for the colorings and antiferromagnetic Potts models

The picture for multi-spin systems (systems with $q > 2$ possible spins for vertices) is much less clear; the above approaches for 2-spin systems do not extend to multi-spin models in a straightforward manner. We aim to establish the analog of the above inapproximability results for the colorings problem, namely, NP-hardness in the tree non-uniqueness region. Our techniques and results generalize to a broad class of antiferromagnetic spin systems.

For the colorings problem, even understanding the uniqueness threshold is challenging. Jonasson [41] established uniqueness when $k \geq \Delta+1$, and it is easy to show non-uniqueness when $k \leq \Delta$ since a fixed coloring on the leaves can "freeze" the internal coloring. As we mentioned earlier, for k-colorings the uniqueness threshold and the semi-translation invariant uniqueness threshold no longer coincide. In particular, Brightwell and Winkler [11] established, for semi-translation invariant measures, uniqueness when $k \geq \Delta$ and non-uniqueness when $k < \Delta$.

We prove, for even k, that it is NP-hard to approximate the number of proper k-colorings (in other words, NP-hard to approximate the partition function) when there are multiple semi-translation invariant Gibbs measures on $\mathbb{T}_\Delta$ (which corresponds to $k < \Delta$). Moreover, our result proves hardness for the class of triangle-free Δ-regular graphs. Hence, our result is particularly interesting in the region $k = \Omega(\Delta/\log\Delta)$ since a seminal result of Johansson [40, 51] shows that all triangle-free graphs are colorable with $O(\Delta/\log\Delta)$ colors. His proof, which uses the nibble method and the Lovász Local Lemma, can be made algorithmic using the constructive proof of [54]. For small values of Δ, one can use that triangle-free graphs are colorable with $3\left\lceil\frac{\Delta+1}{4}\right\rceil$ colors (this bound is also algorithmic[3]), see [50, Chapter 12] or [37, Chapter 4] for references and a thorough account of bounds when stronger restirictions

[3]We give an explicit description of an algorithm based on Catlin's proof [13]. Let $r = \left\lceil\frac{\Delta+1}{4}\right\rceil$. Partition the vertex set of the graph into r sets $V_1, \ldots, V_r$ arbitrarily. Iteratively, if a vertex $v \in V_i$ has more than three neighbors in V_i move it to a set V_j where v has at most three neighbors; such a set exists by the choice of r and the bound Δ on the degree of v. This procedure strictly decreases the aggregate number of edges within the sets $V_1, \ldots, V_r$, so it must end in at most linear number of steps. In the final partition, each V_i induces a triangle-free graph of maximum degree 3, and thus can be 3-colored by Brooks' theorem.

on the girth of the graph are imposed. For general graphs with maximum degree Δ, the interesting region is $k = \Delta - O(\sqrt{\Delta})$, since Molloy and Reed [50] showed, for all fixed constants $\Delta \geq \Delta_0$, determining if a graph G of maximum degree Δ has a $(\Delta + 1 - l)$-coloring is in P for any l such that $l^2 + l \leq \Delta$ (the result is tight up to the value of Δ_0, see [25, 50] for more details). We note that most parts of the proof extend to the odd k case as well, modulo the missing piece in the phase diagram, see Section 1.4.5.1 for a precise description.

Here is the formal statement of our inapproximability result for colorings.

Theorem 12. *For all even $k \geq 3$, all $\Delta \geq 3$, for the k-colorings problem, when $k < \Delta$, unless $\mathsf{NP} = \mathsf{RP}$, there is no FPRAS that approximates the partition function for triangle-free Δ-regular graphs. Moreover, there exists $\varepsilon = \varepsilon(k, \Delta) > 0$ such that, unless $\mathsf{NP} = \mathsf{RP}$, one cannot approximate the partition function within a factor $2^{\varepsilon n}$ for triangle-free Δ-regular graphs (where n is the number of vertices).*

Theorem 12 is proved in Section 6.3. The result extends also to the antiferromagnetic Potts model. As we mentioned earlier, somewhat surprisingly, the uniqueness/non-uniqueness threshold for the infinite tree $\mathbb{T}_\Delta$ is not known for the antiferromagnetic Potts model. We prove that the uniqueness/non-uniqueness threshold for semi-translation invariant Gibbs measures on $\mathbb{T}_\Delta$ is given by $B_c(\Delta) = \frac{\Delta - q}{\Delta}$. We believe this threshold coincides with the uniqueness/non-uniqueness threshold, unlike in the case of colorings. We prove, for even q, that approximating the partition function is NP-hard in the non-uniqueness region for semi-translation invariant measures. The following theorem, which is proved in Section 6.3, is the formal statement of our inapproximability results for the antiferromagnetic Potts model.

Theorem 13. *For all even $q \geq 3$, all $\Delta \geq 3$, for the antiferromagnetic q-state Potts model, for all $B < \frac{\Delta - q}{\Delta}$, unless $\mathsf{NP} = \mathsf{RP}$, there is no FPRAS that approximates the partition function for triangle-free Δ-regular graphs. Moreover, there exists $\varepsilon = \varepsilon(q, B, \Delta) > 0$ such that, unless $\mathsf{NP} = \mathsf{RP}$, one cannot approximate the partition function within a factor $2^{\varepsilon n}$ for triangle-free Δ-regular graphs (where n is the number of vertices).*

In fact, we obtain inapproximability of the partition function for any antiferromagnetic model when there is non-uniqueness of semi-translation invariant measures on $\mathbb{T}_\Delta$ and mild additional conditions. Our results for general models are stated in Section 1.4.3.

1.4.3 Inapproximability for general antiferromagnetic models

The inapproximability results for colorings can be extended to general antiferromagnetic models on bounded degree graphs. The key concept in the general theorem is again the existence of long-range correlations, in the form of semi-translation non-uniqueness on the infinite regular tree.

We use the following definition of antiferromagnetic models, which is in terms of the signature of the interaction matrix $\mathbf{B}$, i.e., the signs of its eigenvalues. Recall that the interaction matrix $\mathbf{B}$ is symmetric and hence its eigenvalues are real. Moreover, it is simple to see that the matrices $\mathbf{B}$ which correspond to non-degenerate models should be irreducible. The Perron-Frobenius theorem then implies that one of the eigenvalues with the largest magnitude is positive. When all the other eigenvalues are negative, we show in Section 5.5.1 (see Corollary 61) that neighboring spins prefer to be different and hence we call such models antiferromagnetic. This notion of antiferromagnetism extends naturally the definition for 2-spin models (see [31, 44, 64]) and also captures antiferromagnetism in the Potts model. A more thorough discussion is given in Section 5.5.1.

Proceeding to the general hardness results, we have already displayed that the second moment argument for random bipartite graphs covers arbitrary models (general interaction matrix $\mathbf{B}$). There are two properties of the gadget G in Section 1.4.5.2 we need to ensure. First, the symmetry breaking between the two sides of the graph, i.e., a typical configuration should have different color frequencies on the two sides of the graph. Second, to be able to quantify the interaction between neighboring gadgets, we need to ensure that phases appear with roughly equal probability and that given the phase, the spins of the vertices with degree $\Delta - 1$ are approximately independent.

The first property can be guaranteed by the absence of translation invariant phases, i.e., maximizers of the function Ψ_1 of the form $(\mathbf{x}, \mathbf{x})$. Since in uniqueness regimes the only maximizer is translation invariant, this can be the case only in non-uniqueness regimes. The second property is subtler and relates to the concentration properties of the random variable $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$. The required concentration is sufficiently strong when the maximizers $\boldsymbol{\alpha},\boldsymbol{\beta}$ (viewed as unorderded pairs) of the function Ψ_1 are (i) Hessian maxima, i.e., the Hessian matrix of Ψ_1 is negative definite when evaluated at $(\boldsymbol{\alpha},\boldsymbol{\beta})$, and (ii) permutation-symmetric, i.e., obtainable from one another by a suitable permutation of the set of spins (equivalently, the coordinates of α_i's and β_j's). We clarify that the permutations must be automorphisms of the interaction matrix $\mathbf{B}$. For example, note that the maxima for the colorings model are permutation-symmetric.

Given these assumptions, the proof approach in Section 1.4.5.2 (with some extra work) can be adapted to give the following general inapproximability result.

Theorem 14. *Let $q \geq 2, \Delta \geq 3$. For an antiferromagnetic q-spin system with interaction matrix $\mathbf{B}$, if the dominant semi-translation invariant Gibbs measures on the tree $\mathbb{T}_\Delta$ are permutation-symmetric and all of them are Hessian dominant and not translation invariant then, unless* NP = RP*, there is no* FPRAS *for approximating the partition function for triangle-free Δ-regular graphs. Moreover, there exists $\varepsilon = \varepsilon(q, \mathbf{B}, \Delta) > 0$ such that, unless* NP = RP*, one cannot approximate the partition function within a factor $2^{\varepsilon n}$ for triangle-free Δ-regular graphs (where n is the number of vertices).*

Theorems 12 and 13 can be obtained as corollaries of Theorem 14 after using the detailed analysis of the dominant phases in the upcoming Theorem 16. We also do the analogous much easier task for antiferromagnetic 2-spin systems in Section 6.1 and thus obtain Theorem 11.

1.4.4 NP-hardness for general spin systems with unique dominant phase

For general models with interaction matrix $\mathbf{B}$, when there is no restriction on the eigenvalues of $\mathbf{B}$, we can prove hardness whenever there is a unique dominant phase which is not translation invariant. The formal statement is along the lines of Theorem 14 and is stated below.

Theorem 15. *Let $q \geq 2, \Delta \geq 3$. For a general q-spin system with interaction matrix $\mathbf{B}$, if there exists a unique dominant semi-translation invariant Gibbs measure on the tree $\mathbb{T}_\Delta$ which is Hessian dominant and not translation invariant then, unless* $\mathsf{NP} = \mathsf{RP}$*, there is no* FPRAS *for approximating the partition function for triangle-free Δ-regular graphs. Moreover, there exists $\varepsilon = \varepsilon(q, \mathbf{B}, \Delta) > 0$ such that, unless* $\mathsf{NP} = \mathsf{RP}$*, one cannot approximate the partition function within a factor $2^{\varepsilon n}$ for triangle-free Δ-regular graphs (where n is the number of vertices).*

The proof of Theorem 15 is given in Section 5.3. We remark that the reduction for Theorem 15 is fairly simple due to the restriction that the dominant phase is unique. In fact, it is a straightforward extension of the known reduction for antiferromagnetic 2-spin systems (note however that the multi-spin setting requires the new results of Section 1.3). We highlight it here because it hints that perhaps a more general inapproximability theorem could be possible which does not require the eigenvalue restriction on the interaction matrix $\mathbf{B}$ (used in Theorem 14).

1.4.5 Overview of the reduction for the k-colorings model

In this section, we overview the reduction for the k-colorings model. We start with analyzing the dominant phases of the model on random bipartite Δ-regular graphs. This reduces to computing the matrix norm $\|\mathbf{B}\|_{\frac{\Delta}{\Delta-1}\to\Delta}$, where $\mathbf{B}$ is the interaction matrix of the k-colorings model. This optimization problem turns out to be quite hard to solve in full generality, i.e., for all values of k, Δ. We solve this problem for all even k and all $\Delta \geq 3$ (see Theorem 16), and this is where the restriction on k comes from in Theorem 12. We should note that if this analysis was extended to k odd, it would immediately extend Theorem 12 to odd k as well.

We also remark that the analysis of $\|\mathbf{B}\|_{\frac{\Delta}{\Delta-1}\to\Delta}$ when $\mathbf{B}$ is the interaction matrix of the antiferromagnetic Potts model turns out to be analogous to the colorings model (modulo the extra parameter B which causes more technical difficulties), so we treat the two models simultaneously in Section 1.4.5.1. Analogous remarks as in the colorings model apply for the restiction on q to be even in Theorem 13.

In Section 1.4.5.2, we describe how to use this information to obtain a gadget for the reduction. This part is similar to [64] for 2-spin systems, modulo of course the technical bottleneck of the analysis of the second moment for random Δ-regular bipartite graphs (which we described thoroughly in previous sections). We then highlight the main points where previous approaches for antiferromagnetic 2-spin systems do not work (roughly, the existence of a large number of dominant phases) and the new elements required.

1.4.5.1 Dominant phases for the antiferromagnetic Potts and colorings models

To obtain Theorems 12 and 13, we need to figure out the dominant phases in a random bipartite regular graph for the antiferromagnetic Potts and colorings models. Recall, the interaction matrix $\mathbf{B}$ for the Potts model is completely determined by a parameter B. The antiferromagnetic regime corresponds to $0 < B < 1$. The coloring model is the zero temperature limit of the Potts model and corresponds to the particular case $B = 0$ in what follows. We should note that in statistical physics terms, the arguments of this section are closely related to the phase diagrams of the models.

Recall that the critical points of Ψ_1 are given by fixpoints of the tree recursions. For the Potts model, the tree recursions (7) can be written as:

$$R_i \propto \Big(BC_i + \textstyle\sum_{j\neq i} C_j\Big)^{\Delta-1}, \quad C_j \propto \Big(BR_j + \textstyle\sum_{i\neq j} R_i\Big)^{\Delta-1}, \tag{18}$$

Using Theorem 2 to connect Jacobian attractive fixpoints and dominant phases together with the function Φ of Theorem 5, we establish the following theorem in Section 6.3.

Theorem 16. *Let* $0 \leq B < 1$ *and* $\Delta \geq 3$.

1. *When* $q \geq \Delta + 1$ *or* $B \geq \frac{\Delta-q}{\Delta}$, *there is a unique fixpoint of* (18) *which is translation invariant. Thus the model is in the semi-translation invariant uniqueness regime of* $\mathbb{T}_\Delta$.

2. *When* $q < \Delta$ *and* $0 \leq B < \frac{\Delta-q}{\Delta}$, *there are multiple fixpoints of* (18). *Thus the model is in the semi-translation invariant non-uniqueness regime of* $\mathbb{T}_\Delta$.

3. *For all even* $q \geq 3$, *for all* $\Delta \geq 3$, *when* $q < \Delta$ *and* $0 \leq B < \frac{\Delta-q}{\Delta}$, *the dominant phases* $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ *are in one-to-one correspondence with subsets* $T \subseteq [q]$ *with* $|T| = q/2$.

Moreover, there exist a, b such that for $T \subseteq [q]$ with $|T| = q/2$, the dominant phase $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ corresponding to T satisfies

$$\alpha_i = a \text{ if } i \in T, \quad \alpha_i = b \text{ if } i \notin T,$$

$$\beta_i = b \text{ if } i \in T, \quad \beta_i = a \text{ if } i \notin T.$$

1.4.5.2 The gadget and its properties

In this section, we give the main elements of our inapproximability results. We start by reviewing the main components of the reduction for 2-spin systems (as carried out in [63, 64]) and in particular the hard-core model. This will allow us to isolate the parts of the argument which do not extend to the multi-spin case and motivate our reduction scheme. To simplify the presentation, we shall focus on the colorings model (k even), but the same ideas can be generalized to the Potts model and (with more technical effort) to arbitrary antiferromagnetic multi-spin models.

The basic gadget in the reduction is a bipartite random graph, which we denote by G. The sides of the bipartition have an equal number of vertices, and the sides are labelled with $+$ and $-$. Most vertices in G have degree Δ but there is also a small number of degree $\Delta-1$ vertices (to allow to make connections between gadgets without creating degree $\Delta+1$ vertices). For $s = \{+, -\}$, let the vertices in the s-side be $U^s \cup W^s$ where the vertices in $U = U^+ \cup U^-$ have degree Δ and the vertices in $W = W^+ \cup W^-$ have degree $\Delta - 1$. The phase of an independent set I is $+$ (resp. $-$) if I has more vertices in U^+ (resp. U^-). Note that the phase depends only on the spins of the "large" portion of the graph, i.e., the spins of vertices in U.

In non-uniqueness regimes, the gadget G has two important properties, both of which can be obtained by building on the second-moment argument we outlined earlier. First, the phase of a random independent set I is equal to $+$ or $-$ with probability roughly equal to $1/2$. Second, conditioned on the phase of a random independent set I, the spins of the vertices in W are approximately independent, i.e., the marginal distribution on W is close to a product distribution. In this product distribution if the phase is $+$ (resp. $-$), a vertex in W^+ is in I with probability p^+ (resp. p^-), while a vertex in W^- is in I with probability

p^- (resp. p^+). The values $p^\pm$ correspond to maxima of the function Ψ_1 and, crucially (as we shall demonstrate shortly), they satisfy $p^+ \neq p^-$.

Using the second-moment analysis and in particular Theorem 4, we can prove that an analogous phenomenon takes place for the k-colorings model in the semi-translation non-uniqueness regime. The main difference is that, instead of two phases, the number of phases is equal to the number of maximizers of the function Ψ_1, cf. Theorem 16. For k even, the phase of a coloring is determined by the dominant set of $k/2$ colors on U^+, i.e., the $k/2$ colors with largest frequencies among vertices of U^+. Each of the $\binom{k}{k/2}$ phases appears with roughly equal probability and given the phase, the marginal distribution on W is close to a product distribution, which we now describe. We can compute explicit values $a = a(k, \Delta), b = b(k, \Delta)$ such that for a phase $T \in \binom{[k]}{k/2}$ the probability mass function $\mathbf{x}$ of a vertex in W^+ has its i-th entry equal to a if $i \in T$ and equal to b if $i \notin T$. Similarly, the probability mass function $\mathbf{y}$ of a vertex in W^- has its i-th entry equal to b if $i \in T$ and equal to a if $i \notin T$. We note here that the values of a, b can be determined using the mapping between the dominant phases (described in Theorem 16) and their respective fixpoints of the tree recursions (see (11)); see also Section 6.3.

Let $\mathcal{Q}$ be the union of the pairs $(\mathbf{x}, \mathbf{y})$ over all phases. Hereafter, we will identify the phases with elements of $\mathcal{Q}$. Note that if $(\mathbf{x}, \mathbf{y}) \in \mathcal{Q}$, then $(\mathbf{y}, \mathbf{x}) \in \mathcal{Q}$ as well. We also denote by $\mathcal{Q}'$ the union of unordered elements of $\mathcal{Q}$. Elements of $\mathcal{Q}'$ are called unordered phases (we use $\mathbf{p}$ to denote unordered phases). Given a phase $\mathbf{p} = \{\mathbf{x}, \mathbf{y}\}$ an ordering of the pair will be called "assigning spin to the phase". The two ordered phases corresponding to the unordered phase $\mathbf{p}$ will be denoted by $\mathbf{p}^+$ and $\mathbf{p}^-$.

The conditional independence property is crucial. It allows to quantify the effect of using vertices of W as terminals to make connections between copies of the gadget G. For example, consider the following type of connection, which we refer to as parallel. Let $v^+ \in W^+$, $v^- \in W^-$ and consider two copies of the gadget G, say G_1, G_2. For $i = 1, 2$ denote by v_i^+, v_i^- the images of v^+, v^- in G_i. Now add the edges (v_1^+, v_2^+) and (v_1^-, v_2^-) and denote the final graph by G_{12}. Thus, a parallel connection corresponds to joining the $+, +$ and $-, -$ sides of two copies of the gadget.

Clearly, random colorings of G_{12} can be generated by first generating random colorings of G_1, G_2 and keeping the resulting coloring if $v_1^{\pm}, v_2^{\pm}$ have different colors. We thus have that the partition function of G_{12} is equal to $(Z_G)^2$ times the probability that $v_1^{\pm}, v_2^{\pm}$ have different colors in random colorings of G_1, G_2. The latter quantity can easily be computed if we condition on the phases $(\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)$ of the colorings in G_1, G_2, and this is equal to $(1 - \mathbf{x}_1^{\intercal}\mathbf{x}_2)(1 - \mathbf{y}_1^{\intercal}\mathbf{y}_2)$.

By taking logarithms, we can assume a parallel connection between gadgets with phases $(\mathbf{x}_1, \mathbf{y}_1)$ and $(\mathbf{x}_2, \mathbf{y}_2)$ incurs an (additive) weight

$$w_p((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) = \ln(1 - \mathbf{x}_1^{\intercal}\mathbf{x}_2) + \ln(1 - \mathbf{y}_1^{\intercal}\mathbf{y}_2).$$

In the hard-core model, parallel connections are sufficient to give hardness. In this case, we have that $\mathcal{Q}' = \{\mathbf{p}\}$ and $\mathcal{Q} = \{\mathbf{p}^+, \mathbf{p}^-\}$ and the respective function $w_p(\cdot, \cdot)$ satisfies

$$w_p(\mathbf{p}^+, \mathbf{p}^+) = w_p(\mathbf{p}^-, \mathbf{p}^-) < w_p(\mathbf{p}^+, \mathbf{p}^-). \tag{19}$$

Thus, in this case, $w_p(\cdot, \cdot)$ takes only two values and neighboring gadgets prefer to have different phases. Now assume that H is an instance of Max-Cut and replace each vertex in H by a copy of the gadget G, while for each edge of H, connect the respective gadgets in parallel. The partition function of the final graph is dominated from phase assignments which correspond to large cuts in H. This intuition is the basis of the reduction in [63, 64].

For the colorings model, reducing from Max-Cut poses an extra challenge. While for every unordered phase $\mathbf{p}$ equation (19) continues to hold, a short calculation shows that the optimal configuration for a triangle of gadgets connected in parallel is to give all three gadgets different phases. To bypass this entanglement, we need to introduce some sort of ferromagnetism in the reduction to enforce gadgets corresponding to vertices of H to use a single (unordered) phase. To achieve this, we use *symmetric* connections, which correspond to having not only $(+,+), (-,-)$ connections of the gadgets, but also $(+,-)$ and $(-,+)$. Thus, a symmetric connection whose endpoints have phases $(\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)$ incurs (additive) weight

$$w_s((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) = w_p((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2)) + w_p((\mathbf{x}_1, \mathbf{y}_1), (\mathbf{y}_2, \mathbf{x}_2)).$$

Symmetric connections will allow us to enforce a single unordered phase to all gadgets, while parallel connections will allow us to recover a maximum-cut partition. To have some modularity in our construction, rather than reducing from MAX-CUT directly, we use the following "phase labeling problem".

COLORINGS PHASE LABELING PROBLEM($\mathcal{Q}$):

INPUT: undirected edge-weighted multigraph $H = (V, E)$ and a partition of the edges $\{E_p, E_s\}$.

OUTPUT: $\textsc{MaxLwt}(H) := \max_{\mathcal{Y}} \textsc{Lwt}_H(\mathcal{Y})$, where the maximization is over all possible phase labelings $\mathcal{Y} : V \to \mathcal{Q}$ and

$$\textsc{Lwt}_H(\mathcal{Y}) := \sum_{\{u,v\}\in E_s} w_s(\mathcal{Y}(u), \mathcal{Y}(v)) + \sum_{\{u,v\}\in E_p} w_p(\mathcal{Y}(u), \mathcal{Y}(v)).$$

Edges in E_p (resp. E_s) correspond to parallel (resp. symmetric) connections and we shall refer to them as parallel (resp. symmetric) edges. The arguments in [64], which we sketched earlier, can easily be adapted to show that an algorithm for approximating the partition function to an arbitrarily small exponential factor yields a PRAS for the phase labeling problem, see the general Lemma 55 in Section 5.1. The harder part of our arguments is to show that a PRAS for the phase labeling problem yields a PRAS for MAX-CUT on 3-regular graphs, see Lemma 58 in Section 5.4.

We conclude by pointing to the general Lemma 56 and its proof in Section 5.5 for the proof of Theorem 14.

1.4.6 AP-reductions for graphs of bounded degree

The sharp transition in the complexity of the hard-core model on graphs with maximum degree Δ is one of the most striking computational dichotomies. It would be interesting thus to explore deeper the boundaries of this dichotomy.

A natural such question is whether the same computational transition remains true for bipartite graphs of maximum degree Δ. One side is trivial: in the regime $\lambda < \lambda_c(\mathbb{T}_\Delta)$, Weitz's algorithm applies invariably irrespectively of the underlying graph structure (provided of course that the maximum degree is Δ). On the hardness side, the things are no longer clear. Even for general bipartite graphs (with unbounded degree) a strong NP-hardness result is considered unlikely, though it is conjectured to not admit an FPRAS. The class of problems whose approximation complexity is equivalent to counting independent sets on bipartite graphs (which we will abbreviate as #BIS) includes many natural problems, see for example [23].

Most #BIS-hardness results have thus far focused on graphs without degree bounds. We make a step for bipartite graphs of bounded degree, first in the case of the hard-core model, proving #BIS-hardness for the hard-core model throughout the non-uniqueness regime of $\mathbb{T}_\Delta$. Thus, if the conjectured hardness of #BIS holds true, the following Theorem 17 is tight (see Section 7.3 for the proof).

Theorem 17. *For all $\Delta \geq 3$, for the hard-core model, for any $\lambda > \lambda_c(\mathbb{T}_\Delta)$, it is* #BIS-hard *to obtain an* FPRAS *that approximates the partition function for bipartite graphs of maximum degree Δ.*

We prove an analogous result for the ferromagnetic Potts model. We utilize an important result of Goldberg and Jerrum [30], who showed that approximating the partition function of the ferromagnetic Potts model on general graphs is BIS-hard. Using bipartite random Δ-regular graphs as gadgets, we show that the hardness is maintained if one restricts to bounded degree bipartite graphs, provided the parameter B is bigger than a natural threshold $\mathfrak{B}_o$ (a function of q, Δ), see Section 6.2 for more details.

Theorem 18. *For all $q \geq 3$, all $\Delta \geq 3$, for the ferromagnetic q-state Potts model, for any $B > \mathfrak{B}_o$, it is* #BIS-hard *to obtain an* FPRAS *that approximates the partition function for bipartite graphs of maximum degree Δ.*

Theorem 18 is proved in Section 7.2. The threshold $\mathfrak{B}_o$, rather than being the uniqueness threshold of the ferromagnetic Potts model on the infinite Δ-regular tree $\mathbb{T}_\Delta$, is the *phase coexistence point* on a random Δ-regular graph (whose existence we establish). When $B < \mathfrak{B}_o$, the uniform vector is the only dominant phase; when $B > \mathfrak{B}_o$, the uniform vector over the spins is not a dominant phase and q permutation-symmetric dominant phases exist; at $B = \mathfrak{B}_o$, the model has $q+1$ dominant phases, the uniform vector and the q permutation-symmetric dominant phases. For more details, see Theorem 69 in Section 6.2.

CHAPTER II

SPIN SYSTEMS ON RANDOM REGULAR BIPARTITE GRAPHS

Random regular bipartite graphs will serve as building blocks for the gadgets we use in our hardness reductions. Thus, a major component in this thesis is the analysis of spin systems on models of random regular graphs. Our primary focus will be on random bipartite Δ-regular graphs; this graph distribution is formally defined in Section 2.1.1. Extensions to random regular graphs are discussed in Chapter 8.

To obtain results for the Gibbs distribution which hold with probability $1-o(1)$ over the choice of the graph, we shall employ a second-moment approach based on the small subgraph conditioning method. We formulate the main technical aspects of the moment approach in Section 2.2.1 and review the small subgraph conditioning method in Section 2.3, all in the spin model setting. Later chapters will progressively establish and refine the various components needed in the analysis.

The reader who is acquainted with similar settings as ours will effectively recognize a good portion of the notions and methods, modulo perhaps the specific technical details. Namely, both the use of moments and the small subgraph conditioning method are well established techniques which have been applied to numerous problems in probabilistic combinatorics and statistical physics. Despite this, the technical details change substantially for each instantation of the method and the purpose of this chapter is to set up a concrete basis for our later investigations.

2.1 Preliminaries

2.1.1 The distribution on Δ-regular Bipartite Graphs

We will use the following simple variant of the standard pairing model (also known as the configuration model [5, 9]) to study random Δ-regular bipartite graphs.

Let V_1, V_2 be disjoint sets of vertices with $|V_1| = |V_2| = n$. For an integer $\Delta \geq 3$, we will denote by $\mathcal{G}_{n,\Delta}$ the graph distribution on the set of bipartite Δ-regular multigraphs generated by the following random process. First, sample uniformly at random Δ perfect matchings between V_1 and V_2. Then, construct a graph G with vertex set $V_1 \cup V_2$ whose edges are given by the union of the Δ perfect matchings. G is *simple* if it does not contain parallel edges.

For the sake of completeness, we next compare the distribution $\mathcal{G}_{n,\Delta}$ with the standard pairing model $\mathcal{P}_{n,\Delta}$ for random Δ-regular bipartite graphs. The conclusion we are going to make is that there is no loss in working with the distribution $\mathcal{G}_{n,\Delta}$. The reader may skip the remainder of the paragraph without impact on what follows. $\mathcal{P}_{n,\Delta}$ is generated by a uniformly random perfect matching between $[V_1] \times [\Delta]$ and $[V_2] \times [\Delta]$ and then naturally projecting the edges of the matching on $V_1 \times V_2$. Let $G \sim \mathcal{P}_{n,\Delta}$ and let $\mathcal{P}^s_{n,\Delta}$ be the conditional distribution on G being simple. It is easy to see that $\mathcal{P}^s_{n,\Delta}$ is the uniform distribution over the set of labelled Δ-regular bicoloured graphs. It is well known that the distributions $\mathcal{P}_{n,\Delta}$ and $\mathcal{P}^s_{n,\Delta}$ are contiguous; a property holds asymptotically almost surely[1] over the distribution $\mathcal{G}_{n,\Delta}$ iff it holds asymptotically almost surely over the distribution $\mathcal{P}^s_{n,\Delta}$ (see [36, Section 9.6] for a formal account of contiguity). It has been proved in [35, 52] that $\mathcal{G}_{n,\Delta}$ is contiguous to $\mathcal{P}_{n,\Delta}$ and hence to $\mathcal{P}^s_{n,\Delta}$.

2.1.2 Recap: spin systems

We briefly recall the spin system framework, for more details see Section 1.1. For an integer $q \geq 2$, a q-spin system is specified by a *symmetric* $q \times q$ interaction matrix $\mathbf{B} = (B_{ij})_{i,j \in [q]}$ with non-negative entries, which specify the strength of the interaction between the spins.

For a finite undirected graph $G = (V, E)$, a q-spin system is a probability distribution μ_G over the space Ω_G of all *configurations*, i.e., spin assignments $\sigma : V \to [q]$. The weight of a configuration $\sigma \in \Omega_G$ is the product of neighboring spin interactions, that is,

$$w_G(\sigma) = \prod_{(u,v) \in E} B_{\sigma(u)\sigma(v)}.$$

The distribution μ_G, known as the *Gibbs distribution* is defined as $\mu_G(\sigma) = w_G(\sigma)/Z_G$

[1] Formally, a property can be interpreted as a sequence of events A_n in underlying probability spaces $(\Omega_n, \mathcal{A}_n, \mathcal{P}_n)$. Then, a property holds asymptotically almost surely if $\mathcal{P}_n(A_n) \to 1$ as $n \to \infty$. We will abbreviate asymptotically almost surely by a.a.s..

where Z_G is the *partition function* defined as $Z_G = \sum_{\sigma \in \Omega_G} w_G(\sigma)$. We drop the subscript G when the graph under consideration is clear.

2.2 A moments' approach for spin systems on random graphs

2.2.1 Expressions for the first and second moments

Let $G \sim \mathcal{G}_{n,\Delta}$ where $\mathcal{G}_{n,\Delta}$ was defined in Section 2.1.1. Briefly, $\mathcal{G}_{n,\Delta}$ is the probability distribution over bipartite graphs with $n+n$ vertices formed by taking the union of Δ random perfect matchings. We will denote the two sides of the bipartition of the graphs as V_1, V_2.

For a q-spin system with interaction matrix $\mathbf{B}$, we will study the Gibbs distribution of the random graph G using the moments of the partition function. We first set up our notation.

We denote by $\triangle_t$ the simplex of t-dimensional probability vectors, i.e.,

$$\triangle_t := \{(x_1, x_2, \ldots, x_t) \in \mathbb{R}^t \,|\, \textstyle\sum_{i=1}^t x_i = 1 \text{ and } x_i \geq 0 \text{ for } i = 1, \ldots, t\}. \tag{20}$$

For a configuration $\sigma : V_1 \cup V_2 \to \{1, \ldots, q\}$, we shall denote the set of vertices assigned color i by $\sigma^{-1}(i)$. For $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \triangle_q \cap \left(\frac{1}{n}\mathbb{Z}^q\right)$, let

$$\Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}} := \{\sigma : V \to \{1, \ldots, q\} \,\big|\, |\sigma^{-1}(i) \cap V_1| = \alpha_i n,\, |\sigma^{-1}(i) \cap V_2| = \beta_j n \text{ for } i = 1, \ldots, q\}, \tag{21}$$

that is, configurations in $\Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ assign $\alpha_i n$ and $\beta_i n$ vertices in V_1 and V_2 the spin value i, respectively. We will be interested in the total weight $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ of configurations in $\Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}$, namely

$$Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}} = \textstyle\sum_{\sigma \in \Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}} w_G(\sigma).$$

Thus, the random variable $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ can be interpreted as a conditional partition function on those configurations with vertex empirical distribution given by $\boldsymbol{\alpha}, \boldsymbol{\beta}$. Note that the partition function Z_G is the sum of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$, over the set of all possible $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \triangle_q$.

We will study $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ by looking at the moments $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]$ and $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]$.

We begin with the first moment. Let $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \triangle_q \cap \left(\frac{1}{n}\mathbb{Z}^q\right)$. For $\sigma \in \Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ and a matching between V_1 and V_2, let nx_{ij} denote the number of edges matching vertices in $\sigma^{-1}(i) \cap V_1$

and $\sigma^{-1}(j) \cap V_2$ and set $\mathbf{x} := (x_{11}, \ldots, x_{qq})$. Note that $\mathbf{x}$ is itself a probability vector in $\triangle_{q^2}$, capturing the edge empirical distribution under the configuration σ. In particular, $\mathbf{x}$ must have the right marginals on the sets V_1, V_2. Precisely, $\mathbf{x} \in \mathcal{M}_1(\boldsymbol{\alpha}, \boldsymbol{\beta}) \cap \left(\frac{1}{n}\mathbb{Z}^{q^2}\right)$, where $\mathcal{M}_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$ is the polytope

$$\mathcal{M}_1(\boldsymbol{\alpha}, \boldsymbol{\beta}) := \left\{ \mathbf{x} : \begin{array}{c} \sum_j x_{ij} = \alpha_i \quad (\forall i \in [q]), \ \sum_i x_{ij} = \beta_j \quad (\forall j \in [q]), \\ x_{ij} \geq 0 \quad (\forall (i,j) \in [q]^2) \end{array} \right\}. \tag{22}$$

To make the following expressions more compact, for $i, j \in [q]$, we denote

$$\mathbf{x}_{i\boldsymbol{\cdot}} := (x_{i1}, \ldots, x_{iq}), \quad \mathbf{x}_{\boldsymbol{\cdot}j} := (x_{1j}, \ldots, x_{qj}).$$

We will also use the following notation for multinomial coefficients. For $\mathbf{z} \in \mathbb{R}^t_{\geq 0}$,

$$\binom{zn}{\mathbf{z}n} := \binom{zn}{z_1 n, \ldots, z_t n} \text{ provided that } z_1 + \ldots + z_t = z.$$

Under the convention that $0^0 \equiv 1$, we then have

$$\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}] = \binom{n}{\boldsymbol{\alpha}n}\binom{n}{\boldsymbol{\beta}n}\left(\sum_{\mathbf{x}}\left[\binom{n}{\mathbf{x}n}^{-1}\prod_i \binom{\alpha_i n}{\mathbf{x}_{i\boldsymbol{\cdot}}n}\prod_j \binom{\beta_j n}{\mathbf{x}_{\boldsymbol{\cdot}j}n}\right]\prod_{i,j} B_{ij}^{nx_{ij}}\right)^{\Delta}, \tag{23}$$

where the sum ranges over $\mathbf{x} \in \mathcal{M}_1(\boldsymbol{\alpha}, \boldsymbol{\beta}) \cap \left(\frac{1}{n}\mathbb{Z}^{q^2}\right)$. Let us briefly explain the derivation of (23). The term $\binom{n}{\boldsymbol{\alpha}n}\binom{n}{\boldsymbol{\beta}n}$ accounts for the cardinality of $\Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}$, while the remainder of the expression is $\mathbf{E}_{\mathcal{G}}[w_G(\sigma)]$ for an arbitrary $\sigma \in \Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}$. Since the weight of a configuration is multiplicative over the edges and the matchings are independent, $\mathbf{E}_{\mathcal{G}}[w_G(\sigma)]$ is the Δ-power of the expected contribution of a single matching. The latter is completely determined by $\mathbf{x}$ and is equal to $\prod_{i,j} B_{ij}^{x_{ij}}$, scaled by the probability that the matching induces the prescribed $\mathbf{x}$, i.e., $\binom{n}{\mathbf{x}n}^{-1}\prod_i \binom{\alpha_i n}{\mathbf{x}_{i\boldsymbol{\cdot}}n}\prod_j \binom{\beta_j n}{\mathbf{x}_{\boldsymbol{\cdot}j}n}$.

The second moment $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]$ is completely analogous, once we introduce the right variables. It will be useful though for our later investigations to study the slightly more general $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}_1,\boldsymbol{\beta}_1} Z_G^{\boldsymbol{\alpha}_2,\boldsymbol{\beta}_2}]$. To do this, for $(\sigma_1, \sigma_2) \in \Sigma^{\boldsymbol{\alpha}_1,\boldsymbol{\beta}_1} \times \Sigma^{\boldsymbol{\alpha}_2,\boldsymbol{\beta}_2}$, we need to compute $\mathbf{E}_{\mathcal{G}}[w_G(\sigma_1) w_G(\sigma_2)]$. Let

$$\gamma_{ik} = |\sigma_1^{-1}(i) \cap \sigma_2^{-1}(k) \cap V_1|/n, \quad \delta_{jl} = |\sigma_1^{-1}(j) \cap \sigma_2^{-1}(l) \cap V_2|/n.$$

The vectors $\boldsymbol{\gamma} := (\gamma_{11}, \ldots, \gamma_{qq})$ and $\boldsymbol{\delta} = (\delta_{11}, \ldots, \delta_{qq})$ capture the overlap of configurations in V_1 and V_2, respectively. In particular, observe that $\boldsymbol{\gamma}$ and $\boldsymbol{\delta}$ satisfy the marginal constraints

$$\boldsymbol{\gamma} \in \mathcal{M}_1(\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2),\ \boldsymbol{\delta} \in \mathcal{M}_1(\boldsymbol{\beta}_1, \boldsymbol{\beta}_2),$$

where the polytope $\mathcal{M}_1(\cdot, \cdot)$ is given by (22). To capture $w_G(\sigma_1)w_G(\sigma_2)$, for a matching between V_1 and V_2, let ny_{ikjl} denote the number of edges matching vertices in $\sigma_1^{-1}(i) \cap \sigma_2^{-1}(k) \cap V_1$ and $\sigma_1^{-1}(j) \cap \sigma_2^{-1}(l) \cap V_2$. The vector $\mathbf{y} = (y_{1111}, \ldots, y_{qqqq})$ must lie in the polytope $\mathcal{M}_2(\boldsymbol{\gamma}, \boldsymbol{\delta})$, where

$$\mathcal{M}_2(\boldsymbol{\gamma}, \boldsymbol{\delta}) := \left\{ \mathbf{y} \ : \ \begin{array}{c} \sum_{j,l} y_{ikjl} = \gamma_{ik} \ \big(\forall (i,k) \in [q]^2\big),\ \sum_{i,k} y_{ikjl} = \delta_{jl} \ \big(\forall (j,l) \in [q]^2\big), \\ y_{ikjl} \geq 0 \ \ \big(\forall (i,k,j,l) \in [q]^4\big). \end{array} \right\} \tag{24}$$

Again, to make the expressions compact, for $i, k, j, l \in [q]$, we denote

$$\mathbf{y}_{ik\cdot} := (y_{ik11}, \ldots, y_{ikqq}), \quad \mathbf{y}_{\cdot jl} := (y_{11jl}, \ldots, y_{qqjl})$$

Under the convention $0^0 \equiv 1$, we then have

$$\begin{aligned} &\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}_1,\boldsymbol{\beta}_2} Z_G^{\boldsymbol{\alpha}_2,\boldsymbol{\beta}_2}] = \\ &\quad \sum_{\boldsymbol{\gamma},\boldsymbol{\delta}} \binom{n}{\boldsymbol{\gamma} n} \binom{n}{\boldsymbol{\delta} n} \Bigg(\sum_{\mathbf{y}} \Bigg[\binom{n}{\mathbf{y} n}^{-1} \prod_{i,k} \binom{\gamma_{ik} n}{\mathbf{y}_{ik\cdot} n} \prod_{j,l} \binom{\delta_{jl} n}{\mathbf{y}_{\cdot jl} n} \Bigg] \prod_{ikjl} (B_{ij} B_{kl})^{n y_{ikjl}} \Bigg)^{\Delta}, \end{aligned} \tag{25}$$

where the sums range over $\boldsymbol{\gamma} \in \mathcal{M}_1(\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2)$, $\boldsymbol{\delta} \in \mathcal{M}_1(\boldsymbol{\beta}_1, \boldsymbol{\beta}_2)$, $\mathbf{y} \in \mathcal{M}_2(\boldsymbol{\gamma}, \boldsymbol{\delta})$. The first line in (25) accounts for the cardinality of $\Sigma^{\boldsymbol{\alpha}_1,\boldsymbol{\beta}_1} \times \Sigma^{\boldsymbol{\alpha}_2,\boldsymbol{\beta}_2}$, while the second line is $\mathbf{E}_{\mathcal{G}}[w_G(\sigma_1)w_G(\sigma_2)]$ for $(\sigma_1, \sigma_2) \in \Sigma^{\boldsymbol{\alpha}_1,\boldsymbol{\beta}_1} \times \Sigma^{\boldsymbol{\alpha}_2,\boldsymbol{\beta}_2}$ with the prescribed $\boldsymbol{\gamma}, \boldsymbol{\delta}$. Since the weight of a configuration is multiplicative over the edges and the matchings are independent, $\mathbf{E}_{\mathcal{G}}[w_G(\sigma_1)w_G(\sigma_2)]$ is the Δ-power of the expected weight of a single matching. The latter is completely determined by $\mathbf{y}$ and is equal to $\prod_{i,k,j,l}(B_{ij}B_{kl})^{y_{ikjl}}$, scaled by the probability that the matching induces the prescribed $\mathbf{y}$.

By restricting to the case $\boldsymbol{\alpha}_1 = \boldsymbol{\alpha}_2 = \boldsymbol{\alpha}$, $\boldsymbol{\beta}_1 = \boldsymbol{\beta}_2 = \boldsymbol{\beta}$ we obtain $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]$. Note that only the range of the sums in (25) change (the regions that $\boldsymbol{\gamma}, \boldsymbol{\delta}$ are allowed to lie); other than that, the expressions are identical. In fact, we have the following simple remark.

Remark 7. *The second moment of the partition function of G is given by* (25)*, where the sum ranges over all $\boldsymbol{\gamma}, \boldsymbol{\delta} \in \triangle_{q^2}$. More generally (and precisely), for $\Sigma \subseteq \triangle_q \times \triangle_q$, denote by*

$$Z_G(\Sigma) = \sum_{(\boldsymbol{\alpha},\boldsymbol{\beta})\in\Sigma} Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}, \text{ so that } \big(Z_G(\Sigma)\big)^2 = \sum_{(\boldsymbol{\alpha}_1,\boldsymbol{\beta}_1),(\boldsymbol{\alpha}_2,\boldsymbol{\beta}_2)\in\Sigma} Z_G^{\boldsymbol{\alpha}_1,\boldsymbol{\beta}_1} Z_G^{\boldsymbol{\alpha}_2,\boldsymbol{\beta}_2}.$$

The second moment $\mathbf{E}_{\mathcal{G}}\big[\big(Z_G(\Sigma)\big)^2\big]$ is given by (25)*, where now the first sum ranges over all*

$$(\boldsymbol{\gamma},\boldsymbol{\delta}) \in \bigcup_{(\boldsymbol{\alpha}_1,\boldsymbol{\beta}_1),(\boldsymbol{\alpha}_2,\boldsymbol{\beta}_2)\in\Sigma} \mathcal{M}_1(\boldsymbol{\alpha}_1,\boldsymbol{\alpha}_2) \times \mathcal{M}_1(\boldsymbol{\beta}_1,\boldsymbol{\beta}_2).$$

The following equally simple remark will be much more crucial. The second moment can be viewed as the first moment of a paired-spin model with interaction matrix $\mathbf{B}\otimes\mathbf{B}$. Indeed, we can interpret $B_{ij}B_{kl}$ as the edge activity between the paired spins (i,k) and (j,l). Since this observation will be a crucial component in our later arguments, we state the following lemma for further expliciteness.

Lemma 19. *For a graph G and an interaction matrix $\mathbf{B}$, let $Z_{G,\mathbf{B}}$ be the partition function of G in the spin system specified by the interaction matrix $\mathbf{B}$. It holds that*

$$\mathbf{E}_{\mathcal{G}}\big[(Z_{G,\mathbf{B}})^2\big] = \mathbf{E}_{\mathcal{G}}\big[Z_{G,\mathbf{B}\otimes\mathbf{B}}\big].$$

Proof. As discussed in Remark 7, we have

$$\mathbf{E}_{\mathcal{G}}\big[(Z_{G,\mathbf{B}})^2\big] = \sum_{\boldsymbol{\gamma},\boldsymbol{\delta}} \binom{n}{\boldsymbol{\gamma}n}\binom{n}{\boldsymbol{\delta}n}\Bigg(\sum_{\mathbf{y}}\Bigg[\binom{n}{\mathbf{y}n}^{-1}\prod_{i,k}\binom{\gamma_{ik}n}{\mathbf{y}_{ik\bullet}n}\prod_{j,l}\binom{\delta_{jl}n}{\mathbf{y}_{\bullet jl}n}\Bigg]\prod_{i,k,j,l}(B_{ij}B_{kl})^{ny_{ikjl}}\Bigg)^{\Delta}, \tag{26}$$

where the sums range over $\boldsymbol{\gamma},\boldsymbol{\delta} \in \triangle_{q^2} \cap \big(\frac{1}{n}\mathbb{Z}^{q^2}\big)$ and $\mathbf{y} \in \mathcal{M}_2(\boldsymbol{\gamma},\boldsymbol{\delta}) \cap \big(\frac{1}{n}\mathbb{Z}^{q^4}\big)$. Similarly,

$$\mathbf{E}_{\mathcal{G}}\big[Z_{G,\mathbf{B}\otimes\mathbf{B}}\big] = \sum_{\overline{\boldsymbol{\alpha}},\overline{\boldsymbol{\beta}}} \binom{n}{\overline{\boldsymbol{\alpha}}n}\binom{n}{\overline{\boldsymbol{\beta}}n}\Bigg(\sum_{\overline{\mathbf{x}}}\Bigg[\binom{n}{\overline{\mathbf{x}}n}^{-1}\prod_{i'}\binom{\overline{\alpha}_{i'}n}{\overline{\mathbf{x}}_{i'\bullet}n}\prod_{j'}\binom{\overline{\beta}_{j'}n}{\overline{\mathbf{x}}_{\bullet j'}n}\Bigg]\prod_{i',j'}(\mathbf{B}\otimes\mathbf{B})_{i'j'}^{n\bar{x}_{i'j'}}\Bigg)^{\Delta}, \tag{27}$$

where the sums range over $\overline{\boldsymbol{\alpha}},\overline{\boldsymbol{\beta}} \in \triangle_{q^2} \cap \big(\frac{1}{n}\mathbb{Z}^{q^2}\big), \overline{\mathbf{x}} \in \mathcal{M}_1(\overline{\boldsymbol{\alpha}},\overline{\boldsymbol{\beta}}) \cap \big(\frac{1}{n}\mathbb{Z}^{q^4}\big)$ and the products over $i',j' \in [q^2]$. The map $f : (i,k) \mapsto q(i-1)+k$ gives a one-to-one correspondence between $(i,k) \in [q]\times[q]$ and $i' \in [q^2]$, and similarly for $(j,l) \in [q]\times[q]$ and $j' \in [q^2]$. Observe also $B_{ij}B_{kl} = (\mathbf{B}\otimes\mathbf{B})_{f(i,k)f(j,l)}$ and that the map f gives a natural correspondence between the entries of $\boldsymbol{\gamma}$ and $\overline{\boldsymbol{\alpha}}$, $\boldsymbol{\delta}$ and $\overline{\boldsymbol{\beta}}$, $\mathbf{y}$ and $\overline{\mathbf{x}}$. We thus have a one-to-one correspondence between $(\overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \overline{\mathbf{x}})$ and $(\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y})$. Thus the expressions (26) and (27) are equal, as wanted. $\square$

2.2.2 The exponential order of the moments

In this section, we study the limits of $\frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big]$ and $\frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2\big]$ as $n \to \infty$. These limits exist under relatively mild conditions as we discuss below (see Remark 8). Roughly, the calculation of these limits involves expanding the factorials using Stirling's approximation and finding the maximum of the resulting function. The intuition here is that the terms in the sums (23) and (25) are typically exponential in n, but have polynomially many terms. This gives that their asymptotic order is determined by a few terms close to their dominating terms.

Under the usual conventions that $\ln 0 \equiv -\infty$ and $0 \ln 0 \equiv 0$, we obtain the following:

$$\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \Psi_1^{\mathbf{B}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) := \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big] = \max_{\mathbf{x} \in \mathcal{M}_1(\boldsymbol{\alpha},\boldsymbol{\beta})} \Upsilon_1(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}), \tag{28}$$

$$\begin{aligned} \text{where} \quad \Upsilon_1(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}) &:= (\Delta - 1) f_1(\boldsymbol{\alpha}, \boldsymbol{\beta}) + \Delta g_1(\mathbf{x}), \\ f_1(\boldsymbol{\alpha}, \boldsymbol{\beta}) &:= \textstyle\sum_i \alpha_i \ln \alpha_i + \sum_j \beta_j \ln \beta_j, \\ g_1(\mathbf{x}) &:= \textstyle\sum_{i,j} x_{ij} \ln B_{ij} - \sum_{i,j} x_{ij} \ln x_{ij}. \end{aligned}$$

And for the second moment:

$$\Psi_2(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \Psi_2^{\mathbf{B}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) := \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2\big] = \max_{\substack{\boldsymbol{\gamma} \in \mathcal{M}_1(\boldsymbol{\alpha},\boldsymbol{\alpha}) \\ \boldsymbol{\delta} \in \mathcal{M}_1(\boldsymbol{\beta},\boldsymbol{\beta})}} \max_{\mathbf{y} \in \mathcal{M}_2(\boldsymbol{\gamma},\boldsymbol{\delta})} \Upsilon_2(\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y}), \tag{29}$$

$$\begin{aligned} \text{where} \quad \Upsilon_2(\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y}) &:= (\Delta - 1) f_2(\boldsymbol{\gamma}, \boldsymbol{\delta}) + \Delta g_2(\mathbf{y}), \\ f_2(\boldsymbol{\gamma}, \boldsymbol{\delta}) &:= \textstyle\sum_{i,k} \gamma_{ik} \ln \gamma_{ik} + \sum_{j,l} \delta_{jl} \ln \delta_{jl}, \\ g_2(\mathbf{y}) &:= \textstyle\sum_{i,k,j,l} y_{ikjl} \ln(B_{ij} B_{kl}) - \sum_{i,k,j,l} y_{ikjl} \ln y_{ikjl}. \end{aligned}$$

Remark 8. *We expand at a lower level of technical detail our discussion on the existence of the limits* $\frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big]$ *and* $\frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2\big]$. *Assuming that the Hessian of these functions is negative definite at the respective maxima, one can use a Gaussian integral around the maximizers to compute the asymptotics of* $\mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big]$ *and* $\mathbf{E}_{\mathcal{G}}\big[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2\big]$, *i.e., the tails of these sums are negligible. The technique has been coined as Laplace's method (see* *for example [15, Chapter 4] or [69, Chapter IX]). This is exactly what we are going to do in Chapter 4 and the reason behind the Hessian condition in Theorem 14. It should be noted however that the limits of the logarithms of the moments may be justified under much milder conditions, see for example [48]. We do not follow this path here since we will need the asymptotics of the moments.*

Remark 9. *The maximization in the first moment depends only on the function $g_1(\mathbf{x})$ which is strictly concave in the convex region over the convex polytope $\mathcal{M}_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$. Hence, for any fixed $\boldsymbol{\alpha}, \boldsymbol{\beta}$, the global maximum of $\Upsilon_1(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x})$ with respect to $\mathbf{x}$ is achieved at a unique point. Similarly, for any fixed $\boldsymbol{\gamma}, \boldsymbol{\delta}$, the maximum of $\Upsilon_2(\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y})$ with respect to $\mathbf{y}$ is achieved at a unique point. Crucially for our considerations in Section 4.3, if $\boldsymbol{\alpha}, \boldsymbol{\beta}$ are global maximizers of Ψ_1, the global maximum of $\Upsilon_2(\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y})$ with respect to $\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y}$ is also achieved at a unique point, see Lemma 31 in Section 3.3.3.*

A notational convention that we have adopted silently so far is perhaps useful to explicitly mention: the indices i, k "point" to the set V_1, while indices j, l "point" to the set V_2.

2.3 Bits on the small subgraph conditioning method

As described in Section 1.3.3, our goal is to obtain concentration for the random variables $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ around their expectation for each dominant phase $(\boldsymbol{\alpha}, \boldsymbol{\beta})$, cf. Definition 6. We will use this information in Chapter 4 to show that the largest contribution to the partition function of a random bipartite regular graph G comes from configurations whose spin frequencies on the two sides of the graph are close to a dominant phase $(\boldsymbol{\alpha}, \boldsymbol{\beta})$. In turn, this will allow us to obtain a gadget to derive our inapproximability results whenever there are more than one dominant phases.

Back to the concentration of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$, we mentioned in Section 1.3.3 that to apply the second moment method, one first needs to compute the ratio $\mathbf{E}_{\mathcal{G}}\big[\big(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big)^2\big]/\big(\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]\big)^2$ and show that it is asymptotically equal to 1. Computing this ratio is highly nontrivial and requires a thorough understanding of the functions Ψ_1, Ψ_2 and the location of their maxima. In this sense, the technical bottleneck of our arguments are proved in Chapter 3. To fast forward slightly, the analysis of the functions Ψ_1 and Ψ_2, together with intensive asymptotic calculations yield that the ratio $\mathbf{E}_{\mathcal{G}}\big[\big(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big)^2\big]/\big(\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]\big)^2$ converges to a constant greater than 1 (cf. equation (13)). This does not allow to obtain concentration of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ around its expectation based on that information alone, as well as to argue the presence of multiple dominant phases in a single graph G.

For probability spaces corresponding to random Δ-regular graphs, there is a well-established method to overcome this failure of the second moment method, known as the small subgraph conditioning method. The small subgraph conditioning method was introduced in [60] to prove that a random Δ-regular contains asymptotically almost surely (a.a.s.) a Hamilton cycle. Roughly, the method provides a way to get a.a.s. results for a random variable Y when the second moment method fails, in the particular case (though common in the random regular graph setting) where the ratio of the second moment of Y to the first moment squared converges to a constant strictly greater than 1. The method tries to attribute the variance of Y to the presence of subgraph structures. In particular, for random Δ-regular graphs the only interesting subgraph structures are short cycles (cycles whose length is constant), which, as it typically turns out, cause all the variance of Y. By taking into account the variance caused by the short cycles, one can obtain strong bounds on Y (in terms of its expectation), as we shall see shortly in the upcoming Theorem 20 and Lemma 21.

The small subgraph conditioning method was first used to analyze spin systems on random regular graphs in [56] and was subsequently used in [63, 27]. As in these works, our goal is to apply the method on $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ for a dominant phase $(\boldsymbol{\alpha},\boldsymbol{\beta})$, and explain the variance of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ by looking at the variance of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ conditioned on the presence of short cycles. In particular, for $i = 1, 2, \ldots$, X_{in} will denote the number of cycles of length i in a random bipartite Δ-regular graph $G \sim \mathcal{G}_{n,\Delta}$.

The method is captured by the following theorem. We note here that there are several versions of the method, we will use a combined version of the respective theorems in [60, 35]. The theorem can be extrapolated from [35], after combining [35, Lemma 1, Remark 4, Remark 9]. The notation $[X]_m$ refers to the m-th order falling factorial of the variable X.

We shall discuss the theorem statement and how it applies to our setting afterwards.

Theorem 20. *For $i = 1, 2, \ldots$, let $\lambda_i > 0$ and $\delta_i > -1$ be constants and assume that for each n there are random variables X_{in}, $i = 1, 2, \ldots$, and Y_n, all defined on the same probability space $\mathcal{G} = \mathcal{G}_n$ such that X_{in} is non-negative integer valued, $Y_n \geq 0$ and $\mathbf{E}\big[Y_n\big] > 0$ (for n sufficiently large). Furthermore, the following hold:*

(A1) $X_{in} \xrightarrow{d} Z_i$ *as $n \to \infty$, jointly for all i, where $Z_i \sim \mathrm{Po}(\lambda_i)$ are independent Poisson random variables;*

(A2) *for every finite sequence $j_1, \ldots, j_m$ of non-negative integers,*

$$\frac{\mathbf{E}_{\mathcal{G}}\big[Y_n[X_{1n}]_{j_1}\cdots[X_{mn}]_{j_m}\big]}{\mathbf{E}_{\mathcal{G}}\big[Y_n\big]} \to \prod_{i=1}^{m}\Big(\lambda_i\big(1+\delta_i\big)\Big)^{j_i} \qquad \text{as } n \to \infty; \tag{30}$$

(A3) $\sum_i \lambda_i \big(\delta_i^{(s)}\big)^2 < \infty$;

(A4) $\mathbf{E}_{\mathcal{G}}\big[Y_n^2\big]/\big(\mathbf{E}_{\mathcal{G}}\big[Y_n\big]\big)^2 \leq \exp\Big(\sum_i \lambda_i \delta_i^2\Big) + o(1)$ *as $n \to \infty$;*

Then:

(C1) *Let $r(n)$ be a function such that $r(n) \to 0$ as $n \to \infty$. It holds that $Y_n > r(n)\mathbf{E}_{\mathcal{G}}\big[Y_n\big]$ asymptotically almost surely.*

(C2) *The following convergence in distribution holds:*

$$\frac{Y_n}{\mathbf{E}_{\mathcal{G}}\big[Y_n\big]} \xrightarrow{d} W = \prod_{i}^{\infty}(1+\delta_i)^{Z_i}\exp\big(-\lambda_i\delta_i\big). \tag{31}$$

This and the convergence in (A1) hold jointly. The infinite product defining W converges a.s. and in L^2, with

$$\mathbf{E}[W] = 1 \text{ and } \mathbf{E}\big[W^2\big] = \lim_{n\to\infty} \mathbf{E}_{\mathcal{G}}\big[Y_n^2\big]/\big(\mathbf{E}_{\mathcal{G}}\big[Y_n\big]\big)^2.$$

Moreover, $W > 0$ a.s..

The conclusion (C1) of Theorem 20 is essentially due to [60], while the conclusion (C2) is due to [35] (note that (C2) implies (C1)). (C1) is generally sufficient when the interest is in proving concentration of a random variable Y_n within a polynomial factor from its expectation (the upper tail can typically be handled using Markov's inequality). (C2) gives the distributional limit of the random variable Y_n and as a consequence gives a handle on the fluctuations from the expectation in terms of the limiting distribution of the short cycle counts.

For our inapproximabilty results, to prove the properties of the gadget used in the reduction we will need (roughly) to argue that the ratio $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}/Z_G$ can be made arbitrary close (within $\pm\varepsilon$ for every $\varepsilon > 0$ and all sufficiently large n) to a fixed constant for all dominant phases $(\boldsymbol{\alpha},\boldsymbol{\beta})$. For this type of approximation, we need sharper bounds on $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$, better than polynomial factors. In this sense, conclusion (C1) will not be sufficient for our purposes. Conclusion (C2) is closer to what we need, but not exactly. Instead, we will use the following lemma, which will allow us to explicitly connect the random variables $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ with the cycle counts X_{in} of the graph G. The lemma is implicit in the arguments of [60] and observed in [35, p.5], where it is discussed without proof in a specific setting, and as such we write and prove a formal statement in the setup of Theorem 20. The proof follows Janson's proof of Theorem 20 but uses a slightly different finish.

Lemma 21. *Assume that the conditions in Theorem 20 hold. For an integer $m > 0$, let*

$$W_{mn} = \prod_{i=1}^{m} \big(1+\delta_i\big)^{X_{in}} \exp\big(-\lambda_i\delta_i\big).$$

Then, for every $\varepsilon > 0$, it holds that

$$\lim_{m\to\infty} \limsup_{n\to\infty} \Pr_{\mathcal{G}_n}\left(\left[\left|\frac{Y_n}{\mathbf{E}_{\mathcal{G}_n}\big[Y_n\big]} - W_{mn}\right| > \varepsilon\right]\right) = 0. \tag{32}$$

Proof of Lemma 21. Wlog we will assume that $\mathbf{E}_{\mathcal{G}_n}\big[Y_n\big] = 1$. We will prove that

$$\limsup_{n\to\infty} \Pr_{\mathcal{G}_n}\big(\big[|Y_n - W_{mn}| > \varepsilon\big]\big) \le \frac{1}{4}\varepsilon^{-2}\Big[\exp\big(\sum_{i=1}^{\infty}\lambda_i\delta_i^2\big) - \exp\big(\sum_{i=1}^{m}\lambda_i\delta_i^2\big)\Big]. \tag{33}$$

This clearly gives the statement of the lemma, since by assumption (A3) of Theorem 20, the lhs is finite and goes to 0 as $m\to\infty$. To prove (33), we follow [35, Proof of Theorem 1] up to a certain point but avoid the use of Skorokhod's theorem in the argument. Janson's proof goes as follows. For a positive integer m define the functions

$$f_n(x_1,\dots,x_m) = \mathbf{E}_{\mathcal{G}_n}[Y_n \mid X_{1n} = x_1,\dots,X_{mn} = x_m],$$
$$f_\infty(x_1,\dots,x_m) = \lim_{n\to\infty} f_n(x_1,\dots,x_m) = \prod_{i=1}^{m}(1+\delta_i)^{x_i} e^{-\lambda_i\delta_i}. \tag{34}$$

The second equality follows by assumption (A2) of Theorem 20 and [35, Lemma 1]. Define also the random variable

$$Y_n^{(m)} = \mathbf{E}_{\mathcal{G}_n}[Y_n \mid X_{1n},\dots,X_{mn}].$$

Using assumptions (A1) and (A2), Fatou's Lemma and that $Y_n^{(m)}$ is a conditional expectation of Y_n, one obtains

$$\limsup_{n\to\infty} \mathbf{E}_{\mathcal{G}_n}\big[|Y_n - Y_n^{(m)}|^2\big] \le \exp\big(\sum_{i=1}^{\infty}\lambda_i\delta_i^2\big) - \exp\big(\sum_{i=1}^{m}\lambda_i\delta_i^2\big),$$

see [35, Equation (5.2)] for details. We now give the main deviation point from Janson's proof, which amounts to proving that for fixed m, we have

$$\lim_{n\to\infty} \Pr_{\mathcal{G}_n}\big(\big[|Y_n^{(m)} - W_{mn}| > \varepsilon\big]\big) = 0. \tag{35}$$

as $n \to \infty$. Fix $M > 0$. By (34), there is N such that for $n \ge N$ it holds that

$$|f_n(x_1,\dots,x_m) - f_\infty(x_1,\dots,x_m)| < \varepsilon \text{ for all integer } x_1,\dots,x_m \in [0,M].$$

It follows that for $n \ge N$, we have

$$\Pr_{\mathcal{G}_n}\big(\big[|Y_n^{(m)} - W_{mn}| > \varepsilon\big]\big) \le \Pr_{\mathcal{G}_n}\Big(\bigcup_{i=1}^{m}\big[X_{in} > M\big]\Big)$$

Note that as $n \to \infty$, the rhs by assumption (A1) converges to $\Pr\big(\bigcup_{i=1}^{m}\big[Z_i > M\big]\big)$. The latter can be made arbitrarily small by letting $M \to \infty$. This proves (35).

The final step is to bound

$$\begin{aligned}
&\limsup_{n\to\infty} \Pr_{\mathcal{G}_n}\big(\big[|Y_n - W_{mn}| > \varepsilon\big]\big)\\
&\qquad\le \limsup_{n\to\infty} \Pr_{\mathcal{G}_n}\big(\big[|Y_n - Y_n^{(m)}| > \varepsilon/2\big]\big) + \limsup_{n\to\infty} \Pr_{\mathcal{G}_n}\big(\big[|Y_n^{(m)} - W_{mn}| > \varepsilon/2\big]\big)\\
&\qquad\le \frac{1}{4}\varepsilon^{-2}\Big[\exp\big(\sum_{i=1}^{\infty}\lambda_i\delta_i^2\big) - \exp\big(\sum_{i=1}^{m}\lambda_i\delta_i^2\big)\Big] + 0,
\end{aligned}$$

which finishes the proof of (33). $\square$

2.4 Notes

We refer the reader to the survey of Wormald [70] for a thorough account of models of random regular graphs. Another standard source is the book of Bollobás [10]. For an account which is targeted towards applications for factor models, look at [49, Chapter 9].

CHAPTER III

SECOND MOMENT ANALYSIS USING MATRIX NORMS

In Chapter 2, for a q-spin system with interaction matrix $\mathbf{B}$, we viewed the partition function of a random Δ-regular bipartite graph as a random variable and calculated its first and second moments. In particular, we saw that the exponential order of the moments is determined by certain functions Ψ_1 and Ψ_2 of the spin frequencies.

In this chapter, we use the first moment to identify the configurations with the largest contribution in expectation; essentially, we characterize in a suitable sense the configurations which are "candidates" to be the modes in the Gibbs distribution. As was demonstrated in [56] (in the case of the hard-core model), these candidate modes correspond to fixpoints of the tree recursions and we shall show a derivation in our setting. More importantly, we will use this information to reformulate the first moment as an induced matrix norm (depending on Δ) of the interaction matrix $\mathbf{B}$.

Of course, there is no a priori reason that the expectation argument of the previous paragraph indeed gives the modes in the Gibbs distribution. We supplement it by showing that the exponential order of the second moment matches the exponential order of the first moment squared. This is done in a surprisingly straightforward way; we observe that the second moment can also be reformulated as the induced matrix norm of $\mathbf{B} \otimes \mathbf{B}$. Using that induced matrix norms are multiplicative over tensor products, the desired alignment between the second moment and the first moment is obtained.

A far more refined variance analysis in Chapter 4 (based on the small subgraph conditioning method) will yield that the identified configurations are indeed the (only) modes in the Gibbs distribution with high probability over the choice of the random graph. To carry out however the variance analysis, we will need more information about the moments and, in particular, the second order behavior of the functions Ψ_1, Ψ_2 around their maxima.

We will connect this problem to the stability of fixpoints to the tree recursions, a connection which will be amply utilized to study the modes for spin models of specific interest in Chapter 6.

3.1 Preliminaries

3.1.1 Recap: first and second moments

For a q-spin system with interaction matrix $\mathbf{B}$, our goal is to understand the Gibbs distribution on a random Δ-regular bipartite graph $G = (V, E)$ (with bipartition $V = V_1 \cup V_2$) by looking at the distribution of spin values in V_1 and V_2. Let $n = |V_1| = |V_2|$. For a configuration $\sigma : V \to [q]$, we denote the set of vertices assigned spin i by $\sigma^{-1}(i)$. For q-dimensional probability vectors $\boldsymbol{\alpha}, \boldsymbol{\beta}$, let

$$\Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}} = \left\{ \sigma : V \to \{1, \dots, q\} \,\middle|\, |\sigma^{-1}(i) \cap V_1| = \alpha_i n,\ |\sigma^{-1}(i) \cap V_2| = \beta_i n \text{ for } i = 1, \dots, q \right\},$$

that is, configurations in $\Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ assign $\alpha_i n$ and $\beta_i n$ vertices in V_1 and V_2 the spin value i, respectively. We will be interested in the total weight $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ of configurations in $\Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}$, namely

$$Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}} = \textstyle\sum_{\sigma \in \Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}} w(\sigma).$$

We study $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ by looking at the moments $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]$ and $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]$, where the expectation is over the distribution of the random Δ-regular bipartite graph, from hereon denoted by $\mathcal{G}$.

Denote the leading term of the first and second moments as:

$$\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \Psi_1^{\mathbf{B}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) := \lim_{n \to \infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big].$$
$$\Psi_2(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \Psi_2^{\mathbf{B}}(\boldsymbol{\alpha}, \boldsymbol{\beta}) := \lim_{n \to \infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\Big[\big(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\big)^2\Big].$$

We will be interested in the global maximizers of Ψ_1, which (at least in expectation) capture the configurations which have the largest contribution in the partition function of a random Δ-regular bipartite graph. We also recall the following definition from Section 1.3.

Definition 6. *For a q-spin system with interaction matrix $\mathbf{B}$, a* dominant phase *is a pair $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ of q-dimensional probability vectors which maximize $\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$.*

3.1.2 Basic definitions: matrix norms

We will reformulate the maxima of the first and second moments in terms of matrix norms. We recall here the basic definitions regarding matrix norms. The usual vector norms are denoted as:

$$\|\mathbf{x}\|_p = (\sum_{i=1}^{n} x_i^p)^{1/p}.$$

We will use the subordinate matrix norm (also known as the induced matrix norm) which will be denoted as $\|\cdot\|_{p\to q'}$ (this is different than the usual notation $\|\cdot\|_{p\to q}$ used in the literature, since we have reserved q for the number of spins) and is defined as:

$$\|\mathbf{A}\|_{p\to q'} = \max_{\|\mathbf{x}\|_p=1} \|\mathbf{A}\,\mathbf{x}\|_{q'}.$$

Note that if $\mathbf{A}$ has non-negative entries then one can restrict the maximization to $\mathbf{x}$ with non-negative entries. A well-known example of an induced norm is the spectral norm $\|\cdot\|_{2\to 2}$.

It was proved by Bennett [7, Proposition 10.1] that induced norms $\|\cdot\|_{p\to q'}$ with $p \le q'$ are multiplicative over Kronecker product. Precisely, for matrices $\mathbf{A}_1, \mathbf{A}_2$ it holds that

$$\|\mathbf{A}_1 \otimes \mathbf{A}_2\|_{p\to q'} = \|\mathbf{A}_1\|_{p\to q'}\|\mathbf{A}_2\|_{p\to q'}. \tag{36}$$

This property will be crucial for our analysis of the second moment of $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$.

We will also need duality of norms. Recall that the norms $\|\cdot\|_p$ and $\|\cdot\|_{p'}$ are dual if $\frac{1}{p} + \frac{1}{p'} = 1$. It then holds

$$\max_{\|\mathbf{y}\|_p=1} |\mathbf{y}^\intercal \mathbf{x}| = \|\mathbf{x}\|_{p'}. \tag{37}$$

Note that (37) is an immediate consequence of Hölder's inequality. For more details, the reader is referred to [33].

3.2 Tree recursions, first moment, and matrix norms

A key component in our arguments is to get a good handle on the function Ψ_1. We displayed in Section 1.3.2 that the important idea is to define a new function Φ which captures in an appropriate way the maximum of Ψ_1. The function Φ allows us to use matrix norms in our analysis of the first and second moments. For the reader's convenience, we recall next the definition of the function Φ.

Let $p = \Delta/(\Delta-1)$. For non-negative q-dimensional vectors $\mathbf{r} := (R_1, \ldots, R_q), \mathbf{c} := (C_1, \ldots, C_q)$, define $\Phi(\mathbf{r}, \mathbf{c})$ by:

$$\exp\big(\Phi(\mathbf{r},\mathbf{c})/\Delta\big) = \frac{\mathbf{r}^\intercal \mathbf{B}\mathbf{c}}{\|\mathbf{r}\|_p \|\mathbf{c}\|_p}.$$

We will show that the critical points of Φ and Ψ_1 match in the sense that there is a one-to-one correspondence between them and their values are equal at the corresponding critical points. The full statement is contained in Theorem 5 given below, but the important element for the analysis of the second moment is captured in the following lemma:

Lemma 22. *For every model on bipartite random regular graphs, it holds that*

$$\max_{\boldsymbol{\alpha},\boldsymbol{\beta} \in \triangle_q} \Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \max_{\mathbf{r},\mathbf{c}} \Phi(\mathbf{r}, \mathbf{c}).$$

Therefore, to determine the dominant phases of Ψ_1, it suffices to study Φ. As we displayed in Setion 1.3.2, the maximum of Φ can be compactly expressed in terms of matrix norms. Recall the one-line derivation:

$$\max_{\mathbf{r},\mathbf{c}} \exp\big(\Phi(\mathbf{r},\mathbf{c})/\Delta\big) = \max_{\mathbf{c}} \max_{\mathbf{r}} \frac{\mathbf{r}^\intercal \mathbf{B}\mathbf{c}}{\|\mathbf{r}\|_p \|\mathbf{c}\|_p} = \max_{\mathbf{c}} \frac{\|\mathbf{B}\mathbf{c}\|_\Delta}{\|\mathbf{c}\|_p} = \|\mathbf{B}\|_{p \to \Delta}, \tag{10}$$

where the second equality follows from matrix norm duality (namely, apply (37) for $\mathbf{x} = \mathbf{B}\mathbf{c}$ and $\mathbf{y} = \mathbf{r}$).

Hence, the dominant phases of Ψ_1 can be expressed in terms of matrix norms:

$$\max_{\boldsymbol{\alpha},\boldsymbol{\beta} \in \triangle_q} \exp\big(\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta})/\Delta\big) = \|\mathbf{B}\|_{\frac{\Delta}{\Delta-1} \to \Delta}. \tag{38}$$

The proof of Lemma 22 is immediate using Theorem 5 stated in Section 1.3.2. We reiterate the theorem for the reader's convenience, and prove it in the remainder of this section. We note that the analysis of the second moment does not need any extra pieces other than Lemma 22.

Theorem 5. *There is a one-to-one correspondence between the fixpoints of the tree recursions and the critical points of Φ (both considered for non-negative $\mathbf{r} = (R_1, \ldots, R_q)$, $\mathbf{c} = (C_1, \ldots, C_q)$ in the projective space, that is, up to scaling by constants as in Remark 3).*

The following transformation $(\mathbf{r},\mathbf{c}) \mapsto (\boldsymbol{\alpha},\boldsymbol{\beta})$ *given by:*

$$\alpha_i = \frac{R_i^{\Delta/(\Delta-1)}}{\sum_i R_i^{\Delta/(\Delta-1)}} \quad \text{and} \quad \beta_j = \frac{C_j^{\Delta/(\Delta-1)}}{\sum_j C_j^{\Delta/(\Delta-1)}} \tag{11}$$

yields a one-to-one correspondence between the critical points of Φ *and the critical points of* Ψ_1 *(in the region defined by* $\alpha_i \ge 0, \beta_j \ge 0$ *and* $\sum_i \alpha_i = 1, \sum_j \beta_j = 1$*).*

Moreover, for the corresponding critical points $(\mathbf{r},\mathbf{c})$ *and* $(\boldsymbol{\alpha},\boldsymbol{\beta})$ *one has*

$$\Phi(\mathbf{r},\mathbf{c}) = \Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta}). \tag{12}$$

Finally, for ergodic $\mathbf{B}$ *(irreducible and ergodic), the local maxima of* Φ *and* Ψ_1 *happen at the critical points (that is, there are no local maxima on the boundary).*

In the following proof of Theorem 5, we give forward references to the ingredients which we will prove next.

Proof of Theorem 5. Lemmas 25 and 26 give the connection between the critical points of Ψ_1 and the fixpoints of the tree recursions. Lemmas 27 and 28 give the connection between the critical points of Ψ_1 and Φ and show that the values agree on the corresponding critical points. Finally, Lemmas 29 and 30 show that the maxima happen in the interior (that is, for $R_i > 0, C_j > 0$ in the case of Φ and for $\alpha_i > 0, \beta_j > 0$ in the case of Ψ_1). $\square$

3.2.1 Preliminaries on maximum-entropy distributions

Let $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ be non-negative vectors in $\mathbb{R}^q$ such that

$$\sum_i \alpha_i = 1 \quad \text{and} \quad \sum_j \beta_j = 1. \tag{39}$$

For $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ that satisfy (39) let

$$g(\alpha_1,\dots,\alpha_q,\beta_1,\dots,\beta_q) = \max \sum_{i=1}^{q}\sum_{j=1}^{q} x_{ij}(\ln(B_{ij}) - \ln x_{ij}), \tag{40}$$

where the maximum is taken over non-negative x_{ij}'s such that

$$\alpha_i = \sum_j x_{ij} \quad \text{and} \quad \beta_j = \sum_i x_{ij}. \tag{41}$$

Lemma 23. *The maximum of the right-hand-side of* (40) *is achieved at unique* x_{ij}*. The* x_{ij} *are given by*

$$x_{ij} = B_{ij} R_i C_j, \tag{42}$$

where $\mathbf{r}$ *and* $\mathbf{c}$ *satisfy*

$$R_i \sum_{j=1}^{q} B_{ij} C_j = \alpha_i \quad \text{and} \quad C_j \sum_{i=1}^{q} B_{ij} R_i = \beta_j, \tag{43}$$

and

$$\begin{aligned} \sum_{j=1}^{q} B_{ij} C_j = 0 &\implies R_i = 0; \\ \sum_{i=1}^{q} B_{ij} R_i = 0 &\implies C_j = 0. \end{aligned} \tag{44}$$

The value of g*, in terms of* R_i*'s and* C_j*'s, is given by*

$$g(\alpha_1, \ldots, \alpha_q, \beta_1, \ldots, \beta_q) = -\sum_{i=1}^{q} \sum_{j=1}^{q} B_{ij} R_i C_j \ln(R_i C_j). \tag{45}$$

Proof. From strict concavity of $-x \ln x$ it follows that the right-hand side of (40) has a unique critical point (if there were two critical points then the segment between the points lies in the linear space defined by (41); the function has a zero derivative on both ends of the segment; and the second derivative of the function is negative on the segment; a contradiction).

Using the method of Lagrange multipliers we obtain that the critical points of the right-hand side of (40) are x_{ij} given by (42) where R_i's and C_j's are solutions of (43). We can make any solution of (43) satisfy (44): if $\sum_{j=1}^{q} B_{ij} C_j = 0$ then set $R_i = 0$ (and symmetrically, if $\sum_{i=1}^{q} B_{ij} R_i = 0$ then set $C_j = 0$). We now argue that this change does not violate (43). Suppose that after the change for some $k \in [q]$ we have

$$R_k \sum_{j=1}^{q} B_{kj} C_j \neq \alpha_k. \tag{46}$$

Then $i = k$ (since only R_i changed) and since $\sum_{j=1}^{q} B_{ij} C_j = 0$ we also have $\alpha_i = 0$, a contradiction (with (46)). Now suppose that after the change for some $j \in [q]$ we have

$$C_j \sum_{k=1}^{q} B_{kj} R_k \neq \beta_j. \tag{47}$$

Then $B_{ij} > 0$ and $C_j > 0$ (otherwise changing R_i would not violate (47)). This then implies $\sum_{j=1}^{q} B_{ij}C_j > B_{ij}C_j > 0$, a contradiction. Thus the change does not violate (43).

Equation (45) is obtained by substituting (43) into (40). □

Remark 10. *Scaling all the R_i's up by the same factor while scaling all the C_j's down by the same factor preserves* (42) *and* (43)*. Modulo such scaling the R_i's and C_j's are unique, since the x_{ij}'s are unique and* (42) *determines the R_i's and C_j's once one value (say R_1) is fixed (here we use the fact that the matrix of the model is ergodic).*

Remark 11. *Note that the condition* (39) *translates (using* (43)*) into the following condition on R_i's and C_j's*

$$\sum_{i=1}^{q}\sum_{j=1}^{q} B_{ij}R_iC_j = 1. \tag{48}$$

Our goal now is to see how the value of (40) changes when we perturb α_i's and β_j's. We are going to view them as functions of a new variable z. All differentiation in this section will be with respect to z. Note that to stay in the subspace defined by (39) we should have, in particular,

$$\sum_i \alpha_i' = \sum_i \alpha_i'' = 0 \quad \text{and} \quad \sum_j \beta_j' = \sum_j \beta_j'' = 0. \tag{49}$$

Differentiating (43) we obtain

$$\sum_{j=1}^{q} B_{ij}(R_iC_j)' = \alpha_i' \quad \text{and} \quad \sum_{i=1}^{q} B_{ij}(R_iC_j)' = \beta_j'. \tag{50}$$

The following ratio of (43) and (50) will be useful later

$$\frac{\alpha_i'}{\alpha_i} = \frac{R_i'}{R_i} + \frac{\sum_{j=1}^{q} B_{ij}C_j'}{\sum_{j=1}^{q} B_{ij}C_j} \quad \text{and} \quad \frac{\beta_j'}{\beta_j} = \frac{C_j'}{C_j} + \frac{\sum_{i=1}^{q} B_{ij}R_i'}{\sum_{i=1}^{q} B_{ij}R_i}. \tag{51}$$

The scaling freedom for R_i's and C_j's (discussed in Remark 10) is equivalent to increasing all R_i'/R_i's by the same (additive) amount and decreasing all C_i'/C_i by the same (additive) amount. We are going to remove this freedom by requiring

$$\sum_{i=1}^{q} \alpha_i \frac{R_i'}{R_i} = \sum_{j=1}^{q} \beta_j \frac{C_j'}{C_j}. \tag{52}$$

(Recall that we study the effect of perturbing g when we change α_i's and β_j's; equation (52) just fixes the corresponding change in R_i's and C_j's.)

Now we compute the derivatives of g.

Lemma 24. *We have*

$$g' = -\sum_{i=1}^{q}(\ln R_i)\alpha_i' - \sum_{j=1}^{q}(\ln C_j)\beta_j', \tag{53}$$

and

$$g'' = -\sum_{i=1}^{q}\frac{R_i'}{R_i}\alpha_i' - \sum_{j=1}^{q}\frac{C_j'}{C_j}\beta_j' - \sum_{i=1}^{q}(\ln R_i)\alpha_i'' - \sum_{j=1}^{q}(\ln C_j)\beta_j''. \tag{54}$$

Proof. Using $(f\ln f)' = (1+\ln f)f'$ and equations (50) and (49) we obtain

$$g' = -\sum_{i=1}^{q}\sum_{j=1}^{q} B_{ij}(1+\ln(R_iC_i))(R_iC_j)' = -\sum_{i=1}^{q}(\ln R_i)\alpha_i' - \sum_{j}(\ln C_j)\beta_j'.$$

Differentiating (53) we obtain (54). □

Note the expressions (53) and (54) are independent of the choice of scaling of R_i's and C_j's (this follows from (49)). The particular tying of R_i'/R_i's and C_j'/C_j's to α_i' and β_j' (given by (52)) will be useful later.

3.2.2 Critical points of Ψ_1 and the tree recursions

In this section we establish the connection between the critical points of Ψ_1 and the fixpoints of the tree recursions.

Lemma 25. *Let $\boldsymbol{\alpha},\boldsymbol{\beta}$ be a critical point of $\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta})$ in the subspace defined by* (39)*. Let* $\mathbf{r},\mathbf{c}$ *be given by* (43)*. Then*

$$\alpha_i \propto R_i^{\Delta/(\Delta-1)} \quad \text{and} \quad \beta_j \propto C_j^{\Delta/(\Delta-1)}. \tag{55}$$

Consequently, $\mathbf{r},\mathbf{c}$ *satisfy the tree recursions stated in the introduction:*

$$R_i \propto \Big(\sum_{j=1}^{q} B_{ij}C_j\Big)^{\Delta-1} \quad \text{and} \quad C_j \propto \Big(\sum_{i=1}^{q} B_{ij}R_i\Big)^{\Delta-1}. \tag{8}$$

Proof. At the critical points of Ψ the first derivative of Ψ has to vanish for all α_i''s and β_j''s from the subspace defined by (49), that is,

$$\begin{aligned}\Psi' &= (\Delta-1)\Big(\sum_{i=1}^{q}(1+\ln\alpha_i)\alpha_i' + \sum_{j=1}^{q}(1+\ln\beta_j)\beta_j'\Big) - \Delta\Big(\sum_{i=1}^{q}(\ln R_i)\alpha_i' + \sum_{j=1}^{q}(\ln C_j)\beta_j'\Big)\\ &= \sum_{i=1}^{q}\big((\Delta-1)(1+\ln\alpha_i) - \Delta\ln R_i\big)\alpha_i' + \sum_{j=1}^{q}\big((\Delta-1)(1+\ln\beta_j) - \Delta\ln C_j\big)\beta_j' = 0,\end{aligned} \tag{56}$$

where the R_i's and C_j's are given by (43). Inspecting (56) we see that $(\Delta-1)(1+\ln\alpha_i)-\Delta\ln R_i$ have the same value. Indeed, if two of them, say with indices i_1, i_2, had different values then we could increase α_{i_1} and decrease α_{i_2} by the same infinitesimal amount and violate (56). Similarly, $(\Delta-1)(1+\ln\beta_j)-\Delta C_j$ have the same value and hence we have (55). Plugging (55) into (43) one obtains (8). $\square$

Lemma 26. *Let* $(\mathbf{r},\mathbf{c})$ *be a solution of the tree recurrences* (8). *Let* $(\boldsymbol{\alpha},\boldsymbol{\beta})$ *be given by* (11). *Then* $(\boldsymbol{\alpha},\boldsymbol{\beta})$ *is a critical point of* $\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta})$ *in the subspace defined by* (39).

Proof. Let

$$Z_R := (\Delta-1)(1+\ln\alpha_i)-\Delta\ln R_i = (\Delta-1)\Big(1-\ln\sum_{i=1}^{q} R_i^{(\Delta+1)/\Delta}\Big),$$

where the second equality follows from (11). Note that Z_R is independent of the choice of i. Similarly let

$$Z_C := (\Delta-1)(1+\ln\beta_j)-\Delta\ln C_j = (\Delta-1)\Big(1-\ln\sum_{j=1}^{q} C_j^{(\Delta+1)/\Delta}\Big).$$

For perturbations of $\boldsymbol{\alpha},\boldsymbol{\beta}$ in the subspace given by (39) we have

$$\begin{aligned}\Psi_1'(\boldsymbol{\alpha},\boldsymbol{\beta}) = \sum_{i=1}^{q}\big((\Delta-1)(1+\ln\alpha_i)-\Delta\ln R_i\big)\alpha_i' + \sum_{j=1}^{q}\big((\Delta-1)(1+\ln\beta_j)-\Delta\ln C_j\big)\beta_j' \\ = Z_R\sum_{i=1}^{q}\alpha_i' + Z_C\sum_{j=1}^{q}\beta_j' = 0,\end{aligned}$$

and hence $(\boldsymbol{\alpha},\boldsymbol{\beta})$ is a critical point. $\square$

3.2.3 Value of Ψ_1 at the critical points

Lemma 27. *Let* $(\boldsymbol{\alpha},\boldsymbol{\beta})$ *be critical point of* $\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta})$. *Let* $(\mathbf{r},\mathbf{c})$ *be given by* (43). *Then*

$$\Phi(\mathbf{r},\mathbf{c}) = \Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta}). \tag{12}$$

Moreover, $(\mathbf{r},\mathbf{c})$ *is a critical point of* $\Phi(\mathbf{r},\mathbf{c})$.

Proof. We have (see equation (45))

$$\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta}) = (\Delta-1)\Big(\sum_{i=1}^{q}\alpha_i\ln\alpha_i + \sum_{j=1}^{q}\beta_j\ln\beta_j\Big) - \Delta\sum_{i=1}^{q}\sum_{j=1}^{q} B_{ij}R_iC_j\ln(R_iC_j). \tag{57}$$

At the critical points we have (see equation (55))

$$\alpha_i = \frac{R_i^{\Delta/(\Delta-1)}}{\sum_{i=1}^q R_i^{\Delta/(\Delta-1)}} \quad \text{and} \quad \beta_j = \frac{C_j^{\Delta/(\Delta-1)}}{\sum_{j=1}^q C_j^{\Delta/(\Delta-1)}}. \tag{58}$$

Plugging (43) into (57) we obtain

$$\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta}) = (\Delta-1)\Big(\sum_{i=1}^q \alpha_i \ln \alpha_i + \sum_{j=1}^q \beta_j \ln \beta_j\Big) - \Delta\Big(\sum_{i=1}^q \alpha_i \ln R_i + \sum_{j=1}^q \beta_j \ln C_j\Big) =$$
$$\sum_{i=1}^q \alpha_i \ln \frac{\alpha_i^{\Delta-1}}{R_i^\Delta} + \sum_{j=1}^q \beta_j \ln \frac{\beta_j^{\Delta-1}}{C_j^\Delta} = -(\Delta-1)\Big[\ln\Big(\sum_{i=1}^q R_i^{\Delta/(\Delta-1)}\Big) + \ln\Big(\sum_{j=1}^q C_j^{\Delta/(\Delta-1)}\Big)\Big], \tag{59}$$

where in the last equality we used (58) and the fact that α_i's and β_j's sum to 1. Recall that

$$\sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j = \sum_{i=1}^q \alpha_i = 1, \tag{60}$$

and hence the following is obtained by adding zero to the right-hand side of (59):

$$\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta}) = \Delta \ln\Big(\sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j\Big) - (\Delta-1)\Big[\ln\Big(\sum_{i=1}^q R_i^{\Delta/(\Delta-1)}\Big) + \ln\Big(\sum_{j=1}^q C_j^{\Delta/(\Delta-1)}\Big)\Big] = \Phi(\mathbf{r},\mathbf{c}).$$

Now we argue that $(\mathbf{r},\mathbf{c})$ is a critical point of $\Phi(\mathbf{r},\mathbf{c})$. We have

$$\frac{\partial}{\partial R_i}\Phi(\mathbf{r},\mathbf{c}) = \Delta \frac{\sum_{j=1}^q B_{ij} C_j}{\sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j} - (\Delta-1)\frac{\frac{\Delta}{\Delta-1} R_i^{1/(\Delta-1)}}{\sum_{i=1}^q R_i^{\Delta/(\Delta-1)}}. \tag{61}$$

Using (58), (43), and (39) we obtain

$$\frac{\partial}{\partial R_i}\Phi(\mathbf{r},\mathbf{c}) = \Delta \frac{\alpha_i}{R_i} - \Delta \frac{\alpha_i}{R_i} = 0.$$

The same argument yields

$$\frac{\partial}{\partial C_j}\Phi(\mathbf{r},\mathbf{c}) = \Delta \frac{\sum_{i=1}^q B_{ij} R_i}{\sum_{i=1}^q \sum_{j=1}^q B_{ij} R_i C_j} - (\Delta-1)\frac{\frac{\Delta}{\Delta-1} C_j^{1/(\Delta-1)}}{\sum_{j=1}^q C_j^{\Delta/(\Delta-1)}} = 0. \tag{62}$$

and hence $\mathbf{r},\mathbf{c}$ is a critical point of Φ. □

Lemma 28. *Let* $(\mathbf{r},\mathbf{c})$ *be a critical point of* $\Phi(\mathbf{r},\mathbf{c})$*. Let* $\boldsymbol{\alpha},\boldsymbol{\beta}$ *be given by* (11)*. Then* $\boldsymbol{\alpha},\boldsymbol{\beta}$ *is a critical point of* $\Psi_1(\boldsymbol{\alpha},\boldsymbol{\beta})$ *in the subspace defined by* (39)*.*

Proof. At a critical point of Φ we have that (61) is zero for $i \in [q]$. Note that the denominators do not depend on i and hence we have

$$R_i^{1/(\Delta-1)} \propto \sum_{j=1}^{q} B_{ij} C_j.$$

Similarly, from (62) we obtain

$$C_j^{1/(\Delta-1)} \propto \sum_{i=1}^{q} B_{ij} R_i.$$

Hence $(\mathbf{r}, \mathbf{c})$ satisfy the tree recursions. Now we use Lemma 26 to conclude that $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ is a critical point of $\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$ in the subspace defined by (39). $\square$

3.2.4 Local maxima of Ψ_1 are in the interior

In this section we show that for models with ergodic (irreducible and aperiodic) interaction matrix $\mathbf{B}$ the maximum of $\Phi(\mathbf{r}, \mathbf{c})$ is achieved in the interior. A symmetric matrix is irreducible if the graph whose edges correspond to non-zero edges of $\mathbf{B}$ is connected. A symmetric matrix is aperiodic if the graph whose edges correspond to non-zero edges of $\mathbf{B}$ has an odd cycle.

Lemma 29. *Assume that* $\mathbf{B}$ *is ergodic. Let* $(\mathbf{r}, \mathbf{c}) \neq 0$ *be a local maximum of* Φ *in the region* $\mathbf{r}, \mathbf{c} \geq 0$*. Then* $R_i > 0$ *for all* $i \in [q]$ *and* $C_j > 0$ *for all* $j \in [q]$*.*

Proof. Suppose not, that is, we have a maximum that has a zero on some coordinate of $\mathbf{r}$ or $\mathbf{c}$. From the ergodicity of $\mathbf{B}$ we have that there exist $i, j \in [q]$ such that i) $R_i = 0$, $C_j > 0$, and $B_{ij} > 0$ or ii) $R_i > 0$, $C_j = 0$, and $B_{ij} > 0$. (Suppose not. Let $Z_R \subseteq [q]$ be the set of i such that $R_i = 0$. Similarly let $Z_C \subseteq [q]$ be the set of j such that $C_j = 0$. If neither i) nor ii) happens then non-zero B_{ij} are possibly between $i \in Z_R$ and $j \in Z_C$ and $i \in [q] \setminus Z_R$ and $j \in [q] \setminus Z_C$. Thus in $\mathbf{B}^2$ the non-zero $(B^2)_{ij}$ are possibly between $i, j \in Z_R$ and $i, j \in [q] \setminus Z_R$. Thus $\mathbf{B}$ is not ergodic.) W.l.o.g. assume that it is the case i) (the case ii) is handled analogously).

The derivative of Φ w.r.t. R_i is (we are using $R_i = 0$)

$$\frac{\partial}{\partial R_i} \Phi(\mathbf{r}, \mathbf{c}) = \Delta \frac{\sum_{j=1}^{q} B_{ij} C_j}{\sum_{i=1}^{q} \sum_{j=1}^{q} B_{ij} R_i C_j} > \Delta \frac{B_{ij} C_j}{\sum_{i=1}^{q} \sum_{j=1}^{q} B_{ij} R_i C_j} > 0,$$

and hence we are not at a maximum, a contradiction. $\square$

Lemma 30. *Assume that* $\mathbf{B}$ *is ergodic. Let* $\boldsymbol{\alpha}, \boldsymbol{\beta} \geq 0$ *be a local maximum of* $\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$ *in the subspace defined by* (39). *Then* $\alpha_i > 0$ *for all* $i \in [q]$ *and* $\beta_j > 0$ *for all* $j \in [q]$.

Proof. It will be useful to view Ψ_1 as a function of $(\mathbf{r}, \mathbf{c})$. Because of Lemma 23 we have $(\mathbf{r}, \mathbf{c})$ satisfying (60) and (44) (and any such $(\mathbf{r}, \mathbf{c})$ yields $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ satisfying (39)). From (59), we have

$$\begin{aligned}\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \sum_{i=1}^{q}\sum_{j=1}^{q} B_{ij}R_iC_j \left((\Delta - 1)\Big(\ln\Big(\sum_{j=1}^{q} B_{ij}C_j\Big) + \ln\Big(\sum_{i=1}^{q} B_{ij}R_i\Big)\Big) - \ln R_i - \ln C_j \right) \\ =: \hat{\Psi}_1(\mathbf{r}, \mathbf{c}).\end{aligned}$$

If $\mathbf{r}$ has a zero coordinate then, by ergodicity of $\mathbf{B}$ there exists $k, \ell \in [q]$ such that (i) $R_k = 0$, $C_\ell > 0$, and $B_{k\ell} > 0$ or (ii) $R_k > 0$, $C_\ell = 0$, and $B_{k\ell} > 0$ (see the argument in the proof of Lemma 29). W.l.o.g. it is the case (i).

Note that we have

$$\frac{\partial}{\partial R_k} \sum_{i=1}^{q}\sum_{j=1}^{q} B_{ij}R_iC_j = \sum_{j=1}^{q} B_{kj}C_j \geq B_{k\ell}C_\ell > 0. \tag{63}$$

We have

$$\begin{aligned}\frac{\partial}{\partial R_k}\hat{\Psi}_1 = &\sum_{j=1}^{q} B_{kj}C_j \Big((\Delta - 1)\ln\Big(\sum_{i=1}^{q} B_{ij}R_i\Big) - \ln C_j \Big) \\ &+ \Big((\Delta - 1)\ln\Big(\sum_{j=1}^{q} B_{kj}C_j\Big) - \ln R_k \Big)\Big(\sum_{j=1}^{q} B_{kj}C_j\Big) + (\Delta - 2)\sum_{j=1}^{q} B_{kj}C_j.\end{aligned} \tag{64}$$

The first sum in (64) is finite since if $C_j > 0$ then $\sum_{i=1}^{q} B_{ij}R_i > 0$ (using (44)); if $C_j = 0$ then the contribution of the term to the sum is zero (we are using the usual convention $0 \ln 0 = 0$). The second term in (64) has value $+\infty$ since $\ln R_k = -\infty$ and (63). Finally, the last term in (64) is finite and hence we have $\frac{\partial}{\partial R_k}\hat{\Psi}_1 = +\infty$.

Recall that $C_\ell > 0$ and hence (using (44)):

$$\frac{\partial}{\partial C_\ell} \sum_{i=1}^{q}\sum_{j=1}^{q} B_{ij}R_iC_j = \sum_{i=1}^{q} B_{i\ell}R_i > 0. \tag{65}$$

Finally, we argue that $\frac{\partial}{\partial C_\ell}\hat{\Psi}_1$ is finite. We have (analogously to (64))

$$\begin{aligned}\frac{\partial}{\partial C_\ell}\hat{\Psi}_1 = &\sum_{i=1}^{q} B_{i\ell}C_i \Big((\Delta - 1)\ln\Big(\sum_{j=1}^{q} B_{ij}C_j\Big) - \ln R_i \Big) \\ &+ \Big((\Delta - 1)\ln\Big(\sum_{i=1}^{q} B_{i\ell}R_i\Big) - \ln C_\ell \Big)\Big(\sum_{i=1}^{q} B_{i\ell}R_i\Big) + (\Delta - 2)\sum_{i=1}^{q} B_{i\ell}R_i.\end{aligned}$$

The first and third terms in (66) are finite by the same argument as for (64). In the second term we use (65) and $C_\ell > 0$.

Now we increase R_k by an infinitesimal amount and change C_ℓ to maintain (48) (and hence (39)). (This is possible because both C_ℓ and R_k change the value of (48), see equations (63) and (65).) This change will increase $\hat{\Psi}_1$ and hence Ψ_1 contradicting the local maximality of $\boldsymbol{\alpha}, \boldsymbol{\beta}$. $\square$

3.3 Second-moment analysis for dominant phases

In this section, we prove Theorem 4, which we restate for convenience here.

Theorem 4. *For any spin system, for all* $\Delta \geq 3$,

$$\max_{\boldsymbol{\alpha},\boldsymbol{\beta}} \Psi_2(\boldsymbol{\alpha}, \boldsymbol{\beta}) = 2 \max_{\boldsymbol{\alpha},\boldsymbol{\beta}} \Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta}).$$

In particular, for every dominant phase $(\boldsymbol{\alpha}, \boldsymbol{\beta})$*, it holds that* $\Psi_2(\boldsymbol{\alpha}, \boldsymbol{\beta}) = 2\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$.

Theorem 4 will eventually allow us to prove strong concentration properties for the random variables $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$, whenever $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ is a dominant phase.

The proof of Theorem 4 resides into two main components: (i) the reformulation of the first moment of any spin system as an induced matrix norm of its interaction matrix $\mathbf{B}$, (ii) viewing the second moment as the first moment of a "paired-spin" system with interaction matrix $\mathbf{B} \otimes \mathbf{B}$. Item (ii) allows us to reformulate the second moment as an induced matrix norm of $\mathbf{B} \otimes \mathbf{B}$. The key component then is to connect the induced matrix norms of $\mathbf{B}$ and $\mathbf{B} \otimes \mathbf{B}$, which we can do by using multiplicative properties of induced matrix norms over tensor product.

3.3.1 The second moment as the first moment of a paired-spin model

We have already discussed in Section 2.2.1 that the second moment can be viewed as the first moment of a paired-spin model with interaction matrix $\mathbf{B} \otimes \mathbf{B}$, see Lemma 19. We briefly review the connection.

For the second moment $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]$, one considers a pair of configurations, say σ and σ', which are both constrained to have marginals $\boldsymbol{\alpha}$ on V_1 and $\boldsymbol{\beta}$ on V_2, where $V = V_1 \cup V_2$.

We capture this constraint using a pair of vectors $\boldsymbol{\gamma}, \boldsymbol{\delta}$ corresponding to the overlap between σ and σ', in particular, γ_{ij} (and δ_{ij}) is the number of vertices in V_1 (and V_2, respectively) with spin i in σ and spin j in σ'. Thus, in the second moment, every vertex in V is assigned a pair of spins (i, k) and the interaction in the paired-spin system is given by $\mathbf{B} \otimes \mathbf{B}$.

Let us now see how this translates to the functions $\Psi_1^{\mathbf{B}}, \Psi_2^{\mathbf{B}}$. Recall, $\Psi_1^{\mathbf{B}}$ indicates the dependence of the function Ψ_1 on the interaction matrix $\mathbf{B}$; to simplify the notation we will drop the exponent if it is $\mathbf{B}$. Extrapolating from Section 2.2.1, we have

$$\Psi_2(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \max_{\boldsymbol{\gamma}, \boldsymbol{\delta}} \Psi_1^{\mathbf{B}\otimes\mathbf{B}}(\boldsymbol{\gamma}, \boldsymbol{\delta}), \tag{66}$$

where the optimization in (66) is constrained to non-negative $\boldsymbol{\gamma}, \boldsymbol{\delta}$ such that

$$\textstyle\sum_i \gamma_{ik} = \alpha_k, \qquad \sum_k \gamma_{ik} = \alpha_i, \qquad \sum_j \delta_{j\ell} = \beta_\ell \quad \text{and} \quad \sum_\ell \delta_{j\ell} = \beta_j. \tag{67}$$

3.3.2 Analyzing the second moment: proof of Theorem 4

We are now ready to prove Theorem 4.

Proof of Theorem 4. Replacing the four constraints in (67) with the weaker constraints $\boldsymbol{\gamma}, \boldsymbol{\delta} \in \triangle_{q^2}$ can only increase the value of (66) and hence

$$\max_{\boldsymbol{\alpha},\boldsymbol{\beta}} \exp(\Psi_2(\boldsymbol{\alpha}, \boldsymbol{\beta})/\Delta) \leq \max_{\boldsymbol{\gamma},\boldsymbol{\delta}\in\triangle_{q^2}} \exp\left(\Psi_1^{\mathbf{B}\otimes\mathbf{B}}(\boldsymbol{\gamma}, \boldsymbol{\delta})/\Delta\right) = \|\mathbf{B}\otimes\mathbf{B}\|_{\frac{\Delta}{\Delta-1}\to\Delta}, \tag{68}$$

where the last equality follows by applying the analogue of (38) for the spin model specified by $\mathbf{B} \otimes \mathbf{B}$. The key fact we now use is that induced norms $\|\cdot\|_{p\to q'}$ with $p \leq q'$ are multiplicative over Kronecker product, see equation (36). Applying (36) in our setting yields

$$\|\mathbf{B}\otimes\mathbf{B}\|_{p\to q'} = \|\mathbf{B}\|_{p\to q'} \|\mathbf{B}\|_{p\to q'}. \tag{69}$$

Therefore,

$$\max_{\boldsymbol{\alpha},\boldsymbol{\beta}} \Psi_2(\boldsymbol{\alpha}, \boldsymbol{\beta}) \leq 2\Delta \log \|\mathbf{B}\|_{\frac{\Delta}{\Delta-1}\to\Delta} = 2 \max_{\boldsymbol{\alpha},\boldsymbol{\beta}} \Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta}). \tag{70}$$

Observe that for any $(\alpha, \boldsymbol{\beta})$ we have $\Psi_2(\boldsymbol{\alpha}, \boldsymbol{\beta}) \geq 2\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$ as a simple implication of $\mathbf{E}[X^2] \geq \mathbf{E}[X]^2$. This yields the equality in (70), as well as the equality for dominant phases. $\square$

3.3.3 Optimal second-moment configuration

We will need more detailed information about the $\boldsymbol{\gamma}, \boldsymbol{\delta}$ which achieve equality in Theorem 4 and equation (66). The following lemma is true whenever $\mathbf{B}$ is regular (and hence for antiferromagnetic models as well, cf. Definition 11).

Lemma 31. *Assume that $\mathbf{B}$ is regular (non-zero determinant). The $\boldsymbol{\gamma}, \boldsymbol{\delta}$ for which the equality in*

$$\max_{\boldsymbol{\alpha},\boldsymbol{\beta}} \max_{\boldsymbol{\gamma},\boldsymbol{\delta} \text{ satisfying (67)}} \Psi_1^{\mathbf{B}\otimes\mathbf{B}}(\boldsymbol{\gamma}, \boldsymbol{\delta}) = \max_{\boldsymbol{\alpha},\boldsymbol{\beta}} \Psi_1^{\mathbf{B}}(\boldsymbol{\alpha}, \boldsymbol{\beta}), \tag{71}$$

is achieved satisfy (for all $i, j, k, l \in [q]$)

$$\gamma_{ik} = \alpha_i \alpha_k \quad \textit{and} \quad \delta_{jl} = \beta_j \beta_l. \tag{72}$$

Proof. We will have to dig in to the proof of (69) and use (67). Bennett's proof of (69) is the following (our particular values are $q' = \Delta$ and $p = \Delta/(\Delta - 1)$):

$$\begin{aligned}
\|(\mathbf{B}\otimes\mathbf{B})\mathbf{r}\|_{q'} &= \left(\sum_k \sum_i \Big| \sum_j B_{ij} \sum_l B_{kl} R_{jl} \Big|^{q'} \right)^{1/q'} \\
&\le \|\mathbf{B}\|_{p\to q'} \left(\sum_k \left(\sum_j \Big| \sum_l B_{kl} R_{jl} \Big|^{p} \right)^{q'/p} \right)^{1/q'} \\
&\le \|\mathbf{B}\|_{p\to q'} \left(\sum_j \left(\sum_k \Big| \sum_l B_{kl} R_{jl} \Big|^{q'} \right)^{p/q'} \right)^{1/p} \\
&\le \|\mathbf{B}\|^2_{p\to q'} \left(\sum_{j,l} R^p_{jl} \right)^{1/p}.
\end{aligned}$$

Note that in the last inequality one uses $\|\mathbf{B}\ \mathbf{r}\|_{q'} \le \|\mathbf{B}\|_{p\to q'} \|\mathbf{r}\|_p$, applied to the vectors $\mathbf{r}'_j := (R_{j1}, R_{j2}, \ldots, R_{jq})$, for $j = 1, \ldots, q$. Thus if $\mathbf{r}$ is a maximizer of

$$\max_{\mathbf{r}} \frac{\|(\mathbf{B}\otimes\mathbf{B})\mathbf{r}\|_{q'}}{\|\mathbf{r}\|_p}, \tag{73}$$

then the vectors $\mathbf{r}'_j$ are maximizers of

$$\max_{\mathbf{r}'} \frac{\|\mathbf{B}\,\mathbf{r}'\|_{q'}}{\|\mathbf{r}'\|_p}. \tag{74}$$

The same, by symmetry, applies to $\mathbf{r}''_l := (R_{1l}, R_{2l}, \ldots, R_{ql})$, for $l = 1, \ldots, q$.

The second inequality in Bennett's proof is Minkowski's inequality applied to vectors $\mathbf{Br}'_1, \ldots, \mathbf{Br}'_q$. The equality is achieved only if $\mathbf{Br}'_1, \ldots, \mathbf{Br}'_q$ generate space of dimension one, and since $\mathbf{B}$ is regular we have also that $\mathbf{r}'_1, \ldots, \mathbf{r}'_q$ generate space of dimension one. Hence, for a maximizer $\mathbf{r}$ of (73) we have $\mathbf{r} = \mathbf{r}' \otimes \mathbf{r}''$, where $\mathbf{r}'$ and $\mathbf{r}''$ are maximizers of (74). By Theorem 5 (equation (11)) we then have

$$\gamma_{ik} = \alpha'_i \alpha''_k, \tag{75}$$

for the corresponding maximizers of $\Psi_1^{\mathbf{B}\otimes\mathbf{B}}(\boldsymbol{\gamma})$ and $\Psi_1^{\mathbf{B}}(\boldsymbol{\alpha})$. Equation (75) together with constraints

$$\sum_i \gamma_{ik} = \alpha_k \quad \text{and} \quad \sum_k \gamma_{ik} = \alpha_i,$$

implies $\gamma_{ik} = \alpha_i \alpha_k$ (since $\alpha_k = \sum_i \gamma_{ik} = \sum_i \alpha'_i \alpha''_k = \alpha''_k$). The proof of $\delta_{jl} = \beta_j \beta_l$ is analogous. $\square$

3.4 Tree recursions and critical points of the first moment

The second moment results of the previous section will be used to establish that the Gibbs distribution is concentrated at the global maxima of $\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$. To simplify the analysis of the local maxima of Ψ_1 we connect them to attractive fixpoints of the associated tree recursions.

In [56], it was observed that the critical points of Ψ_1 correspond to fixpoints of the following tree recursions (8), which we recall for convenience here:

$$\hat{R}_i \propto \Big(\sum_{j=1}^{q} B_{ij} C_j\Big)^{\Delta-1} \quad \text{and} \quad \hat{C}_j \propto \Big(\sum_{i=1}^{q} B_{ij} R_j\Big)^{\Delta-1}. \tag{8}$$

The fixpoints are those R_i's and C_j's such that $\hat{R}_i \propto R_i$ and $\hat{C}_j \propto C_j$, for all $i, j \in [q]$. In [56], this connection was established for the case for hard-core model, so we give a derivation in our setting in Section 3.2.2.

We call a fixpoint x of a function f a *Jacobian attractive fixpoint* if the Jacobian of f at x has spectral radius less than 1. We say that a critical point $\boldsymbol{\alpha}, \boldsymbol{\beta}$ is a *Hessian local maximum* if the Hessian of Ψ_1 at $\boldsymbol{\alpha}, \boldsymbol{\beta}$ is negative definite. (Note this is a sufficient condition for $\boldsymbol{\alpha}, \boldsymbol{\beta}$ to be a local maximum.)

The purpose of this section is to prove Theorem 2, which we recall for the convenience of the reader here (cf. Theorem 5 for a more detailed version of the below connection.)

Theorem 2. *Jacobian attractive fixpoints of the depth-two tree recursions* (8) *correspond to Hessian local maxima of* Ψ_1. *Moreover, if* $\mathbf{B}$ *is ergodic (irreducible and aperiodic), Hessian local maxima of* Ψ_1 *correspond to Jacobian attractive fixpoints of the depth-two tree recursions.*

Theorem 2 is important for analyzing the global maxima of Ψ_1 for colorings and anti-ferromagnetic Potts model (see Section 6.3). The rest of this section is devoted to the proof of Theorem 2.

3.4.1 Maximum entropy configurations on random regular bipartite graphs

We analyze the critical points by looking at the second derivative. Using $(f \ln f)'' = (f')^2/f + (1 + \ln f)f''$ we have

$$\begin{aligned}
\Psi_1''(\boldsymbol{\alpha}, \boldsymbol{\beta}) &= (\Delta - 1)\sum_{i=1}^{q}\Big((\alpha_i')^2/\alpha_i + (1+\ln\alpha_i)\alpha_i''\Big) - \Delta\sum_{i=1}^{q}\Big(\alpha_i'\frac{R_i'}{R_i} + (\ln R_i)\alpha_i''\Big)\\
&\quad +(\Delta - 1)\sum_{j=1}^{q}\Big((\beta_j')^2/\beta_j + (1+\ln\beta_j)\beta_j''\Big) - \Delta\sum_{j=1}^{q}\Big(\beta_j'\frac{C_j'}{C_j} + (\ln C_j)\beta_j''\Big)\\
&= (\Delta - 1)\sum_{i=1}^{q}(\alpha_i')^2/\alpha_i - \Delta\sum_{i=1}^{q}\alpha_i'\frac{R_i'}{R_i} + \sum_{i=1}^{q}\alpha_i''\Big((\Delta-1)(1+\ln\alpha_i) - \Delta\ln R_i\Big) \qquad (76)\\
&\quad +(\Delta - 1)\sum_{j=1}^{q}(\beta_j')^2/\beta_j - \Delta\sum_{j=1}^{q}\beta_j'\frac{C_j'}{C_j} + \sum_{j=1}^{q}\beta_j''\Big((\Delta-1)(1+\ln\beta_j) - \Delta\ln C_j\Big)\\
&= (\Delta - 1)\sum_{i=1}^{q}(\alpha_i')^2/\alpha_i - \Delta\sum_{i=1}^{q}\alpha_i'\frac{R_i'}{R_i} + (\Delta - 1)\sum_{j=1}^{q}(\beta_j')^2/\beta_j - \Delta\sum_{j=1}^{q}\beta_j'\frac{C_j'}{C_j},
\end{aligned}$$

where the last equality follows from (56) (replacing α_i' by α_i'' and β_j' by β_j''; note that they are both from the same subspace (49)).

Plugging (51) into (76) we obtain

$$\Psi_1''(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \sum_{i=1}^{q}\alpha_i'\left((\Delta-1)\frac{\sum_{j=1}^{q}B_{ij}C_j'}{\sum_{j=1}^{q}B_{ij}C_j} - \frac{R_i'}{R_i}\right) + \sum_{j=1}^{q}\beta_j'\left((\Delta-1)\frac{\sum_{i=1}^{q}B_{ij}R_i'}{\sum_{i=1}^{q}B_{ij}R_i} - \frac{C_j'}{C_j}\right). \tag{77}$$

We are going to use the second partial derivative test (which gives a sufficient condition) to establish maxima of Ψ_1. We will use the following terminology for local maxima established using this method.

Definition 9. *A critical point x of a function $f : \mathcal{M} \to \mathbb{R}$ is called* **Hessian local maximum** *if the Hessian of f at x is negative definite.*

Let $\mathbf{L}$ be the (matrix of) linear map $(r_1, \ldots, r_q, c_1, \ldots, c_q) \mapsto (\hat{r}_1, \ldots, \hat{r}_q, \hat{c}_1, \ldots, \hat{c}_q)$ given by

$$\hat{r}_i = \sum_j \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} c_j \quad \text{and} \quad \hat{c}_j = \sum_i \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} r_i. \tag{78}$$

In the following, we denote by $\mathbf{I}$ the identity matrix of dimension $2q \times 2q$.

Lemma 32. *A critical point $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ is a Hessian local maximum of $\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$ in the subspace defined by* (49) *if and only if* $\mathbf{w}^\intercal(\mathbf{I}+\mathbf{L})((\Delta-1)\mathbf{L}-\mathbf{I})\mathbf{w} < 0$ *for all* $\mathbf{w} = (r_1, \ldots, r_q, c_1, \ldots, c_q)^\intercal$ *such that*

$$\sum_{i=1}^{q} \sqrt{\alpha_i} r_i = 0 \quad \textit{and} \quad \sum_{j=1}^{q} \sqrt{\beta_j} c_j = 0. \tag{79}$$

Proof. To check whether we are at a Hessian local maximum of $\Psi(\boldsymbol{\alpha}, \boldsymbol{\beta})$ we have to have (77) negative for non-zero α_i''s and β_j''s from the subspace defined by (49) and (52).

Let $r_i = \sqrt{\alpha_i} R_i'/R_i$ and $c_j = \sqrt{\beta_j} C_j'/C_j$. Using (51) we have

$$\begin{aligned}
\Psi'' = & \sum_i \alpha_i \left(\frac{R_i'}{R_i} + \frac{\sum_j B_{ij} C_j'}{\sum_j B_{ij} C_j} \right) \left((\Delta - 1) \frac{\sum_j B_{ij} C_j'}{\sum_j B_{ij} C_j} - \frac{R_i'}{R_i} \right) \\
& + \sum_j \beta_j \left(\frac{C_j'}{C_j} + \frac{\sum_i B_{ij} R_i'}{\sum_i B_{ij} R_i} \right) \left((\Delta - 1) \frac{\sum_i B_{ij} R_i'}{\sum_i B_{ij} R_i} - \frac{C_j'}{C_j} \right) \\
= & \sum_i \left(r_i + \sum_j \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} c_j \right) \left(\sum_j (\Delta - 1) \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} c_j - r_i \right) \\
& + \sum_j \left(c_j + \sum_i \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} r_i \right) \left(\sum_i (\Delta - 1) \frac{B_{ij} R_i C_j}{\sqrt{\alpha_i \beta_j}} r_i - c_j \right).
\end{aligned}$$

Let $\mathbf{w} = (r_1, \ldots, r_q, c_1, \ldots, c_q)^\intercal$. In terms of $\mathbf{L}$ and $\mathbf{w}$ we have

$$\Phi'' = \mathbf{w}^\intercal (\mathbf{I} + \mathbf{L})((\Delta - 1)\mathbf{L} - \mathbf{I})\mathbf{w}. \tag{80}$$

We have to examine when (80) is in the subspace defined by (49) and (52), which in terms of r_i's and c_j's become

$$\sum_i \alpha_i' = \sum_j \beta_j' = \sum_i \sqrt{\alpha_i} r_i + \sum_j \sqrt{\beta_j} c_j = 0, \tag{81}$$

$$\sum_i \alpha_i \frac{R_i'}{R_i} - \sum_j \beta_j \frac{C_j'}{C_j} = \sum_i \sqrt{\alpha_i} r_i - \sum_j \sqrt{\beta_j} c_j = 0. \tag{82}$$

We give more detail on the derivation of (81) below. We have

$$\begin{aligned}\sum_i \alpha_i' &= \sum_i \alpha_i \frac{\alpha_i'}{\alpha_i} = \sum_i \alpha_i \left(\frac{R_i'}{R_i} + \frac{\sum_j B_{ij} C_j'}{\sum_j B_{ij} C_j} \right) = \sum_i r_i \sqrt{\alpha_i} + \sum_i \sum_j B_{ij} R_i C_j' \\ &= \sum_i r_i \sqrt{\alpha_i} + \sum_j \frac{c_j}{\sqrt{\beta_j}} \sum_i B_{ij} R_i C_j = \sum_i r_i \sqrt{\alpha_i} + \sum_j c_j \sqrt{\beta_j},\end{aligned}$$

the derivation for $\sum_j \beta_j'$ is analogous. □

3.4.2 Attractive fixpoints of tree recursions

The variables R_i, C_j, α_i, β_j in this section refer to a priori different quantities as the variables in Section 3.4.1. We feel that this conflict is justified since we will establish that they coincide.

For convenience we repeat the tree recursions as stated in the introduction:

$$\hat{R}_i \propto \left(\sum_{j=1}^{q} B_{ij} C_j \right)^{\Delta-1} \quad \text{and} \quad \hat{C}_j \propto \left(\sum_{i=1}^{q} B_{ij} R_j \right)^{\Delta-1}. \tag{8}$$

We are interested in the **fixpoints** of the tree recursions, that is, R_i's and C_j's such that

$$\hat{R}_i \propto R_i \quad \text{and} \quad \hat{C}_j \propto C_j$$

for all $i, j \in [q]$. Note that the fixpoints correspond to the critical points of Ψ_1 (using Theorem 5).

Next we examine the stability of fixpoints. For a continuously differentiable map a sufficient condition for a fixpoint to be attractive is if the spectral radius of the derivative is less than one at the fixpoint. We will use the following terminology for fixpoints whose attractiveness is established using this method.

Definition 10. *A fixpoint x of a function $f : \mathcal{M} \to \mathcal{M}$ is called* **jacobian attractive fixpoint** *if the Jacobian of f at x has spectral radius less than* 1.

Lemma 33. *Let* $(\mathbf{r},\mathbf{c})$ *be a fixpoint of the tree recursions. Let* $\alpha_i = \sum_{j=1}^q B_{ij}R_iC_j$ *and* $\beta_j = \sum_{i=1}^q B_{ij}R_iC_j$ *and let* $\mathbf{L}$ *be the (matrix of the) map defined by* (78)*. We have that* $(\mathbf{r},\mathbf{c})$ *is jacobian attractive if and only if* $(\Delta-1)\mathbf{L}$ *has spectral radius less than* 1 *in the subspace of* $\mathbf{w} = (r_1,\dots,r_q,c_1,\dots,c_q)$ *that satisfy*

$$\sum_{i=1}^q \sqrt{\alpha_i}r_i = 0 \quad \textit{and} \quad \sum_{j=1}^q \sqrt{\beta_j}c_j = 0. \tag{79}$$

Proof. W.l.o.g. we can assume that $(\mathbf{r},\mathbf{c})$ is scaled so that

$$\sum_{i=1}^q\sum_{j=1}^q B_{ij}R_iC_j = 1. \tag{83}$$

Note that the scaling does not affect the value of $\mathbf{L}$ nor does it affect the constraint (79).

When we perturb the R_i's and C_j's and apply one step of the tree recursion we obtain

$$\frac{\hat{R}_i'}{\hat{R}_i} = (\Delta-1)\frac{\sum_{j=1}^q B_{ij}C_j\frac{C_j'}{C_j}}{\sum_{j=1}^q B_{ij}C_j} \quad \text{and} \quad \frac{\hat{C}_j'}{\hat{C}_j} = (\Delta-1)\frac{\sum_{i=1}^q B_{ij}R_i\frac{R_i'}{R_i}}{\sum_{i=1}^q B_{ij}R_i}. \tag{84}$$

We can rewrite (84) as follows

$$\frac{\hat{R}_i'}{\hat{R}_i} = (\Delta-1)\frac{\sum_{j=1}^q B_{ij}R_iC_j\frac{C_j'}{C_j}}{\alpha_i} \quad \text{and} \quad \frac{\hat{C}_j'}{\hat{C}_j} = (\Delta-1)\frac{\sum_{i=1}^q B_{ij}R_iC_j\frac{R_i'}{R_i}}{\beta_j}. \tag{85}$$

The perturbation that scales all R_i's by the same factor does not change the messages (since they are in the projective space) and hence we need to exclude it when studying local stability of (84). Similarly scaling all C_j's by the same factor does not change the messages. We need to locate an invariant subspace of (85) whose complement corresponds to the scaling. We obtain the following subspace (it corresponds to preserving (83)):

$$\sum_{i=1}^q \alpha_i\frac{R_i'}{R_i} = 0 \quad \text{and} \quad \sum_{j=1}^q \beta_j\frac{C_j'}{C_j} = 0. \tag{86}$$

Now we check that (86) is invariant under the map (85), indeed,

$$\sum_{i=1}^q \alpha_i\frac{\hat{R}_i'}{\hat{R}_i} = (\Delta-1)\sum_{i=1}^q\sum_{j=1}^q B_{ij}R_iC_j\frac{C_j'}{C_j} = (\Delta-1)\sum_{j=1}^q \beta_j\frac{C_j'}{C_j} = 0; \tag{87}$$

the argument for $\sum_{j=1}^q \beta_j\frac{\hat{C}_j'}{\hat{C}_j} = 0$ is analogous.

A fixpoint $(R_1,\dots,R_q,C_1,\dots,C_q)$ is Jacobian attractive if the linear transformation

$$\left(\frac{R_1'}{R_1},\dots,\frac{R_q'}{R_q},\frac{C_1'}{C_1},\dots,\frac{C_q'}{C_q}\right) \mapsto \left(\frac{\hat{R}_1'}{\hat{R}_1},\dots,\frac{\hat{R}_q'}{\hat{R}_q},\frac{\hat{C}_1'}{\hat{C}_1},\dots,\frac{\hat{C}_q'}{\hat{C}_q}\right)$$

given by (84) has spectral radius less than 1 in the subspace defined by (86).

Let $r_i = \sqrt{\alpha_i}R_i'/R_i$, $c_j = \sqrt{\beta_j}C_j'/C_j$, $\hat{r}_i = \sqrt{\alpha_i}\hat{R}_i'/\hat{R}_i$, and $\hat{c}_j = \sqrt{\beta_j}\hat{C}_j'/\hat{C}_j$. This linear transformation of variables turns (85) into

$$\hat{r}_i = (\Delta-1)\sum_{j=1}^{q}\frac{B_{ij}R_iC_j}{\sqrt{\alpha_i\beta_j}}c_j \quad\text{and}\quad \hat{c}_j = (\Delta-1)\sum_{i=1}^{q}\frac{B_{ij}R_iC_j}{\sqrt{\alpha_i\beta_j}}r_i. \tag{88}$$

Note that (88) is $(\Delta-1)\mathbf{J}$ where $\mathbf{J}$ is the map defined by (78). The constraint (86) becomes (79). $\square$

3.4.3 Connecting attractive fixpoints to maximum entropy configurations

Now we are ready to prove Theorem 2.

Proof of Theorem 2. Let S be the linear subspace defined by (79) (note that (81) together with (82) define the same subspace). The constraint for the fixpoint to be jacobian attractive is that $(\Delta-1)\mathbf{L}$ on S has spectral radius less than 1. The constraint for the critical point to be Hessian maximum is that the eigenvalues of $(\mathbf{I}+\mathbf{L})((\Delta-1)\mathbf{L}-\mathbf{I})$ on S are negative (see equation (80)).

Note that $\mathbf{L}$ is symmetric and it is a result of tensor product with the matrix $\left(\begin{smallmatrix}0&1\\1&0\end{smallmatrix}\right)$. Hence $\mathbf{L}$ has symmetric real spectrum (symmetry means that if a is an eigenvalue then so is $-a$). Note that S is invariant under $\mathbf{L}$ and hence the spectrum of $\mathbf{L}$ on S is a subset of the spectrum of $\mathbf{L}$ (it is still symmetric real; the restriction wiped out a pair of eigenvalues -1 and 1).

The constraint for the fixpoint to be jacobian attractive, in terms of eigenvalues, is: for each eigenvalue x of $\mathbf{L}$ on S

$$-1 < (\Delta-1)x < 1. \tag{89}$$

The constraint for the critical point to be Hessian maximum, in terms of eigenvalues, is: for each eigenvalue x of $\mathbf{L}$ on S

$$(1+x)\big((\Delta-1)x-1\big) < 0 \quad\text{and}\quad (1-x)\big(-(\Delta-1)x-1\big) < 0, \tag{90}$$

where the second constraint comes from the symmetry of the spectrum (thus $-x$ is an eigenvalue). Note that conditions (89) and (90) are equivalent (since $(1+x)\big((\Delta-1)x-1\big)$ is negative for $-1 < x < 1/(\Delta-1)$). $\square$

CHAPTER IV

THE GADGET

In this chapter, we construct the gadget which we will use to derive our NP-hardness results. The chapter is then devoted to proving the properties of the gadget.

The gadget is a random graph from a graph distribution $\mathcal{G}^r_{n,\Delta}$, which closely resembles the graph distribution $\mathcal{G}_{n,\Delta}$. The main difference is that $\mathcal{G}^r_{n,\Delta}$ is supported on bipartite graphs whose vertices have degrees both Δ and $\Delta - 1$, with r controlling the number of vertices with degree $\Delta - 1$. The vertices of degree $\Delta - 1$ will be crucial to allow us to make connections between the gadgets without exceeding the degree bound Δ.

We will study the random graph distribution $\mathcal{G}^r_{n,\Delta}$ for sufficiently small values of r, allowing us to directly transfer the results from our analysis of the random graph distribution $\mathcal{G}_{n,\Delta}$. In particular, the asymptotics of the first and second moments of the partition function for the two graph distributions will be essentially the same up to easily computable correction factors.

4.1 Construction

Let $\Delta \geq 3$ and n, r be integers with $n > r > 0$. The graph distribution $\mathcal{G}^r_n := \mathcal{G}^r_{n,\Delta}$ is defined as follows.

1. $\mathcal{G}^r_n$ is supported on bipartite graphs. The two parts of the bipartite graph are denoted by $+, -$ and each is partitioned as $U^s \cup W^s$ where $|U^s| = n$, $|W^s| = r$ for $s = \{+, -\}$. U denotes the set $U^+ \cup U^-$ and similarly W denotes the set $W^+ \cup W^-$.

2. To sample $G \sim \mathcal{G}^r_n$, sample uniformly and independently Δ matchings: (i) $(\Delta - 1)$ perfect matchings between $U^+ \cup W^+$ and $U^- \cup W^-$, (ii) a n-matching between U^+ and U^-. The edge set of G is the union of the Δ matchings. Thus, vertices in U have degree Δ, while vertices in W have degree $\Delta - 1$.

The case $r = 0$ will also be critical for our arguments, in which case there are no vertices of degree $\Delta - 1$ and hence the graph distribution is identical to the distribution $\mathcal{G}_n := \mathcal{G}_{n,\Delta}$, which we studied thoroughly in Chapters 2 & 3. We will thus write $\mathcal{G}_n$ instead of $\mathcal{G}_n^0$.

Let $G \sim \mathcal{G}_n^r$ and denote by μ_G the Gibbs distribution on G with interaction matrix $\mathbf{B}$. We let $\mathcal{Q}$ denote the set of dominant phases on a random Δ-regular bipartite graph, that is, the union of the pairs $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ which maximize $\Psi_1^{\mathbf{B}}(\boldsymbol{\alpha}, \boldsymbol{\beta})$ (recall that $\boldsymbol{\alpha}, \boldsymbol{\beta}$ are q-dimensional probability vectors, i.e., $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \triangle_q$ where $\triangle_q$ is the standard $(q-1)$-simplex). We will use $\mathbf{p}$ to denote a dominant phase $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ and for the purposes of this section $\mathbf{p}^+, \mathbf{p}^-$ will be used to denote the vectors $\boldsymbol{\alpha}, \boldsymbol{\beta}$ respectively.

For a configuration $\sigma : U \cup W \to [q]$, the footprint of σ is defined as a pair of q-dimensional probability vectors $(\mathbf{p}_\sigma^+, \mathbf{p}_\sigma^-)$, i.e., $\mathbf{p}_\sigma^+, \mathbf{p}_\sigma^- \in \triangle_q$. The vectors $\mathbf{p}_\sigma^+, \mathbf{p}_\sigma^-$ count the frequencies of the colors in σ in the parts U^+, U^-. Formally, for $s \in \{+, -\}$, the i-th entry of $\mathbf{p}_\sigma^s$ is equal to $|\sigma^{-1}(i) \cap U^s|/n$. The phase of a configuration σ is denoted by $Y(\sigma)$ and is equal to

$$Y(\sigma) := \arg\min_{\mathbf{p} \in \mathcal{Q}} \left\| \mathbf{p}_\sigma^+ - \mathbf{p}^+ \right\|_1 + \left\| \mathbf{p}_\sigma^- - \mathbf{p}^- \right\|_1 . \tag{91}$$

If there are more than one dominant phases that achieve the minimum in (91), any tie breaking criterion may be used, e.g., the lowest indexed phase. Note that the phase of σ depends only on the spins of vertices in U. For $\sigma : U \cup W \to [q]$, denote by σ_W the restriction of σ to vertices in W.

The exact marginal distribution of μ_G on the vertices in W is quite intricate. However, we shall display shortly that, conditioned on the phase of the configuration, it can be well approximated by an appropriate product measure. To do this, recall that every dominant phase $\mathbf{p} = (\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathcal{Q}$ corresponds to a fixpoint $(\mathbf{r}, \mathbf{c})$ of the tree recursions (8) (see Theorem 5). Let $\mathbf{r} = (R_1, \dots, R_q)$ and $\mathbf{c} = (C_1, \dots, C_q)$ and denote by $\hat{\mathbf{r}} = (\hat{R}_1, \dots, \hat{R}_q), \hat{\mathbf{c}} = (\hat{C}_1, \dots, \hat{C}_q)$ the scaled versions of $\mathbf{r}, \mathbf{c}$ respectively so that $\sum_i \hat{R}_i = 1$ and $\sum_i \hat{C}_i = 1$. To avoid overloading the notation, we do not explicitly index the R_i's and C_j's by the phase $\mathbf{p}$. We next define the relevant product measure $\nu_{\mathbf{p}}^{\otimes}(\cdot)$ on the space of spin assignments to vertices in W for a phase $\mathbf{p} \in \mathcal{Q}$. For $\sigma : W \to [q]$ and $\mathbf{p} \in \mathcal{Q}$, let

$$\nu_{\mathbf{p}}^{\otimes}(\sigma) = \prod_{i \in [q]} (\hat{R}_i)^{|\sigma^{-1}(i) \cap W^+|} \prod_{j \in [q]} (\hat{C}_j)^{|\sigma^{-1}(j) \cap W^-|}. \tag{92}$$

We can now state formally the properties of the gadget $G \sim \mathcal{G}_n^r$ that we will need. Let $\Delta \geq 3$. We impose the following conditions on the spin system specified by $\mathbf{B}$:

(H1) $\mathbf{B}$ is regular (non-zero determinant) and ergodic (irreducible and aperiodic).

(H2) Every dominant phase $\mathbf{p} \in \mathcal{Q}$ is a Hessian maximum of $\Psi_1^{\mathbf{B}}(\boldsymbol{\alpha}, \boldsymbol{\beta})$. Equivalently, by Theorem 2, the fixpoint $(\mathbf{r}, \mathbf{c})$ of the tree recursions (8) corresponding to each $\mathbf{p} \in \mathcal{Q}$ is Jacobian stable (note that the validity of this assumption is allowed to vary with the degree Δ).

(H3) The dominant phases $\mathbf{p} \in \mathcal{Q}$ are permutation symmetric. That is, if $\mathbf{p}_1 = (\boldsymbol{\alpha}_1, \boldsymbol{\beta}_1)$ and $\mathbf{p}_2 = (\boldsymbol{\alpha}_2, \boldsymbol{\beta}_2)$, then there exists a $q \times q$ permutation matrix $\mathbf{P}$ such that $\mathbf{B} = \mathbf{PB}$ and $(\boldsymbol{\alpha}_1, \boldsymbol{\beta}_1) = (\mathbf{P}\boldsymbol{\alpha}_2, \mathbf{P}\boldsymbol{\beta}_2)$ or $(\boldsymbol{\alpha}_1, \boldsymbol{\beta}_1) = (\mathbf{P}\boldsymbol{\beta}_2, \mathbf{P}\boldsymbol{\alpha}_2)$. In other words, the dominant phases can be obtained from each other by permuting the spins in a way that $\mathbf{B}$ is left invariant.

The condition (H1) allows to apply in full strength the results of Chapter 3. The condition (H2) allows to compute the asymptotics of the moments using Laplace's method. The condition (H3) allows to argue that the dominant phases have roughly the same contribution to the partition function of $G \sim \mathcal{G}_n^r$.

Remark 12. *Whenever (H3) is true, i.e., the phases are permutation symmetric, it should be intuitively clear that it suffices to check (H2) for a single dominant phase* $\mathbf{p} \in \mathcal{Q}$*; the Hessian condition for the remaining dominant phases follows from symmetry. This is indeed true, see the upcoming Lemma 38, where the condition (H2) is captured explicitly by the magnitude of the eigenvalues of an explicit matrix.*

The rest of this chapter is devoted to the proof of the following theorem, which formally states the properties of the gadget.

Theorem 34. *Let $\Delta \geq 3$ and suppose that the interaction matrix $\mathbf{B}$ satisfies conditions (H1), (H2), (H3). Let r be a fixed constant. Then, for every $\varepsilon > 0$, a random graph $G \sim \mathcal{G}^r_{n,\Delta}$ satisfies with probability $1 - o(1)$ as $n \to \infty$ all of the following:*

1. *For each $\mathbf{p} \in \mathcal{Q}$, $(1-\varepsilon)/|\mathcal{Q}| \leq \mu_G(Y(\sigma) = \mathbf{p}) \leq (1+\varepsilon)/|\mathcal{Q}|$. That is, the phases in the graph G appear with roughly equal probability.*

2. *For each $\mathbf{p} \in \mathcal{Q}$, for all $\eta : W \to [q]$, $\mu_G(\sigma_W = \eta \,|Y(\sigma) = \mathbf{p})/\nu^{\otimes}_{\mathbf{p}}(\eta) \in [1-\varepsilon, 1+\varepsilon]$. That is, conditioned on the phase $\mathbf{p}$ of the configuration, the spins of the vertices in W are roughly independent and the marginal measure on them can be approximated by the measure $\nu^{\otimes}_{\mathbf{p}}(\cdot)$.*

3. *There is no edge between W^+ and W^-. Moreover, there is no vertex in G which has two neighbors in $W^+ \cup W^-$.*

Moreover, G is simple with asymptotically positive probability and the above continue to hold with probability $1 - o(1)$ conditioned on G being simple.

Remark 13. *When r is allowed to vary moderately with n, say $r = o(n^{1/4})$, Items 1 and 3 still hold, however Item 2 is no longer true pointwise. The reason is that the number of configurations η grows with n, which does not allow to conclude Item 2 for every $\eta : W \to [q]$. Instead, one needs to slightly weaken Item 2 to the following statement. For every $\varepsilon > 0$, for every $\mathbf{p} \in \mathcal{Q}$,*

$$\lim_{n\to\infty} \sup_{\eta\in[q]^W} \Pr_{\mathcal{G}^r_n}\left(\frac{\mu_G(\sigma_W = \eta \,|Y(\sigma) = \mathbf{p})}{\nu^{\otimes}_{\mathbf{p}}(\eta)} \notin (1 \pm \varepsilon)\right) = 0. \tag{93}$$

This version will only be used for the AP-reductions in Chapter 7 (along the lines of [63]), where we do need r to vary with n.

4.2 Proving the properties of the gadget

The goal of this section is to give the proof of Theorem 34. To be able to conclude (93) in the case that r varies moderately with n, we will work under the slightly weaker assumption $r = o(n^{1/4})$.

Let $G \sim \mathcal{G}_n^r$. To get a handle on Items 1 and 2 of Theorem 34, we first define the partitioned functions conditioned on a phase $\mathbf{p} \in \mathcal{Q}$. Similar definitions appear in [63]. For a configuration $\eta : W \to [q]$ and $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \triangle_q$, define

$$\begin{aligned} Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}(\eta) &:= \sum_{\sigma;\ \sigma_W = \eta} w_G(\sigma) \mathbf{1}\{(\mathbf{p}_\sigma^+, \mathbf{p}_\sigma^-) = (\boldsymbol{\alpha}, \boldsymbol{\beta})\}, \\ Z_G^{\mathbf{p}}(\eta) &:= \sum_{\boldsymbol{\alpha},\boldsymbol{\beta}} Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}(\eta) \mathbf{1}_{\arg\min_{\mathbf{p}'} \|(\boldsymbol{\alpha},\boldsymbol{\beta}) - (\mathbf{p}'^+, \mathbf{p}'^-)\|_1 = \mathbf{p}}, \\ Z_G^{\mathbf{p}} &:= \sum_{\eta} Z_G^{\mathbf{p}}(\eta). \end{aligned} \tag{94}$$

Let us explain briefly these definitions. $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}(\eta)$ is the contribution from configurations which agree with η on W and their footprint on U is $(\boldsymbol{\alpha}, \boldsymbol{\beta})$. To picture the definition of $Z_G^{\mathbf{p}}(\eta)$, view the phases $\mathcal{Q}$ as modes in the Gibbs distribution which "attract" the configurations closest to them. Note that distances here are only with respect to the spin assignments on U. Therefore, the spin assignment on W induces a partition on the attraction space of the phase $\mathbf{p}$. In this view, $Z_G^{\mathbf{p}}(\eta)$ is simply the contribution of the set indexed by η in the partition. Finally, $Z_G^{\mathbf{p}}$ is the total contribution of the configurations which are attracted by $\mathbf{p}$.

The following equalities display the relevance of these quantities to Theorem 34.

$$\mu_G(Y(\sigma) = \mathbf{p}) = \frac{Z_G^{\mathbf{p}}}{\sum_{\mathbf{p} \in \mathcal{Q}} Z_G^{\mathbf{p}}}, \quad \mu_G(\sigma_W = \eta \,|\, Y(\sigma) = \mathbf{p}) = \frac{Z_G^{\mathbf{p}}(\eta)}{Z_G^{\mathbf{p}}}. \tag{95}$$

It will also be useful to explicitly state how the definitions in (96) degenerate in the case $r = 0$. In this setting there are no vertices of degree $\Delta - 1$ (and hence no set W), so the graph distribution $\mathcal{G}_n^0$ is identical to the graph distribution $\mathcal{G}(n, \Delta)$. The definitions in (94) extend to this setting by simply dropping the argument η. The conditioned partition functions when $r = 0$ are thus given by

$$Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}} = \sum_{\substack{\sigma; \\ (\mathbf{p}_\sigma^+, \mathbf{p}_\sigma^-) = (\boldsymbol{\alpha}, \boldsymbol{\beta})}} w_G(\sigma), \qquad Z_G^{\mathbf{p}} = \sum_{\substack{\boldsymbol{\alpha},\boldsymbol{\beta}; \\ \mathbf{p} = \arg\min_{\mathbf{p}'} \left\|(\boldsymbol{\alpha},\boldsymbol{\beta}) - (\mathbf{p}'^+, \mathbf{p}'^-)\right\|}} Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}. \tag{96}$$

To start, we are going to show that Items 1 and 2 of Theorem 34 hold in expectation. This is the scope of the following lemma. Note that $o(1)$ refers to quantities that tend to 0 as $n \to \infty$.

Lemma 35. *Let r be a fixed constant and let $\mathbf{p}$ be a Hessian dominant phase, i.e., $\mathbf{p} \in \mathcal{Q}$. There exists a constant $C(\mathbf{p})$ such that for every $\eta : W \to [q]$, it holds that*

$$\mathbf{E}_{\mathcal{G}_n^r}\left[Z_G^{\mathbf{p}}(\eta)\right] = \big(1+o(1)\big)C^r \nu_{\mathbf{p}}^{\otimes}(\eta)\mathbf{E}_{\mathcal{G}_n}\left[Z_G^{\mathbf{p}}\right], \text{ and thus } \sup_{\eta}\left|\frac{\mathbf{E}_{\mathcal{G}_n^r}\left[Z_G^{\mathbf{p}}(\eta)\right]}{\mathbf{E}_{\mathcal{G}_n^r}\left[Z_G^{\mathbf{p}}\right]} - \nu_{\mathbf{p}}^{\otimes}(\eta)\right| = o(1). \tag{97}$$

Moreover, when the phases $\mathcal{Q}$ are permutation-symmetric, $\mathbf{E}_{\mathcal{G}_n}\left[Z_G^{\mathbf{p}}\right] = \big(1+o(1)\big)\mathbf{E}_{\mathcal{G}_n}\left[Z_G^{\mathbf{p}'}\right]$ for any phases $\mathbf{p}, \mathbf{p}' \in \mathcal{Q}$ and the constant C in (97) *does not depend on the particular phase $\mathbf{p}$. Consequently, for $\mathbf{p}, \mathbf{p}' \in \mathcal{Q}$*

$$\mathbf{E}_{\mathcal{G}_n^r}\left[Z_G^{\mathbf{p}}\right] = \big(1+o(1)\big)\,\mathbf{E}_{\mathcal{G}_n^r}\left[Z_G^{\mathbf{p}'}\right], \text{ and thus } \frac{\mathbf{E}_{\mathcal{G}_n^r}\left[Z_G^{\mathbf{p}}\right]}{\sum_{\mathbf{p}\in\mathcal{Q}}\mathbf{E}_{\mathcal{G}_n^r}\left[Z_G^{\mathbf{p}}\right]} = \big(1+o(1)\big)\frac{1}{|\mathcal{Q}|}. \tag{98}$$

Proof. The second equalities in each of (97) and (98) follow immediately from the first. The latter may be proved by explicit calculations following the same arguments as in [63, Lemma 3.3] and essentially reduce to arguing that certain sums are dominated by their maximum terms.

It is worthy to note that the first part of the lemma is true even if the phases are not permutation-symmetric, which is not necessarily true for the second part. □

In light of Equations (95), (97) and (98), the path to obtain Items 1 and 2 of Theorem 34 is now paved: it suffices to show that the conditioned partition functions $Z_G^{\mathbf{p}}(\eta)$ are (with positive probability) arbitrarily close to their expectations for large n. Note that we want this to be simultaneously true for all $\mathbf{p}$ and η, that is, for the same graph G. This in turn requires using in full strength a theorem by Janson [35], which is an extension of the small subgraph conditioning method introduced by Robinson and Wormald [60].

We do a quite extensive, and hopefully illuminating, exposition of these theorems and their application in our setting in the next section. For satisfying the reader who is more interested in the proof of Theorem 34, the following lemma is a stripped-down version of the results in Section 2.3, yet at the same point containing some important bits which will allow us to motivate it. Note that we have already discussed a special case of the following lemma, cf. Lemma 6 in Section 1.3.3. Note that the latter is obtained by setting $r = 0$ and replacing the random variables $Z_G^{\mathbf{p}}$ with $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$.

Lemma 36. *Let $r \geq 0$ be a constant not depending on n. Let $G \sim \mathcal{G}_n^r$ and denote by X_{in}, $i = 1, 2, \ldots,$ the number of cycles of length $2i$ in G. There exist random variables $W_{mn}^{\mathbf{p}}$, a deterministic function of $X_{1n}, X_{2n}, \ldots, X_{mn}$, such that for every $\varepsilon > 0$*

$$\lim_{m\to\infty} \limsup_{n\to\infty} \Pr_{\mathcal{G}_n^r}\Big(\bigcup_{\mathbf{p}}\bigcup_{\eta}\Big[\Big|\frac{Z_G^{\mathbf{p}}(\eta)}{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]} - W_{mn}^{\mathbf{p}}\Big| > \varepsilon\Big]\Big) = 0, \tag{99}$$

There also exists a positive constant $c > 0$ such that $W_{mn}^{\mathbf{p}} > c$ uniformly in m, n. Moreover, when the phases $\mathcal{Q}$ are permutation-symmetric, the variables $W_{mn}^{\mathbf{p}}$ do not depend on the phase $\mathbf{p}$.

Remark 14. *When r is allowed to vary with n, say $r = o(n^{1/4})$, (99) must be modified into*

$$\lim_{m\to\infty} \limsup_{n\to\infty} \sup_{\mathbf{p},\eta} \Pr_{\mathcal{G}_n^r}\Big(\Big[\Big|\frac{Z_G^{\mathbf{p}}(\eta)}{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]} - W_{mn}^{\mathbf{p}}\Big| > \varepsilon\Big]\Big) = 0. \tag{100}$$

Lemma 36 provides a straightforward proof of Theorem 34, so we shall elucidate its most important aspects in an attempt to demystify its rather unintuitive statement. Equation (99) says that for all sufficiently large m, n the random variables $Z_G^{\mathbf{p}}(\eta)/\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]$ are well-approximated by the variables $W_{mn}^{\mathbf{p}}$, with large probability. To get a feeling about this statement, it is well known fact that a random Δ-regular graph is locally tree-like and its girth diverges as $n \to \infty$. That is, as n grows large, for any positive integer t, for all but $o(n)$ vertices, the t-depth neighborhood of a vertex is eventually isomorphic to the first t levels of the infinite Δ-regular tree. This is in alignment with the fact that $\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]$ is determined by the Gibbs measure on the infinite Δ-regular tree associated to the phase $\mathbf{p}$. On the other hand, a graph $G \sim \mathcal{G}_n^r$ does have $o(n)$ vertices which are contained in constant sized cycles. Thus, it is reasonable to expect that $Z_G^{\mathbf{p}}(\eta)$ fluctuates from its expectation. It is equally reasonable to expect the fluctuations to depend on the presence of small cycles which occur with small but non-zero probability.

The surprising aspect of (99), a consequence of applying the conditioning method, is that there is an explicit handle on these fluctuations, a crucial component in the proof of Theorem 34. This handle is given by the variables $W_{mn}^{\mathbf{p}}$, which are a deterministic function of the small cycle counts in G. Crucially for our proof of Theorem 34, when the phases are permutation-symmetric, the fluctuations from the expectation are captured by a single random variable, which allows us to control them uniformly over all the phases $\mathbf{p}$ and configurations η.

We should point out that the notation $W^{\mathbf{p}}_{mn}$ should not be confused by any means to the labeling of the degree $\Delta - 1$ vertices in G, i.e., the set of vertices W.

Proof of Theorem 34. We consider first the case that r is fixed, and then discuss how to modify the argument for $r = o(n^{1/4})$. Let $\varepsilon' > 0$ be sufficiently smaller than ε, to be picked later, and consider also an arbitrary $\varepsilon'' > 0$.

By Lemma 36, for all m, n sufficiently large the random variables $Z^{\mathbf{p}}_G(\eta)/\mathbf{E}_{\mathcal{G}^r_n}\big[Z^{\mathbf{p}}_G(\eta)\big]$ are well approximated by $W^{\mathbf{p}}_{mn}$ with large probability. That is, there exist $M(\varepsilon'')$, $N(\varepsilon'')$ such that for $m \geq M$ and $n \geq N$, it holds with probability $1 - \varepsilon''$ over the choice of the graph G that, for every phase $\mathbf{p}$ and every configuration $\eta : W \to [q]$,

$$Z^{\mathbf{p}}_G(\eta) = (W^{\mathbf{p}}_{mn} \pm \varepsilon')\mathbf{E}_{\mathcal{G}^r_n}\big[Z^{\mathbf{p}}_G(\eta)\big]. \tag{101}$$

We will show that whenever this is the case (for an appropriate choice of ε'), Items 1 and 2 hold. To do this, sum (101) over η to obtain that for each phase $\mathbf{p}$, it holds

$$Z^{\mathbf{p}}_G = (W^{\mathbf{p}}_{mn} \pm \varepsilon')\mathbf{E}_{\mathcal{G}^r_n}\big[Z^{\mathbf{p}}_G\big], \tag{102}$$

Using that $W^{\mathbf{p}}_{mn}$ are uniformly bounded by the positive constant c in Lemma 36, we obtain that for ε' sufficiently smaller than c, the ratio $Z^{\mathbf{p}}_G(\eta)/Z^{\mathbf{p}}_G$ is within a multiplicative $(1 \pm \varepsilon)$ from $\mathbf{E}_{\mathcal{G}^r_n}\big[Z^{\mathbf{p}}_G(\eta)\big]/\mathbf{E}_{\mathcal{G}^r_n}\big[Z^{\mathbf{p}}_G\big]$. This gives Item 2 of the theorem, when used in conjuction with (95) and (98). Note that this part of the argument did not use that the phases $\mathbf{p}$ are permutation-symmetric.

To obtain Item 1, we have to use that the phases $\mathbf{p}$ are permutation-symmetric. Then $W^{\mathbf{p}}_{mn} =: W_{mn}$ by the last assertion in Lemma 36. Thus, a summation of (102) over $\mathbf{p} \in \mathcal{Q}$ gives $Z^{\mathbf{p}}_G = (W_{mn} \pm \varepsilon')\mathbf{E}_{\mathcal{G}^r_n}\big[Z^{\mathbf{p}}_G\big]$. Exactly the same reasoning yields the thesis.

It is a standard union bound to show that Item 3 holds with probability $1 - o(1)$ over the choice of the graph G, essentially because G is an expander. Perhaps the second assertion there requires a brief proof sketch. Let $v \in U^+ \cup W^+$, $w_1, w_2 \in W^-$ and let E_i be the event that (v, w_i) is an edge of G. The events E_1, E_2 are negatively correlated since v

has a fixed number of edges incident to it, either Δ or $\Delta - 1$. It is also easy to see that $\Pr_{\mathcal{G}_n^r}(E_i) \le 1 - (1 - 1/n)^{\Delta} = O(1/n)$, so that $\Pr_{\mathcal{G}_n^r}(E_1 \cap E_2) = O(1/n^2)$. A union bound over the roughly $nr^2 = o(n^2)$ choices for the vertices v, w_1, w_2 gives the desired bound.

Thus, a graph $G \sim \mathcal{G}_n^r$ satisfies Items 1, 2 and 3 with probability $1-\varepsilon''$ for all sufficiently large n. Since ε'' was arbitrary, this gives the first part of the theorem. The second part of the theorem follows immediately by contiguity, see [35, Section 2]. $\square$

4.2.1 Application of the small subgraph conditioning method

The application of Theorem 20, and similarly Lemma 21, requires a verification of its assumptions. This check is routine for the most part, but it is nevertheless technically arduous, mainly because of condition (A3), which requires precise calculation of the moments' asymptotics. We suppress the verification in the following lemma whose proof is given later in this section. The lemma includes some details on a few quantities which will be relevant in the proof of Lemma 36.

Lemma 37. *Let $G \sim \mathcal{G}_n^r$ and X_{in} be the number of cycles of length $2i$ appearing in G, $i = 1, 2, \ldots$. Let $S = \{(\mathbf{p}, \eta) \mid \mathbf{p} \in \mathcal{Q},\ \eta : W \to [q]\}$ and for $s \in S$ with $s = (\mathbf{p}, \eta)$, set $Y_n^{(s)} = Z_G^{\mathbf{p}}(\eta)$. In the setting of Theorem 14, the assumptions of Theorem 20 hold.*

Further, for $s \in S$ with $s = (\mathbf{p}, \eta)$, for $i = 1, 2, \ldots$, $\delta_i^{(s)}$ satisfies (i) *$\delta_i^{(s)} > 0$,* (ii) *$\delta_i^{(s)}$ depends on $\mathbf{p}$ but not on η,* (iii) *$\sum_i \lambda_i \delta_i^{(s)} < \infty$,* (iv) *if the phases are permutation-symmetric, $\delta_i^{(s)}$ depends on the spin model but not on the particular phase $\mathbf{p}$.*

Further, all of the above hold if $S = \{(\mathbf{p}, \eta) \mid \mathbf{p} \in \mathcal{Q},\ \eta : W \to [q]\}$ and $Y_n^{(s)} = Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}(\eta)$.

Using Lemmas 21 and 37, we are ready to prove Lemmas 6 and 36.

Proof of Lemmas 6 and 36. To see (14) and (99), note that the $W_{mn}^{(s)}$ of Lemma 21 depend on the particular s only through $\delta_i^{(s)}$. By Item (ii) of Lemma 37, these depend only on $\mathbf{p}$ in general and specifically for the permutation-symmetric case, only on the spin model by Item (iv).

It remains to prove that $W_{mn}^{\mathbf{p}}$ are lower bounded uniformly in $\mathbf{p}$ by a positive constant. Since the number of phases $\mathbf{p}$ is bounded by a constant depending only on the spin model, it suffices to show that this is the case for a fixed phase $\mathbf{p}$. Using Item (i) of Lemma 37 and that the random variables X_{in} are non-negative integer valued, we have everywhere the bound

$$W^{\mathbf{p}}_{mn} = \prod_{i=1}^{m} \big(1+\delta^{\mathbf{p}}_i\big)^{X_{in}} \exp\big(-\lambda_i\delta^{\mathbf{p}}_i\big) \geq \prod_{i=1}^{m} \exp\big(-\lambda_i\delta^{\mathbf{p}}_i\big) > \prod_{i=1}^{\infty} \exp\big(-\lambda_i\delta^{\mathbf{p}}_i\big).$$

Note that we have identified the $\delta^{(s)}_i$'s with the respective $\delta^{\mathbf{p}}_i$'s, this is justified by Item (ii) of Lemma 37. The last quantity is finite and positive by Item (iii) in Lemma 37. □

We next prove Lemma 37 which amounts to checking the validity of the assumptions (A1)-(A4) of Theorem 20 for $Z^{\mathbf{p}}_G(\eta)$ for every phase $\mathbf{p}$ and configuration $\eta : W \to [q]$. Recall that a phase $\mathbf{p}$ corresponds to a pair of vectors $(\boldsymbol{\alpha},\boldsymbol{\beta})$ which is a global maximum of Φ, which in turn corresponds a triple $(\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{x})$, a global maximum of Υ_1, see Section 2.2.1 for details. We shall also check the validity of the assumptions for the random variables $Z^{\boldsymbol{\alpha},\boldsymbol{\beta}}_G$ when $G \sim \mathcal{G}_{n,\Delta}$. This will make the arguments crispier and easier to extend to the slightly more peculiar random variables $Z^{\mathbf{p}}_G(\eta)$.

To begin, the following lemma puts together some relevant quantities and information which have appeared in Section 3.4. For vectors $\mathbf{z}_i \in \mathbb{R}^{m_i}$, $i = 1,\dots,t$ we denote by $[\mathbf{z}_1,\dots,\mathbf{z}_t]^{\intercal}$ the $\mathbb{R}^{\sum_i m_i}$ vector which is the concatenation of the vectors $\mathbf{z}_1,\dots,\mathbf{z}_t$. For a vector $\mathbf{z} = [z_1,\dots,z_n]^{\intercal}$ and a function $f : \mathbb{R} \to \mathbb{R}$, we denote by $f(\mathbf{z})$ the vector $[f(z_1),\dots,f(z_n)]^{\intercal}$, provided that the $f(z_i)$'s are well defined.

Lemma 38. *Suppose that $(\boldsymbol{\alpha},\boldsymbol{\beta})$ is a Hessian local maximum of Ψ_1 and let $\mathbf{p}$ be the phase corresponding to the pair $(\boldsymbol{\alpha},\boldsymbol{\beta})$. Define the vector $\mathbf{x} = (x_{ij})_{i,j\in[q]}$ as in Lemma 23.*

Let $\mathbf{J}$ be the matrix $\begin{bmatrix} \mathbf{0} & \mathbf{L} \\ \mathbf{L}^{\intercal} & \mathbf{0} \end{bmatrix}$, where $\mathbf{L}$ is the $q \times q$ matrix whose ij-entry is given by $x_{ij}/\sqrt{\alpha_i}\sqrt{\beta_j}$. It holds that:

1. *The spectrum of $\mathbf{J}$ is*

$$\pm 1, \pm\lambda_1, \dots, \pm\lambda_{q-1},$$

where the λ_i's are positive and satisfy $\max_i \lambda_i < \frac{1}{\Delta-1}$.

2. *The eigenvectors of* $\mathbf{J}$ *associated to the (simple) eigenvalues* ± 1 *are given by the vectors* $[\sqrt{\boldsymbol{\alpha}}, \pm\sqrt{\boldsymbol{\beta}}]^{\intercal}$*. Moreover, any other eigenvector of* $\mathbf{J}$ *is of the form* $[\mathbf{v}_1, \mathbf{v}_2]^{\intercal}$*, where* $\sqrt{\boldsymbol{\alpha}}^{\intercal}\mathbf{v}_1 = \sqrt{\boldsymbol{\beta}}^{\intercal}\mathbf{v}_2 = 0$*.*

Relevant to Lemma 37, observe that if the phases $\mathbf{p}$ *are permutation-symmetric, then the* λ_i*'s are* common for all phases*.*

Recall that $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ is the weight of configurations in $\Sigma^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ for a random Δ-regular bipartite graph $G \sim \mathcal{G}_{n,\Delta}$. Let X_i be the number of cycles of even length i in G. We have the following lemmas.

Note that we prefer to reserve the notation λ_i for the eigenvalues of the matrix $\mathbf{J}$ of Lemma 38, so that we make the slight change of notation in Theorem 20 from λ_i to μ_i; no confusion should arise.

Lemma 39 (Lemma 7.3 in [56])**.** *Condition 1 of Theorem 20 holds for even* i *with*

$$\mu_i = \frac{r(\Delta, i)}{i} = \frac{(\Delta-1)^i + (-1)^i(\Delta-1)}{i},$$

where $r(\Delta, i)$ *is the number of ways to properly edge color a cycle of length* i *with* Δ *colors.*

The proof of Lemma 39 is given in [56] and is omitted.

Lemma 40. *In the notation and setting of Lemma 38, for every fixed* $r \geq 0$*, every* $\eta : W \to [q]$ *and even* $i \geq 2$*, it holds that*

$$\frac{\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\, X_i]}{\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]}, \frac{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)\, X_i]}{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]} \to \mu_i(1+\delta_i) \text{ as } n \to \infty, \text{ where } \delta_i = \sum_{j=1}^{q-1} \lambda_j^i. \tag{103}$$

In particular, δ_i *is positive.*

The proof of Lemma 40 is given in Section 4.2.2.

Lemma 41. *In the setting of Lemma 40, for every fixed* $r \geq 0$*, every* $\eta : W \to [q]$ *and for every finite sequence* $m_1, \ldots, m_k$ *of nonnegative integers, it holds that*

$$\frac{\mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}\, [X_2]_{m_1} \cdots [X_{2k}]_{m_k}\big]}{\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]}, \frac{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)\, [X_2]_{m_1} \cdots [X_{2k}]_{m_k}]]}{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]}$$
$$\to \prod_{i=1}^{k} \big(\mu_i(1+\delta_i)\big)^{m_i} \text{ as } n \to \infty.$$

Once we give the proof of Lemma 40, the proof of Lemma 41 is identical to [56, Proof of Lemma 7.5] and is omitted.

Lemma 42. *In the notation and setting of Lemma 38, it holds that*

$$\exp\Big(\sum_{\textit{even } i\geq 2}\mu_i\delta_i^2\Big)=\prod_{i=1}^{q-1}\prod_{j=1}^{q-1}\big(1-(\Delta-1)^2\lambda_i^2\lambda_j^2\big)^{-1/2}\prod_{i=1}^{q-1}\prod_{j=1}^{q-1}\big(1-\lambda_i^2\lambda_j^2\big)^{-(\Delta-1)/2}.$$

Moreover, $\sum_i\mu_i\delta_i<\infty$.

The proof of Lemma 42 is given in Section 4.2.2.

Finally, we find the asymptotics of the second moment over the first moment squared.

Lemma 43. *In the notation and setting of Lemma 38, for every fixed* $r\geq 0$*, every* $\eta: W\to[q]$*, it holds that*

$$\begin{aligned}\lim_{n\to\infty}\frac{\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]}{\big(\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]\big)^2}&=\lim_{n\to\infty}\frac{\mathbf{E}_{\mathcal{G}_n^r}\big[\big(Z_G^{\mathbf{p}}(\eta)\big)^2\big]}{\big(\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]\big)^2}\\&=\prod_{i=1}^{q-1}\prod_{j=1}^{q-1}\big(1-(\Delta-1)^2\lambda_i^2\lambda_j^2\big)^{-1/2}\prod_{i=1}^{q-1}\prod_{j=1}^{q-1}\big(1-\lambda_i^2\lambda_j^2\big)^{-(\Delta-1)/2}.\end{aligned}$$

The proof of Lemma 43 is quite extensive. In Section 4.3, we reduce the asymptotics to determinants of relevant Hessian matrices. These determinants are computed in Section 4.4.2, where also the proof of Lemma 43 is given.

We are now ready to give the proof of Lemma 37.

Proof of Lemma 37. Lemmas 39–43 verify assumptions (A1)–(A4) of Theorem 20. This proves the first part of the lemma. The second part, that is, Items (i) to (iv), follow from Lemmas 38, 42. □

4.2.2 Proofs

We now give the proofs of Lemmas 40 and 42 which are the most crucial components for the success of the method in our case.

Proof of Lemma 40. The proof is close to [56, Proof of Lemma 7.4], the approach is minorly different to account for the q-spin setting. We make the minor notation change from X_i to X_ℓ. We first give the proof for $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$ for $G\sim\mathcal{G}_n$.

Let $\mathcal{S} = \{S_1, \ldots, S_q\}$ and $\mathcal{T} = \{T_1, \ldots, T_q\}$ be partitions of V_1 and V_2 respectively such that $|S_i| = \alpha_i n$ and $|T_j| = \beta_j n$ for all $i, j \in [q]$. Denote by $Y_{\mathcal{S},\mathcal{T}}$ the weight of the configuration σ that $\mathcal{S}, \mathcal{T}$ induce, i.e. for a vertex $v \in V_1$, $\sigma(v) = i$ iff $v \in S_i$ and similarly for vertices in V_2.

Fix a specific pair of $\mathcal{S}, \mathcal{T}$. By symmetry,

$$\frac{\mathbf{E}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}} X_\ell]}{\mathbf{E}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]} = \frac{\mathbf{E}[Y_{\mathcal{S},\mathcal{T}} X_\ell]}{\mathbf{E}[Y_{\mathcal{S},\mathcal{T}}]}. \tag{104}$$

We now decompose X_ℓ as follows:

- ξ will denote a proper Δ-edge colored, rooted and oriented ℓ-cycle ($r(\Delta, \ell)$ possibilities), in which the vertices are colored with $\{Y_1, \ldots, Y_q, G_1, \ldots, G_q\}$ and edges are colored with $\{1, \ldots, \Delta\}$.

  A vertex colored with Y_i (resp. G_i) for some $i \in [q]$ will be loosely called yellow (resp. green) and signifies that the vertex belongs to S_i (resp. T_i). Since a yellow vertex belongs to V_1, and a green vertex belongs to V_2, a vertex coloring is consistent with the bipartiteness of the random graph if adjacent vertices of the cycle are not both yellow or green, that is, the vertex assignments which are prohibited for neighboring vertices in the cycle are (Y_i, Y_j) and (G_i, G_j), $\forall (i, j) \in [q]^2$. Note here that we do not expicitly prohibit assignments (Y_i, G_j) in the presence of a hard constraint $B_{ij} = 0$; this will be accounted otherwise. The color of the edges will prescribe which of the Δ perfect matchings an edge of a (potential) cycle will belong to.

- Given ξ, ζ denotes a position that an i-cycle can be, i.e., the exact vertices it traverses in order, such that the prescription of the vertex colors of ξ is satisfied.

- $\mathbf{1}_{\xi,\zeta}$ is the indicator function whether a cycle specified by ξ, ζ is present in the graph G.

Note that each possible cycle corresponds to exactly 2ℓ different configurations ξ (the number of ways to root and orient the cycle). For each of those ξ, the respective sets of configurations ζ are the same. Hence, we may write

$$X_\ell = \frac{1}{2\ell} \sum_{\xi} \sum_{\zeta} \mathbf{1}_{\xi,\zeta}.$$

Let $p_1 := \Pr[\mathbf{1}_{\xi,\zeta} = 1]$. It follows that

$$\mathbf{E}[Y_{\mathcal{S},\mathcal{T}} X_\ell] = \frac{1}{2\ell} \sum_{\xi} \sum_{\zeta} p_1 \cdot \mathbf{E}[Y_{\mathcal{S},\mathcal{T}} | \mathbf{1}_{\xi,\zeta} = 1].$$

In light of (104), we need to study the ratio $\mathbf{E}[Y_{\mathcal{S},\mathcal{T}} | \mathbf{1}_{\xi,\zeta} = 1]/\mathbf{E}[Y_{\mathcal{S},\mathcal{T}}]$. At this point, to simplify notation, we may assume that ξ, ζ are fixed.

We have shown in Section 2.2.1 that

$$\mathbf{E}[Y_{\mathcal{S},\mathcal{T}}] = \left(\sum_{\mathbf{x}} \left[\binom{n}{\mathbf{x}n}^{-1} \prod_i \binom{\alpha_i n}{\mathbf{x}_{i\cdot} n} \prod_j \binom{\beta_j n}{\mathbf{x}_{\cdot j} n} \right] \prod_{i,j} B_{ij}^{x_{ij} n} \right)^{\Delta}, \tag{105}$$

where the variables $\mathbf{x} = (x_{11}, \ldots, x_{qq})$ denote the number of edges between $\mathcal{S}, \mathcal{T}$ in one matching. In particular nx_{ij} is the number of edges between the sets S_i and T_j.

To calculate $\mathbf{E}[Y_{\mathcal{S},\mathcal{T}} \mid \mathbf{1}_{\xi,\zeta} = 1]$, we need to introduce some notation. For colors $c_1, c_2 \in \{Y_1, \ldots, Y_q, G_1, \ldots, G_q\}$, we say that an edge is of type $\{c_1, c_2\}$ if its endpoints have colors c_1, c_2. Let y_i, g_j denote the number of vertices colored with Y_i, G_j respectively. For $k = 1, \ldots, \Delta$, let $a_{ij}(k)$ denote the number of edges of color k and type $\{Y_i, G_j\}$, and

$$a_{i\perp}(k) := \sum_j a_{ij}(k),\ a_{\perp j}(k) := \sum_j a_{ij}(k),\ a(k) := \sum_{i,j} a_{ij}(k),\ \mathbf{a}(k) := (a_{11}, \ldots, a_{qq}).$$

Finally, for $i, j \in [q]$ let $a_{ij} := \sum_k a_{ij}(k)$. By considering the sum of the degrees of vertices colored Y_i, the sum of the degrees of vertices colored G_j and the total number of edges of the cycle, we obtain the following equalities.

$$\textstyle\sum_j a_{ij} = 2y_i,\ \sum_i a_{ij} = 2g_j,\ \sum_{i,j} a_{ij} = 2\ell. \tag{106}$$

We are almost set to compute $\mathbf{E}[Y_{\mathcal{S},\mathcal{T}} | \mathbf{1}_{\xi,\zeta} = 1]$. For the k-th matching, denote by $\mathbf{x}$ the the total number of edges between sets S_i and T_j in the k-th matching. This number includes the $a_{ij}(k)$ edges prescribed by ξ, ζ. Set $E = \mathbf{E}[Y_{\mathcal{S},\mathcal{T}} | \mathbf{1}_{\xi,\zeta} = 1]$. We have

$$E = \prod_{k=1}^{\Delta} \left(\sum_{\mathbf{x}} \left[\binom{n - a(k)}{\mathbf{x}n - \mathbf{a}(k)}^{-1} \prod_i \binom{\alpha_i n - a_{i\perp}(k)}{\mathbf{x}_{i\cdot} n - \mathbf{a}_{i\cdot}(k)} \prod_j \binom{\beta_j n - a_{\perp j}(k)}{\mathbf{x}_{\cdot j} n - \mathbf{a}_{\cdot j}} \right] \prod_{i,j} B_{ij}^{n x_{ij}} \right),$$

where we remind the reader the notation $\mathbf{a}_{i\cdot} = (a_{i1}, \ldots, a_{iq})$ and $\mathbf{a}_{\cdot j} = (a_{1j}, \ldots, a_{qj})$.

Standard approximations of binomial coefficients, see for example [27, Lemma 27], give

$$\frac{\binom{\alpha_i n - a_{i\perp}(k)}{\mathbf{x}_{i\cdot} n - \mathbf{a}_{i\cdot}(k)}}{\binom{\alpha_i n}{\mathbf{x}_{i\cdot} n}} \sim \frac{\prod_j (x_{ij})^{a_{ij}(k)}}{\alpha_i^{a_{i\perp}(k)}},\ \frac{\binom{\beta_j n - a_{\perp j}(k)}{\mathbf{x}_{\cdot j} n - \mathbf{a}_{\cdot j} n}}{\binom{\beta_j n}{\mathbf{x}_{\cdot j}}} \sim \frac{\prod_i (x_{ij})^{a_{ij}(k)}}{\beta_j^{a_{\perp j}(k)}},\ \frac{\binom{n - a(k)}{\mathbf{x}n - \mathbf{a}(k)}}{\binom{n}{\mathbf{x}n}} \sim \prod_{i,j} (x_{ij})^{a_{ij}(k)}.$$

Thus, we obtain

$$\frac{\mathbf{E}[Y_{\mathcal{S},\mathcal{T}}|\mathbf{1}_{\xi,\zeta}=1]}{\mathbf{E}[Y_{\mathcal{S},\mathcal{T}}]} \sim \frac{\prod_{i,j}\left(x_{ij}\right)^{a_{ij}}}{\prod_i \alpha_i^{\sum_j a_{ij}} \prod_j \beta_j^{\sum_i a_{ij}}}.$$

Clearly $p_1 \sim n^{-\ell}$ and for given ξ, the number of possible ζ is asymptotic to $n^\ell \prod_i \alpha_i^{y_i} \prod_j \beta_j^{g_j}$. Thus, for the given ξ, we have

$$\frac{\sum_\zeta p_1 \mathbf{E}[Y_{\mathcal{S},\mathcal{T}} \,|\, \mathbf{1}_{\xi,\zeta}=1]}{\mathbf{E}[Y_{\mathcal{S},\mathcal{T}}]} \sim \frac{\prod_i \alpha_i^{y_i} \prod_j \beta_j^{g_j} \prod_{i,j}\left(x_{ij}\right)^{a_{ij}}}{\prod_i \alpha_i^{\sum_j a_{ij}} \prod_j \beta_j^{\sum_i a_{ij}}} = \prod_{i,j}\Big(\frac{x_{ij}}{\sqrt{\alpha_i\beta_j}}\Big)^{a_{ij}}$$

Note that the rhs evaluates to 0 whenever there exist i, j such that $B_{ij} = 0$ but $a_{ij} \neq 0$, since then we have $x_{ij} = 0$. This is in complete accordance with the fact that the configuration induced by the partition $\{\mathcal{S}, \mathcal{T}\}$ has zero weight. Thus, by (104), we may write

$$\frac{\mathbf{E}[YX_\ell]}{\mathbf{E}[Y]} = \frac{r(\Delta,\ell)}{2\ell} \cdot \sum_\xi N_{\mathbf{a}} \Big(\frac{x_{ij}}{\sqrt{\alpha_i\beta_j}}\Big)^{a_{ij}},$$

where $\mathbf{a} = \{a_{11}, \ldots, a_{qq}\}$ and $N_{\mathbf{a}}$ is the number of possible ξ with a_{ij} edges having assignment (Y_i, G_j). To analyze this sum, we employ a technique given in [35]. The idea is to define a weighted transition matrix and view it as a weighted graph. The powers of the matrix count the (multiplicative) weight of walks in the graph and a closed walk in this graph will correspond to a specification ξ. By defining the weights appropriately, one can also ensure that each closed walk will correctly capture the weight of the specification ξ.

In our setting, the transition matrix is simply the matrix $\mathbf{J}$ of Lemma 38. The first q rows and q columns correspond to the colors Y_i and the remaining rows and columns to colors G_j. The total weight of closed walks of length ℓ is given by $\mathrm{Tr}(\mathbf{J}^\ell)$. Using the description of the eigenvalues given in Item 1 of Lemma 38, we obtain that for even ℓ, $\mathrm{Tr}(\mathbf{J}^\ell) = 2\Big(1 + \sum_{i=1}^{q-1} \lambda_i^\ell\Big)$. This concludes the proof for the variables $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}$.

We next turn to the variables $Z_G^{\mathbf{p}}(\eta)$ when $G \sim \mathcal{G}_n^r$. For a fixed $\eta : W \to [q]$, our goal is to compute $\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)\, X_\ell]/\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]$. Denote by $\boldsymbol{\eta}^\pm$ the q-dimensional vectors whose i-th entries are given by $|\sigma^{-1} \cap W^\pm|$ (thus $\boldsymbol{\eta}^+$ is the analog of $\boldsymbol{\alpha}n$ and $\boldsymbol{\eta}^-$ is the analog of $\boldsymbol{\beta}n$).

For $\mathbf{x} \in \mathcal{M}_1(\boldsymbol{\alpha} + \boldsymbol{\eta}^+/n, \boldsymbol{\beta} + \boldsymbol{\eta}^-/n)$, let

$$\kappa_G^{\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{x}}(\eta) := \binom{n+r}{\mathbf{x}n}^{-1} \prod_i \binom{\alpha_i n + \eta_i^+}{\mathbf{x}_{i\cdot}} \prod_j \binom{\beta_j n + \eta_j^-}{\mathbf{x}_{\cdot j}} \prod_{i,j} (B_{ij})^{n x_{ij}},$$

and for $\mathbf{x} \in \mathcal{M}_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$, let

$$\kappa_G^{\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{x}} := \binom{n}{\mathbf{x}n}^{-1} \prod_i \binom{\alpha_i n}{\mathbf{x}_{i\boldsymbol{\cdot}}} \prod_j \binom{\beta_j n}{\mathbf{x}_{\boldsymbol{\cdot}j}} \prod_{i,j} (B_{ij})^{nx_{ij}},$$

We have that

$$\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)] = \sum_{\boldsymbol{\alpha},\boldsymbol{\beta}} \binom{n}{\boldsymbol{\alpha}n}\binom{n}{\boldsymbol{\beta}n} \left(\sum_{\mathbf{x}} \kappa_G^{\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{x}}(\eta)\right)^{\Delta-1} \left(\sum_{\mathbf{x}} \kappa_G^{\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{x}}\right), \tag{107}$$

where the sum over $\boldsymbol{\alpha}, \boldsymbol{\beta}$ ranges in $\Sigma^{\mathbf{p}}$ and the sum over $\mathbf{x}$'s range in $\mathcal{M}_1(\boldsymbol{\alpha}+\boldsymbol{\eta}^+/n, \boldsymbol{\beta}+\boldsymbol{\eta}^-/n)$ and $\mathcal{M}_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$ respectively.

To compute $\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)\, X_\ell]$, we again decompose X_ℓ as before, so that

$$\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)\, X_\ell] = \frac{1}{2\ell} \sum_{\xi} \sum_{\zeta} p_1 \mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta) \,|\, \mathbf{1}_{\xi,\zeta}].$$

It will suffice thus to compute $\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta) \,|\, \mathbf{1}_{\xi,\zeta}]/\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]$. Keeping the same definitions for $\mathbf{a}(k), \mathbf{a}_{i\boldsymbol{\cdot}}, \mathbf{a}_{\boldsymbol{\cdot}j}, a_{i\perp}(k), a_{\perp j}(k), a(k)$, for $k = 1, \ldots, \Delta - 1$, let

$$\hat{\kappa}_{G,k}^{\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{x}}(\eta) = \binom{n + r - a(k)}{\mathbf{x}n - \mathbf{a}(k)}^{-1} \prod_i \binom{\alpha_i n + \eta_i^+ - a_{i\perp}(k)}{\mathbf{x}_{i\boldsymbol{\cdot}} n - \mathbf{a}_{i\boldsymbol{\cdot}}(k)} \prod_j \binom{\beta_j n + \eta_j^- - a_{\perp j}(k)}{\mathbf{x}_{\boldsymbol{\cdot}j} n - \mathbf{a}_{\boldsymbol{\cdot}i}(k)} \prod_{i,j} B_{ij}^{nx_{ij}},$$

while for $k = \Delta$, let

$$\hat{\kappa}_{G,\Delta}^{\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{x}} = \binom{n - a(k)}{\mathbf{x}n - \mathbf{a}(k)}^{-1} \prod_i \binom{\alpha_i n - a_{i\perp}(k)}{\mathbf{x}_{i\boldsymbol{\cdot}} n - \mathbf{a}_{i\boldsymbol{\cdot}}(k)} \prod_j \binom{\beta_j n - a_{\perp j}(k)}{\mathbf{x}_{\boldsymbol{\cdot}j} n - \mathbf{a}_{\boldsymbol{\cdot}i}(k)} \prod_{i,j} B_{ij}^{nx_{ij}}.$$

We have that

$$\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)] = \sum_{\boldsymbol{\alpha},\boldsymbol{\beta}} \binom{n}{\boldsymbol{\alpha}n}\binom{n}{\boldsymbol{\beta}n} \left[\prod_{k=1}^{\Delta-1} \left(\sum_{\mathbf{x}} \hat{\kappa}_{G,k}^{\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{x}}(\eta)\right)\right] \left(\sum_{\mathbf{x}} \hat{\kappa}_{G,\Delta}^{\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{x}}\right). \tag{108}$$

The previous calculations show that

$$\frac{\hat{\kappa}_{G,\Delta}^{\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{x}}}{\kappa_G^{\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{x}}} \sim \frac{\prod_{i,j} (x_{ij})^{a_{ij}(\Delta)}}{\prod_i \alpha_i^{\sum_j a_{ij}(\Delta)} \prod_j \beta_j^{\sum_i a_{ij}(\Delta)}}$$

Moreover, we have

$$\frac{\binom{\alpha_i n + \eta_i^+ - a_{i\perp}(k)}{\mathbf{x}_{i\boldsymbol{\cdot}} n - \mathbf{a}_{i\boldsymbol{\cdot}}(k)}}{\binom{\alpha_i n + \eta_i^+}{\mathbf{x}_{i\boldsymbol{\cdot}} n}} \sim \frac{\prod_j x_{ij}^{a_{ij}(k)}}{\alpha_i^{a_{i\perp}(k)}}, \quad \frac{\binom{\beta_j n + \eta_j^- - a_{\perp j}(k)}{\mathbf{x}_{\boldsymbol{\cdot}j} n - \mathbf{a}_{\boldsymbol{\cdot}j}(k)}}{\binom{\beta_j n + \eta_j^-}{\mathbf{x}_{\boldsymbol{\cdot}j} n}} \sim \frac{\prod_i x_{ij}^{a_{ij}(k)}}{\beta_j^{a_{\perp j}(k)}}, \quad \frac{\binom{n + r - a(k)}{\mathbf{x}n - \mathbf{a}(k)}}{\binom{n+r}{\mathbf{x}n}} \sim \prod_{i,j} x_{ij}^{a(k)}.$$

yielding

$$\frac{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta) \,|\, \mathbf{1}_{\xi,\zeta}]}{\mathbf{E}_{\mathcal{G}_n^r}[Z_G^{\mathbf{p}}(\eta)]} \sim \frac{\prod_{i,j} (x_{ij})^{a_{ij}}}{\prod_i \alpha_i^{\sum_j a_{ij}} \prod_j \beta_j^{\sum_i a_{ij}}},$$

exactly as before. Hence, the rest of the proof for the variables $Z_G^{\mathbf{p}}(\eta)$ is exactly the same as for the variables $Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}(\eta)$. $\square$

Proof of Lemma 42. Using Lemma 39, we have

$$\sum_{\text{even } i\geq 2} \mu_i \delta_i^2 = \sum_{\text{even } i\geq 2} \frac{r(\Delta, i)}{i} \cdot \left(\sum_{j=1}^{q-1} \lambda_j^i\right)^2 = \sum_{\text{even } i\geq 2} \frac{(\Delta-1)^i + (\Delta-1)}{i} \cdot \left(\sum_{j=1}^{q-1}\sum_{j'=1}^{q-1} \lambda_j^i \lambda_{j'}^i\right).$$

Observe that $\sum_{j\geq 1} \frac{x^{2j}}{2j} = -\frac{1}{2}\ln(1-x^2)$ for all $|x| < 1$. By Item 1 of Lemma 38, it holds that $(\Delta-1)\lambda_j < 1$ for all j, so that $(\Delta-1)\lambda_j\lambda_{j'} < 1$ for all j, j'. It follows that

$$\sum_{\text{even } i\geq 2} \mu_i \delta_i^2 = -\frac{1}{2}\Big(\sum_{i,j} \ln\big(1-(\Delta-1)^2\lambda_i^2\lambda_j^2\big) + (\Delta-1)\sum_{i,j}\ln\big(1-\lambda_i^2\lambda_j^2\big)\Big),$$

thus proving the first part of the lemma. The proof of $\sum_i \mu_i\delta_i < \infty$ is completely analogous.

□

4.3 Moment Asymptotics

In this section we compute the asymptotics of the moments. For a dominant phase $\mathbf{p} = (\boldsymbol{\alpha}, \boldsymbol{\beta})$ (cf. Definition 6), we will be interested in the asymptotics of

$$\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}],\ \mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2],$$

the asymptotics of $\mathbf{E}_{\mathcal{G}}[Z_G^{\mathbf{p}}],\ \mathbf{E}_{\mathcal{G}}[(Z_G^{\mathbf{p}})^2]$ can be treated completely analogously.

The main idea behind the calculation of the moment asymptotics is that the sums in (23) and (25) are determined by the terms with the largest contribution. For example, to compute the asymptotics of $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]$, we need to consider the vector $\mathbf{x}^*$ which maximizes $\Upsilon_1(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x})$ and integrate over the vectors $\mathbf{x}$ which are close to $\mathbf{x}^*$. Provided that the Hessian of $\Upsilon_1(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x})$ at $\mathbf{x}^*$ is negative definite, we obtain a Gaussian integral which can be computed in terms of the determinant of the Hessian. A similar approach applies to compute the asymptotics of $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]$; there, we need to consider $(\boldsymbol{\gamma}^*, \boldsymbol{\delta}^*, \mathbf{y}^*)$ which maximize $\Upsilon_2(\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y})$ and integrate over $(\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y})$ which are close to $(\boldsymbol{\gamma}^*, \boldsymbol{\delta}^*, \mathbf{y}^*)$. Note that the $\mathbf{x}^*$ and $(\boldsymbol{\gamma}^*, \boldsymbol{\delta}^*, \mathbf{y}^*)$ are unique; see Remark 9 for details. To lighten notation, we will drop the stars from the maximizers when clear from context.

To perform the integrations accurately, we will need to ensure that the integration variables lie in a full dimensional space. This requires a slightly different view of the domain of the functions Υ_1 and Υ_2, which we refer to as the full dimensional representation. We will also have to take care of a minor technical detail relative to the presence of hard constraints. Such considerations were circumvented in the maximization of Υ_1 and Υ_2, but here it will be cleaner to explicitly deal with them in a slightly more combinatorial fashion. We do this in the context of the calculations needed for the asymptotics of $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})]$ and $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]$, respectively.

4.3.1 Full-dimensional representations

For the first moment, $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})]$ is a sum over $\mathbf{x}$ while $\boldsymbol{\alpha}, \boldsymbol{\beta}$ are fixed. Let

$$P_1 = \big\{(i,j) \in [q]^2 \,\big|\, B_{ij} > 0\big\}. \tag{109}$$

In the presence of a hard constraint $B_{ij} = 0$, edge assignments (i,j) have zero weight and hence correspond to a non-permissible configuration. In the maximization of Υ_1, the hard constraint $B_{ij} = 0$ was not directly relevant, since for $x_{ij} > 0$ the function Υ_1 evaluates to $-\infty$. Indeed, we found that the optimal x_{ij} is of the form $B_{ij}R_iC_j$ and hence zero. However, the asymptotics in which we are interested include products of the optimal values of the x_{ij} and to correctly capture them, we need to explicitly rule out the zero values.

To do so, in the formulation (22), we hard-code $x_{ij} = 0$ for a pair $(i,j) \notin P_1$ and hence the variables $\mathbf{x}$ are restricted to the space

$$\begin{aligned} \textstyle\sum_j x_{ij} &= \alpha_i \quad \big(\forall i \in [q]\big), & \textstyle\sum_i x_{ij} &= \beta_j \quad \big(\forall j \in [q]\big), \\ x_{ij} &= 0 \quad \big(\forall (i,j) \in [q]^2 \backslash P_1\big), & x_{ij} &\geq 0 \quad \big(\forall (i,j) \in P_1\big). \end{aligned} \tag{110}$$

Note that the dimension of the polytope (110) is $|P_1| - (2q-1)$ since the matrix $\mathbf{B}$ is irreducible. To get affinely independent variables $\mathbf{x}$, we use the equalities in (110) and substitute an appropriate set of $(q-1)^2 - |P_1|$ variables. We will not need to understand these substitutions till Section 4.4.1, yet in the integrations which follow it is preferable to have integration variables rather than integrate over subspaces.

After this process, we are going to have $|P_1| - (2q-1)$ variables lying in a full dimensional space. We refer to this set of variables as the full dimensional $\mathbf{x}$. We still use the notation $\mathbf{x}$ for these variables and refer to x_{ij} even if x_{ij} is not in the representation $\mathbf{x}$, under the understanding that this is just a shorthand for the substituted expression. Using these conventions, note that $\Upsilon_1(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x})$ is now a function of the full dimensional $\mathbf{x}$, and we will refer to this setup as the full dimensional representation of Υ_1.

For the second moment, $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]$ is a sum over $\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y}$ while $\boldsymbol{\alpha}, \boldsymbol{\beta}$ are fixed. Let

$$P_2 = \big\{(i,k,j,l) \in [q]^4 \,\big|\, B_{ikjl} > 0\big\}. \tag{111}$$

Since $B_{ikjl} = B_{ij}B_{kl}$, we clearly have $|P_2| = |P_1|^2$. As above, we now restrict $\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y}$ to lie in the space

$$\begin{gathered}
\textstyle\sum_k \gamma_{ik} = \alpha_i \quad \big(\forall i \in [q]\big), \quad \sum_l \delta_{jl} = \beta_j \quad \big(\forall j \in [q]\big), \quad \sum_{j,l} y_{ikjl} = \gamma_{ik} \quad \big(\forall (i,k) \in [q]^2\big) \\
\textstyle\sum_i \gamma_{ik} = \alpha_k \quad \big(\forall k \in [q]\big), \quad \sum_j \delta_{jl} = \beta_l \quad \big(\forall l \in [q]\big), \quad \sum_{i,k} y_{ikjl} = \delta_{jl} \quad \big(\forall (j,l) \in [q]^2\big) \\
\gamma_{ik} \geq 0 \quad \big(\forall (i,k) \in [q]^2\big), \quad \delta_{jl} \geq 0 \quad \big(\forall (j,l) \in [q]^2\big), \quad y_{ikjl} \geq 0 \quad \big(\forall (i,k,j,l) \in P_2\big), \\
y_{ikjl} = 0 \quad \big(\forall (i,k,j,l) \in [q]^4 \backslash P_2\big).
\end{gathered} \tag{112}$$

Once again, note the extra equality constraints relative to the original formulation (24). The same remarks and conventions apply as in the case of the first moment. Therefore, we obtain full dimensional $(\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y})$ and a full dimensional representation of $\Upsilon_2(\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y})$. We point out that in the full dimensional $(\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y})$, there are $(q-1)^2$ variables γ_{ik}, $(q-1)^2$ variables δ_{jl} and $|P_2| - (2q^2 - 1)$ variables y_{ikjl}.

The following two lemmas express the asymptotics of the moments in terms of suitable determinants. Thus, to compute the asymptotics of the ratio $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2] / \big(\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]\big)^2$, it suffices to compute the determinants appearing in the above two lemmas. This is a long and twisted road and is deferred to Section 4.4.2, where also the proof of Lemma 43 is completed.

Lemma 44. *Suppose that $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ maximize Ψ_1. Let $\mathbf{x}$ be the (unique) maximizer of Υ_1. Denote by $\mathbf{H}_1$ be the Hessian of the full dimensional representation of $\Upsilon_1(\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x})$ scaled by $1/\Delta$ and by $\mathbf{H}_{1,\mathbf{x}}$ the square submatrix of $\mathbf{H}_1$ corresponding to rows and columns indexed by $\mathbf{x}$. Then*

$$\lim_{n\to\infty} \frac{(2\pi n)^{q-1} \mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]}{e^{n\Upsilon_1(\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{x})}} = \frac{\Big(\prod_i \alpha_i \prod_j \beta_j\Big)^{(\Delta-1)/2} \Big(\prod_{(i,j)\in P_1} x_{ij}\Big)^{-\Delta/2}}{\big(\mathrm{Det}(-\mathbf{H}_{1,\mathbf{x}})\big)^{\Delta/2}}.$$

Lemma 45. *Suppose that* $(\boldsymbol{\alpha},\boldsymbol{\beta})$ *maximize* Ψ_1. *Let* $(\boldsymbol{\gamma},\boldsymbol{\beta},\mathbf{y})$ *be the (unique) maximizer of* Υ_2. *Denote by* $\mathbf{H}_2$ *be the Hessian of the full dimensional representation of* $\Upsilon_2(\boldsymbol{\gamma},\boldsymbol{\delta},\mathbf{y})$ *scaled by* $1/\Delta$ *and* $\mathbf{H}_{2,\mathbf{y}}$ *the square submatrix of* $\mathbf{H}_2$ *corresponding to rows and columns indexed by* $\mathbf{y}$. *Then*

$$\lim_{n\to\infty}\frac{(2\pi n)^{2(q-1)}\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]}{e^{n\Upsilon_2(\boldsymbol{\gamma},\boldsymbol{\delta},\mathbf{y})}}=\frac{\Big(\prod_{i,k}\gamma_{ik}\prod_{j,l}\delta_{jl}\Big)^{(\Delta-1)/2}\Big(\prod_{(i,k,j,l)\in P_2}y_{ikjl}\Big)^{-\Delta/2}}{\Delta^{(q-1)^2}\big(\mathrm{Det}(-\mathbf{H}_2)\big)^{1/2}\big(\mathrm{Det}(-\mathbf{H}_{2,\mathbf{y}})\big)^{(\Delta-1)/2}}.$$

The proofs of Lemmas 44 and 45 are similar; we give the proof of Lemma 45 which is a bit tricky due to a two step integration.

Proof of Lemma 45. We assume a full dimensional representation of $(\boldsymbol{\gamma},\boldsymbol{\delta},\mathbf{y})$. We denote by $(\boldsymbol{\gamma}^*,\boldsymbol{\delta}^*,\mathbf{y}^*)$ the optimal vector which maximizes the full dimensional representation of $\Upsilon_2(\boldsymbol{\gamma},\boldsymbol{\delta},\mathbf{y})$. We have that $\gamma^*_{ik},\delta^*_{jl}>0$ for all i,k,j,l and $y^*_{ikjl}>0$ for $(i,k,j,l)\in P_2$. For all δ sufficiently small it holds that:

$$\|(\boldsymbol{\gamma},\boldsymbol{\delta},\mathbf{y})-(\boldsymbol{\gamma}^*,\boldsymbol{\delta}^*,\mathbf{y}^*)\|_2\le\delta\text{ implies }\begin{cases}\gamma_{ik},\delta_{jl}>0\text{ for all }i,k,j,l\in[q],\\ y_{ikjl}>0\text{ for }(i,k,j,l)\in P_2.\end{cases}$$

Since Υ_2 has a unique global maximum $(\boldsymbol{\gamma}^*,\boldsymbol{\delta}^*,\mathbf{y}^*)$ in the space (112), standard compactness arguments imply that there exists $\varepsilon(\delta)>0$ such that $\|(\boldsymbol{\gamma},\boldsymbol{\delta},\mathbf{y})-(\boldsymbol{\gamma}^*,\boldsymbol{\delta}^*,\mathbf{y}^*)\|\ge\delta$ implies $\Upsilon_2(\boldsymbol{\gamma}^*,\boldsymbol{\delta}^*,\mathbf{y}^*)-\Upsilon_2(\boldsymbol{\gamma},\boldsymbol{\delta},\mathbf{y})\ge\varepsilon$. It follows that the contribution of terms with $\|(\boldsymbol{\gamma},\boldsymbol{\delta},\mathbf{y})-(\boldsymbol{\gamma}^*,\boldsymbol{\delta}^*,\mathbf{y}^*)\|\ge\delta$ to $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]$ may be safely ignored. Hence we may restrict our attention to $\boldsymbol{\gamma},\delta,\mathbf{y}$ satisfying $\|(\boldsymbol{\gamma},\boldsymbol{\delta},\mathbf{y})-(\boldsymbol{\gamma}^*,\boldsymbol{\delta}^*,\mathbf{y}^*)\|<\delta$. Moreover, using Taylor's expansion, we may choose δ small enough such that Υ_2 decays quadratically in a δ-ball around $(\boldsymbol{\gamma}^*,\boldsymbol{\delta}^*,\mathbf{y}^*)$.

Utilizing the choice of δ and Stirling's approximation for factorials, we thus obtain

$$\frac{\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]}{e^{n\Upsilon_2(\boldsymbol{\gamma}^*,\boldsymbol{\delta}^*,\mathbf{y}^*)}}=\big(1+O(n^{-1})\big)\frac{1}{(2\pi n)^{2(q-1)}}\sum_{\boldsymbol{\gamma},\boldsymbol{\delta}}\Big(\frac{1}{\sqrt{2\pi n}}\Big)^{2(q-1)^2}\Big(\prod_{i,k}\gamma_{ik}\prod_{j,l}\delta_{jl}\Big)^{(\Delta-1)/2}$$
$$\Bigg[\sum_{\mathbf{y}}\Big(\frac{1}{\sqrt{2\pi n}}\Big)^{|P_2|-(2q^2-1)}\Big(\prod_{(i,k,j,l)\in P_2}\frac{1}{\sqrt{y_{ikjl}}}\Big)e^{n\big(\Upsilon_2(\boldsymbol{\gamma},\boldsymbol{\delta},\mathbf{y})-\Upsilon_2(\boldsymbol{\gamma}^*,\boldsymbol{\delta}^*,\mathbf{y}^*)\big)/\Delta}\Bigg]^{\Delta}.$$

We now compute

$$L=\lim_{n\to\infty}\frac{(2\pi n)^{2(q-1)}\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}})^2]}{e^{n\Upsilon_2(\boldsymbol{\gamma}^*,\boldsymbol{\delta}^*,\mathbf{y}^*)}}.$$

Standard techniques of rewriting sums as integrals and an application of the dominated convergence theorem (see for example [36, Section 9.4]) ultimately give

$$L = \frac{1}{(2\pi)^{2(q-1)}} \Big(\prod_{i,k} \gamma_{ik}^* \prod_{j,l} \delta_{jl}^* \Big)^{(\Delta-1)/2} \Big(\prod_{(i,k,j,l)\in P_2} y_{ikjl}^* \Big)^{-\Delta/2} \tag{113}$$
$$\Big(\frac{1}{\sqrt{2\pi}}\Big)^{2(q-1)^2} \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} \bigg[\Big(\frac{1}{\sqrt{2\pi}}\Big)^{|P_2|-(2q^2-1)} \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} e^{\frac{1}{2}(\gamma,\delta,\mathbf{y})\cdot \mathbf{H} \cdot (\gamma,\delta,\mathbf{y})^{\intercal}} d\mathbf{y} \bigg]^{\Delta} d\boldsymbol{\gamma} d\boldsymbol{\delta},$$

where $\mathbf{H}$ denotes the Hessian matrix of Υ_2 evaluated at $(\boldsymbol{\gamma}^*, \boldsymbol{\delta}^*, \mathbf{y}^*)$ scaled by $1/\Delta$ and the operator $\cdot$ stands for matrix multiplication.

We thus focus on computing the integral in (113). We begin with the inner integration. Let

$$I_1 = \bigg[\Big(\frac{1}{\sqrt{2\pi}}\Big)^{|P_2|-(2q^2-1)} \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} e^{\frac{1}{2}(\gamma,\delta,\mathbf{y})\cdot \mathbf{H} \cdot (\gamma,\delta,\mathbf{y})^{\intercal}} d\mathbf{y} \bigg]^{\Delta}.$$

To calculate I_1, we first decompose the exponent to isolate the terms involving $\mathbf{y}$. We obtain

$$\frac{1}{2}(\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y}) \cdot \mathbf{H} \cdot (\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y})^{\intercal} = \frac{1}{2}(\boldsymbol{\gamma}, \boldsymbol{\delta}) \cdot \mathbf{H}_{\boldsymbol{\gamma},\boldsymbol{\delta}} \cdot (\boldsymbol{\gamma}, \boldsymbol{\delta})^{\intercal} - \frac{1}{2}\mathbf{y} \cdot (-\mathbf{H}_{\mathbf{y}}) \cdot \mathbf{y}^{\intercal} + \mathbf{T} \cdot \mathbf{y}^{\intercal},$$

where $\mathbf{H} = \begin{bmatrix} \mathbf{H}_{\gamma,\delta} & \mathbf{H}_{\boldsymbol{\gamma}\boldsymbol{\delta},\mathbf{y}} \\ \mathbf{H}_{\boldsymbol{\gamma}\boldsymbol{\delta},\mathbf{y}}^{\intercal} & \mathbf{H}_{\mathbf{y}} \end{bmatrix}$ and $\mathbf{T} = (\boldsymbol{\gamma}, \boldsymbol{\delta}) \cdot \mathbf{H}_{\boldsymbol{\gamma}\boldsymbol{\delta},\mathbf{y}}$. Specifically:

- $\mathbf{H}_{\boldsymbol{\gamma},\boldsymbol{\delta}}$ is the square submatrix of $\mathbf{H}$ corresponding to the rows indexed by $\boldsymbol{\gamma}, \boldsymbol{\delta}$ and the columns indexed by $\boldsymbol{\gamma}, \boldsymbol{\delta}$,

- $\mathbf{H}_{\mathbf{y}}$ is the square submatrix of $\mathbf{H}$ corresponding to the rows indexed by $\mathbf{y}$ and the columns indexed by $\mathbf{y}$,

- $\mathbf{T} = (\boldsymbol{\gamma}, \boldsymbol{\delta}) \cdot \mathbf{H}_{\boldsymbol{\gamma}\boldsymbol{\delta},\mathbf{y}}$, where $\mathbf{H}_{\boldsymbol{\gamma}\boldsymbol{\delta},\mathbf{y}}$ is the submatrix of $\mathbf{H}$ corresponding to the rows indexed by $\boldsymbol{\gamma}, \boldsymbol{\delta}$ and the columns indexed by $\mathbf{y}$.

Note that $\mathbf{H}_{\mathbf{y}}$ is the Hessian of $g_2(\mathbf{y})$ evaluated at $\mathbf{y}^*$. Since $g_2(\mathbf{y})$ is concave, we have that $\mathbf{H}_{\mathbf{y}}$ is negative definite. Utilizing this decomposition, we obtain

$$\begin{aligned} I_1 &= e^{\frac{\Delta}{2}(\boldsymbol{\gamma},\boldsymbol{\delta})\cdot\mathbf{H}_{\gamma,\delta}\cdot(\boldsymbol{\gamma},\boldsymbol{\delta})^{\intercal}} \bigg[\Big(\frac{1}{\sqrt{2\pi}}\Big)^{|P_2|-(2q^2-1)} \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} e^{-\frac{1}{2}\mathbf{y}\cdot(-\mathbf{H}_{\mathbf{y}})\cdot\mathbf{y}^{\intercal} + \mathbf{T}\cdot\mathbf{y}^{\intercal}} d\mathbf{y} \bigg]^{\Delta} \\ &= \frac{1}{\big(\mathrm{Det}(-\mathbf{H}_{\mathbf{y}})\big)^{\Delta/2}} e^{\frac{\Delta}{2}\left(\mathbf{T}\cdot(-\mathbf{H}_{\mathbf{y}})^{-1}\cdot\mathbf{T}^{\intercal} + (\boldsymbol{\gamma},\boldsymbol{\delta})\cdot\mathbf{H}_{\gamma,\delta}\cdot(\boldsymbol{\gamma},\boldsymbol{\delta})^{\intercal}\right)}. \end{aligned}$$

We are left with the task of computing the integral

$$I_2 = \Big(\frac{1}{\sqrt{2\pi}}\Big)^{2(q-1)^2} \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} e^{\frac{\Delta}{2}\big(\mathbf{T}\cdot(-\mathbf{H_y})^{-1}\cdot\mathbf{T}^{\intercal}+(\boldsymbol{\gamma},\boldsymbol{\delta})\cdot\mathbf{H}_{\boldsymbol{\gamma},\boldsymbol{\delta}}\cdot(\boldsymbol{\gamma},\boldsymbol{\delta})^{\intercal}\big)} d\boldsymbol{\gamma} d\boldsymbol{\delta}. \tag{114}$$

Using the definition of $\mathbf{T}$, we have

$$\mathbf{T}\cdot(-\mathbf{H_y})^{-1}\cdot\mathbf{T}^{\intercal}+(\boldsymbol{\gamma},\boldsymbol{\delta})\cdot\mathbf{H}_{\boldsymbol{\gamma},\boldsymbol{\delta}}\cdot(\boldsymbol{\gamma},\boldsymbol{\delta})^{\intercal} = (\boldsymbol{\gamma},\boldsymbol{\delta})\cdot\big(\mathbf{H}_{\boldsymbol{\gamma},\boldsymbol{\delta}} - \mathbf{H}_{\boldsymbol{\gamma}\boldsymbol{\delta},\mathbf{y}}\cdot\mathbf{H}_{\mathbf{y}}^{-1}\cdot\mathbf{H}_{\boldsymbol{\gamma}\boldsymbol{\delta},\mathbf{y}}^{\intercal}\big)\cdot(\boldsymbol{\gamma},\boldsymbol{\delta})^{\intercal}.$$

The matrix $\mathbf{M} = \mathbf{H}_{\boldsymbol{\gamma},\boldsymbol{\delta}} - \mathbf{H}_{\boldsymbol{\gamma}\boldsymbol{\delta},\mathbf{y}}\cdot\mathbf{H}_{\mathbf{y}}^{-1}\cdot\mathbf{H}_{\boldsymbol{\gamma}\boldsymbol{\delta},\mathbf{y}}^{\intercal}$ is the Schur complement of the block $\mathbf{H_y}$ of $\mathbf{H}$. In fact, we have the identity $\mathrm{Det}(\mathbf{H}) = \mathrm{Det}(\mathbf{H_y})\mathrm{Det}(\mathbf{M})$ and in particular $\mathbf{M}$ is negative definite. A Gaussian integration then yields

$$I_2 = \Big(\frac{1}{\Delta^{2(q-1)^2}\mathrm{Det}(-\mathbf{M})}\Big)^{1/2} = \Big(\frac{\mathrm{Det}(-\mathbf{H_y})}{\Delta^{2(q-1)^2}\mathrm{Det}(-\mathbf{H})}\Big)^{1/2}. \tag{115}$$

Combining equations (113), (114), (115), we obtain the statement of the lemma. □

4.4 The determinants

This section addresses the computation of the determinants of the Hessians in Lemmas 44 and 45. The calculations are quite complex since one has to make a choice of free variables, do the substitutions, differentiate, and then hope that the structure of the problem will prevail in the determinants. Pushing this procedure in our setting leads to complications since the choice of free variables takes away much of the combinatorial structure of the problem. We follow a different path, which amongst other things, reveals that the determinants, via the matrix-tree theorem, correspond to counting weighted trees in appropriate graphs.

The proof has two parts. The first part connects different formulations of the Hessian of a constrained maximization in an abstract setting. Essentially, this puts together well known concepts from optimization in a way that will allow to stay as close as possible to the combinatorial structure of the determinants. The second part specialises the work of the first part to compute the required determinants and is unavoidably more computational.

4.4.1 Hessian formulations for constrained problems

The setting of this section is the following: we are given Υ, a function of $\mathbf{z} \in \mathbb{R}^n$, subject to the linear constraints $\mathbf{Az} = \mathbf{b}$, where $\mathbf{A} \in \mathbb{R}^{m\times n}$. The assumption of linear constraints stems from the setting of Lemmas 44 and 45, yet the arguments extend to other constraints as well by considering gradients of these constraints at the point $\mathbf{z}_0$ and implicit functions. Wlog, we will also assume that $\mathbf{b} = \mathbf{0}$.

We are interested in the Hessian $\mathbf{H}^f$ of a full dimensional representation of Υ. A full dimensional representation of Υ consists essentially of substituting an appropriate subset of the variables $\mathbf{z}$ using the constraints $\mathbf{Az} = \mathbf{0}$. Note that the representation is not as much tied to Υ as it is tied to the space $\mathbf{Az} = \mathbf{0}$. Specifically, assume that the row rank of $\mathbf{A}$ is r. In all the relevant constrained functions we consider, the constraints are not linearly independent so such an assumption is necessary. A full dimensional representation of Υ is specified by two submatrices of $\mathbf{A}$ denoted by $(\mathbf{A}_f, \mathbf{A}_{fs})$. The matrix $\mathbf{A}_f$ is a submatrix of $\mathbf{A}$ consisting of r linearly independent rows of $\mathbf{A}$, so that $\mathbf{Az} = \mathbf{0}$ iff $\mathbf{A}_f\,\mathbf{z} = \mathbf{0}$. Then, $\mathbf{A}_{fs}$ is an $r \times r$ submatrix of $\mathbf{A}_f$ which is invertible. The variables corresponding to columns of $\mathbf{A}_{fs}$ are denoted by $\mathbf{z}_s$. The remaining variables $\mathbf{z}_f$ are called free and $\mathbf{A}_{ff}$ is the submatrix of $\mathbf{A}_f$ induced by the columns indexed by $\mathbf{z}_f$. Renaming if needed, the equation $\mathbf{A}_f\,\mathbf{z} = \mathbf{0}$ may be naturally decomposed as

$$\begin{bmatrix} \mathbf{A}_{ff} & \mathbf{A}_{fs} \end{bmatrix} \begin{bmatrix} \mathbf{z}_f \\ \mathbf{z}_s \end{bmatrix} = \mathbf{0}, \text{ so that } \mathbf{z} = \begin{bmatrix} \mathbf{z}_f \\ \mathbf{z}_s \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ -(\mathbf{A}_{fs})^{-1}\mathbf{A}_{ff} \end{bmatrix} \mathbf{z}_f.$$

Thus, we can now think of Υ as a function which is completely determined by the variables $\mathbf{z}_f$ which, in contrast with the variables $\mathbf{z}$, span a full dimensional space.

Denote by $\mathbf{H}$ the unconstrained Hessian of Υ with respect to the variables $\mathbf{z}$ and by $\mathbf{H}^f$ the Hessian of the full dimensional representation of Υ with respect to the variables $\mathbf{z}_f$. The Hessians $\mathbf{H}$, $\mathbf{H}^f$ are connected by the following equation, which follows by straightforward matrix calculus and its proof is omitted.

$$\mathbf{H}^f = \mathbf{S}^\intercal\,\mathbf{H}\,\mathbf{S}, \text{ where } \mathbf{S} = \begin{bmatrix} \mathbf{I} \\ -(\mathbf{A}_{fs})^{-1}\mathbf{A}_{ff} \end{bmatrix}. \tag{116}$$

Note that $\mathbf{H}^f$ is different, though closely related, from the constrained Hessian $\mathbf{H}^c$ of Υ in the subspace $\mathbf{Az} = \mathbf{0}$, see for example [46, Chapter 10]. The constrained Hessian $\mathbf{H}^c$ has infinitely many matrix representations, all of which correspond to similar matrices, that is, matrices with the same set of eigenvalues. A matrix representation may be obtained by first picking an orthonormal basis of the $(n-r)$-dimensional space $\{\mathbf{z} \,|\, \mathbf{A}\mathbf{z} = \mathbf{0}\}$. Let $\mathbf{E}$ denote the $n \times (n-r)$ matrix whose columns are the vectors in the basis. Then a matrix representation of $\mathbf{H}^c$ is given by

$$\mathbf{H}^c = \mathbf{E}^\intercal \, \mathbf{H} \, \mathbf{E}, \tag{117}$$

where $\mathbf{H}$ is as before the unconstrained Hessian of Υ with respect to the variables $\mathbf{z}$. We are ready to prove the following. It is useful to recall here that congruent matrices have the same number of negative, zero and positive eigenvalues.

Lemma 46. $\mathbf{H}^f$ *is congruent to any matrix representation of* $\mathbf{H}^c$*. Moreover, it holds that*

$$\operatorname{Det}\big(\mathbf{H}^f\big) = \operatorname{Det}\big(\mathbf{H}^c\big)\operatorname{Det}\big(\mathbf{A}_f\mathbf{A}_f^\intercal\big)\Big/\operatorname{Det}\big(\mathbf{A}_{fs}\big)^2 .$$

Proof of Lemma 46. The columns of the matrix $\mathbf{S}$ defined in equation (116) form a basis of the space $\{\mathbf{z} \,|\, \mathbf{A}\mathbf{z} = \mathbf{0}\}$. Indeed, $\mathbf{S}$ has clearly full column rank and also $\mathbf{A}_f\,\mathbf{S} = \mathbf{0}$ implying $\mathbf{A}\,\mathbf{S} = \mathbf{0}$ as well. For future use, by a direct evaluation

$$\mathbf{S}^\intercal\mathbf{S} = \mathbf{I} + \mathbf{A}_{ff}^\intercal\big(\mathbf{A}_{fs}\mathbf{A}_{fs}^\intercal\big)^{-1}\mathbf{A}_{ff}, \text{ so } \operatorname{Det}(\mathbf{S}^\intercal\mathbf{S}) = \operatorname{Det}\Big(\mathbf{I} + \mathbf{A}_{ff}\mathbf{A}_{ff}^\intercal\big(\mathbf{A}_{fs}\mathbf{A}_{fs}^\intercal\big)^{-1}\Big),$$

where the latter equality uses Sylvester's determinant theorem. This clearly yields

$$\operatorname{Det}(\mathbf{S}^\intercal\mathbf{S}) = \operatorname{Det}\big(\mathbf{A}_f\mathbf{A}_f^\intercal\big)\Big/\operatorname{Det}(\mathbf{A}_{fs})^2 . \tag{118}$$

Comparing (116) and (117), the only difference is that $\mathbf{S}$ does not necessarily encode an orthonormal basis. Nevertheless, there clearly exists an invertible matrix $\mathbf{P}$ such that $\mathbf{S}\,\mathbf{P}$ consists of orthonormal columns, for example by the Gram-Schmidt process on the columns of $\mathbf{S}$. It follows that $\mathbf{P}^\intercal\,\mathbf{H}^f\,\mathbf{P}$ is a matrix representation of $\mathbf{H}^c$. This proves the first part of the lemma and also gives $\operatorname{Det}\big(\mathbf{H}^c\big) = \operatorname{Det}\big(\mathbf{H}^f\big)\operatorname{Det}(\mathbf{P})^2$.

For the second part, the selection of $\mathbf{P}$ implies that $(\mathbf{S}\,\mathbf{P})^\intercal\mathbf{S}\,\mathbf{P}$ is the identity matrix and hence $\operatorname{Det}(\mathbf{S}^\intercal\mathbf{S})\operatorname{Det}(\mathbf{P})^2 = 1$. The desired equality follows. $\square$

Lemma 46 allows us to focus on the determinant of $\mathbf{H}^c$ or equivalently the product of its eigenvalues. The latter may be handled using bordered Hessians. Specifically, let $\mathbf{A}_f$ be any submatrix of $\mathbf{A}$ induced by r linearly independent rows. Then, λ is an eigenvalue of $\mathbf{H}^c$ iff it is a root of the polynomial

$$p(\lambda) = \mathrm{Det}\left(\begin{bmatrix} \mathbf{0} & \mathbf{A}_f \\ -\mathbf{A}_f^\intercal & \mathbf{H} - \lambda \mathbf{I}_n \end{bmatrix}\right). \tag{119}$$

In our case, deleting rows of $\mathbf{A}$ to obtain $\mathbf{A}_f$ would cause undesirable complications. In the following, we circumvent such deletions by adding suitable "perturbations". We will also allow for certain degrees of freedom to select the perturbations which will be exploited in the computations. We first prove the following.

For a polynomial $p(s)$, $[s^t]p(s)$ denotes the coefficient of s^t in $p(s)$.

Lemma 47. *Let $\mathbf{M} \in \mathbb{R}^{m\times m}$ be a symmetric matrix with rank r and let μ_i, $i = 1, \ldots, m$ be the eigenvalues of $\mathbf{M}$ with corresponding unit eigenvectors $\mathbf{v}_i$, where $\{\mathbf{v}_1, \ldots, \mathbf{v}_m\}$ is an orthonormal basis of $\mathbb{R}^m$. Then, for any symmetric matrix $\mathbf{T} \in \mathbb{R}^{m\times m}$, it holds that*

$$[\varepsilon^{m-r}]\,\mathrm{Det}\big(\varepsilon\mathbf{T} + \mathbf{M}\big) = \prod_{i;\,\mu_i\neq 0} \mu_i \prod_{i;\,\mu_i = 0} \mathbf{v}_i^\intercal\,\mathbf{T}\,\mathbf{v}_i, \tag{120}$$

In particular, if $\mathbf{T}$ is positive semidefinite and $[\mathbf{T}\ \mathbf{M}]$ has full row rank, the rhs of (120) is non-zero.

Proof of Lemma 47. Let $\mathbf{M}(\varepsilon) = \varepsilon\mathbf{T}+\mathbf{M}$ and denote by $\mu_i(\varepsilon), \mathbf{v}_i(\varepsilon)$ the eigenvalues and unit eigenvectors of $\mathbf{M}(\varepsilon)$. Rellich's theorem asserts that $\mu_i(\varepsilon)$ and $\mathbf{v}_i(\varepsilon)$ are analytic functions of ε around $\varepsilon = 0$. By Hadamard's first variation formula, we have $\dfrac{\partial\mu_i}{\partial\varepsilon} = \mathbf{v}_i^\intercal \dfrac{\partial\mathbf{M}}{\partial\varepsilon}\mathbf{v}_i$. At $\varepsilon = 0$, $\mathbf{M}$ has rank r and hence exactly $m - r$ eigenvalues are zero. Thus, for small enough ε,

$$\mathrm{Det}(\mathbf{M}(\varepsilon)) = \varepsilon^{m-r} \prod_{i\colon \mu_i\neq 0} \mu_i \prod_{i\colon \mu_i=0} \mathbf{v}_i^\intercal\mathbf{T}\mathbf{v}_i + O(\varepsilon^{m-r+1}).$$

Hence, $[\varepsilon^{m-r}]\mathrm{Det}(\mathbf{M}(\varepsilon)) \neq 0$ if for every $\mathbf{v}_i \neq \mathbf{0}$ such that $\mathbf{M}\mathbf{v}_i = \mathbf{0}$, we have $\mathbf{v}_i^\intercal\mathbf{T}\mathbf{v}_i \neq 0$. The latter is true. Otherwise, using the positive semidefiniteness of $\mathbf{T}$, we obtain $\mathbf{v}_i^\intercal[\mathbf{T}\ \mathbf{M}] = \mathbf{0}$, contradicting that $[\mathbf{T}\ \mathbf{M}]$ has full row rank. □

The following lemma gives the promised extension of (119).

Lemma 48. *Suppose that* $\mathbf{T}$ *is a diagonal positive semidefinite* $m \times m$ *matrix such that* $[\mathbf{T}\ \mathbf{A}]$ *has full row rank. Let* $\mathbf{H}$ *(resp.* $\mathbf{H}^c$*) be the unconstrained (resp. constrained) Hessian of* Υ *evaluated at a point* $\mathbf{z}_0$*. Then,* λ *is an eigenvalue of* $\mathbf{H}^c$ *iff it is a root of the polynomial*

$$p(\lambda) = [\varepsilon^{m-r}]\,\mathrm{Det}(\mathbf{H}_\lambda) \text{ where } \mathbf{H}_\lambda = \begin{bmatrix} \varepsilon\mathbf{T} & \mathbf{A} \\ -\mathbf{A}^\intercal & \mathbf{H} - \lambda\mathbf{I}_n \end{bmatrix}. \tag{121}$$

Further, if $\mathbf{H}$ *is invertible, then* $\mathrm{Det}\big(\mathbf{H}^c\big) = (-1)^n \mathrm{Det}(\mathbf{H}) \dfrac{[\varepsilon^{m-r}]\,\mathrm{Det}\big(\varepsilon\mathbf{T} + \mathbf{A}\mathbf{H}^{-1}\mathbf{A}^\intercal\big)}{[\varepsilon^{m-r}]\,\mathrm{Det}\big(\varepsilon\mathbf{T} - \mathbf{A}\mathbf{A}^\intercal\big)}$.

Proof of Lemma 48. Let $\mathbf{T} = (t_{i,j})_{i,j\in[m]}$ and $\mathbf{H}_\lambda = (h_{i,j})_{i,j\in[m+n]}$. Let $\mathcal{W} = \binom{[m]}{m-r}$ and for $W \in \mathcal{W}$ let $P_W = \{\sigma \in S_{m+n} \,|\, \{i \in [m] \,|\, \sigma(i) = i\} = W\}$. Since $\mathbf{T}$ is diagonal, by Leibniz's formula,

$$p(\lambda) = [\varepsilon^{m-r}]\mathrm{Det}(\mathbf{H}_\lambda) = \sum_{W\in\mathcal{W}} \prod_{i\in W} t_{i,i} \sum_{\sigma\in P_W} \mathrm{sgn}(\sigma) \prod_{i\in[m+n]\setminus W} h_{i,\sigma(i)}. \tag{122}$$

Let $\mathbf{A}_{[m]\setminus W}$ be the $r \times n$ submatrix of $\mathbf{A}$ which is obtain by excluding the rows indexed by W. Identifying permutations in P_W with permutations of $[n+r]$ in the natural way, we obtain

$$\sum_{\sigma\in P_W} \mathrm{sgn}(\sigma) \prod_{i\in[m+n]\setminus W} h_{i,\sigma(i)} = \mathrm{Det}\left(\begin{bmatrix} \mathbf{0}_r & \mathbf{A}_{[m]\setminus W} \\ -\mathbf{A}^\intercal_{[m]\setminus W} & \mathbf{H} - \lambda\mathbf{I}_n \end{bmatrix}\right) \equiv q_W(\lambda). \tag{123}$$

If $\mathbf{A}_{[m]\setminus W}$ has row rank $< r$, then $q_W(\lambda)$ is 0. Otherwise, the roots of $q_W(\lambda)$ are the eigenvalues of $\mathbf{H}^c$, c.f. (119). By (122), this is also the case for $p(\lambda)$, provided it is not identically zero.

To prove that $p(\lambda)$ is nonzero, we prove that the leading coefficient of $p(\lambda)$ is nonzero. Starting from (123) and plugging into (122), this can easily be seen to equal

$$[\varepsilon^{m-r}]\mathrm{Det}\left(\begin{bmatrix} \varepsilon\mathbf{T} & \mathbf{A} \\ -\mathbf{A}^\intercal & -\mathbf{I}_n \end{bmatrix}\right) = [\varepsilon^{m-r}](-1)^n \mathrm{Det}(\varepsilon\mathbf{T} - \mathbf{A}\mathbf{A}^\intercal),$$

where in the latter equality we used the Schur complement of the block $-\mathbf{I}_n$. This is non-zero by Lemma 47.

The determinant of $\mathbf{H}^c$ is the product of its eigenvalues. This equals $p(0)$ divided by the leading coefficient of $p(\lambda)$. The latter has already been computed. The former, using the Schur complement of the invertible $\mathbf{H}$, is equal to $[\varepsilon^{m-r}]\mathrm{Det}(\mathbf{H})\,\mathrm{Det}\big(\varepsilon\mathbf{T}-\mathbf{A}\mathbf{H}^{-1}\mathbf{A}^\intercal\big)$. This concludes the proof. □

Finally, we combine the above lemmas to obtain the following.

Lemma 49. *Let Υ be a function of $\mathbf{z}\in\mathbb{R}^n$ subject to the linear constraints $\mathbf{A}\mathbf{z}=\mathbf{b}$, where $\mathbf{A}\in\mathbb{R}^{m\times n}$ where $\mathbf{A}$ has rank r. Let $\big(\mathbf{A}_f,\mathbf{A}_{fs}\big)$ specify a full dimensional representation of Υ and let $\mathbf{H}^f$ be the corresponding Hessian of Υ evaluated at a point $\mathbf{z}_0$.*

Suppose $\mathbf{T}$ is a positive semidefinite diagonal matrix with dimensions $m\times m$ such that $[\mathbf{T}\ \mathbf{A}]$ has full row rank. Let $\mathbf{H}$ (resp. $\mathbf{H}^c$) be the unconstrained Hessian of Υ evaluated at $\mathbf{z}_0$. If $\mathbf{H}$ is invertible, then

$$\mathrm{Det}\big(-\mathbf{H}^f\big)=\frac{L\big(\mathbf{A}_f,\mathbf{A},\mathbf{T}\big)}{\mathrm{Det}\big(\mathbf{A}_{fs}\big)^2}\,\mathrm{Det}(\mathbf{H})\,[\varepsilon^{m-r}]\,\mathrm{Det}\big(\varepsilon\mathbf{T}-\mathbf{A}\mathbf{H}^{-1}\mathbf{A}^\intercal\big), \tag{124}$$

where $L\big(\mathbf{A}_f,\mathbf{A},\mathbf{T}\big)=\mathrm{Det}\big(\mathbf{A}_f\mathbf{A}_f^\intercal\big)\Big/[\varepsilon^{m-r}]\,\mathrm{Det}\big(\varepsilon\mathbf{T}-\mathbf{A}\mathbf{A}^\intercal\big)$.

Proof of Lemma 49. Just combine Lemmas 46 and 48. The minor sign change $-\mathbf{H}^f$ in the statement can easily be accounted by applying the lemmas to the function $-\Upsilon$. □

The rhs of (124) has two qualitatively different factors: the factor $\frac{L(\mathbf{A}_f,\mathbf{A},\mathbf{T})}{\mathrm{Det}(\mathbf{A}_{fs})^2}$ depends on the specific full dimensional representation, while the remaining factor is tied to the Hessian of Υ. The technical convenience of Lemma 49 is dual: first, it gives an explicit formula for $\mathrm{Det}\big(-\mathbf{H}^f\big)$ without doing substitutions which would hinder the combinatorial view of the constraints $\mathbf{A}$; second, it isolates the deletions of rows of $\mathbf{A}$ in the factor $L(\mathbf{A}_f,\mathbf{A})$ and leaves untouched the more complicated matrix $\mathbf{A}\mathbf{H}^{-1}\mathbf{A}^\intercal$.

4.4.2 The computations

In this section, we utilize Lemma 49 to compute the determinants in Lemmas 44 and 45.

Notation: For a vector $\mathbf{z}\in\mathbb{R}^n$, $\mathbf{z}^D$ denotes the $n\times n$ diagonal matrix $\mathbf{diag}\{z_1,\ldots,z_n\}$. For vectors $\mathbf{z}_i\in\mathbb{R}^{m_i}$, $i=1,\ldots,t$ we denote by $[\mathbf{z}_1,\ldots,\mathbf{z}_t]^\intercal$ the $\mathbb{R}^{\sum_i m_i}$ vector which is the concatenation of the vectors $\mathbf{z}_1,\ldots,\mathbf{z}_t$. For matrices $\mathbf{A}$ and $\mathbf{B}$, $\mathbf{A}\otimes\mathbf{B}$ will denote the Kronecker product of $\mathbf{A},\mathbf{B}$, while $\mathbf{A}\oplus\mathbf{B}$ is the direct sum of $\mathbf{A},\mathbf{B}$, that is, the block diagonal matrix $\mathbf{diag}\{\mathbf{A},\mathbf{B}\}$. The expression $\oplus_2\mathbf{A}$ is a shorthand for $\mathbf{A}\oplus\mathbf{A}$. Further, $\mathbf{I}_n$

denotes the identity matrix of dimensions $n \times n$. Finally, $\mathbf{1}_n, \mathbf{0}_n$ denote the all-one and all-zero n-dimensional vector.

To start, the equality constraints in (110) and (112) may be written in the form

$$\mathbf{A}_1 \mathbf{x} = \mathbf{0}, \qquad \mathbf{A}_2 \left[\boldsymbol{\gamma},\, \boldsymbol{\delta},\, \mathbf{y}\right]^\intercal = \left[\boldsymbol{\alpha},\, \boldsymbol{\alpha},\, \boldsymbol{\beta},\, \boldsymbol{\beta},\, \mathbf{0}_{2q^2}\right]^\intercal.$$

The matrices $\mathbf{A}_1$ and $\mathbf{A}_2$ have dimensions $2q \times (q^2 - |P_1|)$ and $(4q + 2q^2) \times (2q^2 + q^4 - |P_2|)$, respectively, where P_1, P_2 are defined in (109), (111). Note that we exclude from consideration variables x_{ij} and y_{ikjl} which are hard-coded to zero. This is done to ensure that the unconstrained Hessians are invertible, so that Lemma 49 applies directly. It will be useful to decompose the matrix $\mathbf{A}_2$ as

$$\mathbf{A}_2 = \begin{bmatrix} \mathbf{A}_{2,\boldsymbol{\gamma}\boldsymbol{\delta}} & \mathbf{0} \\ -\mathbf{I}_{2q^2} & \mathbf{A}_{2,\mathbf{y}} \end{bmatrix}, \tag{125}$$

where $\mathbf{A}_{2,\boldsymbol{\gamma}\boldsymbol{\delta}}, \mathbf{A}_{2,\mathbf{y}}$ have dimensions $4q \times 2q^2$ and $2q^2 \times (q^4 - |P_2|)$, respectively.

The easiest way to handle the matrices $\mathbf{A}_1, \mathbf{A}_{2,\boldsymbol{\gamma}\boldsymbol{\delta}}, \mathbf{A}_{2,\mathbf{y}}$ is as incidence matrices of appropriate bipartite graphs. This view will simplify the arguments significantly. In particular, for an undirected graph G, we denote by $\mathbf{A}_G$ the 0,1 incidence matrix of G, by $\mathbf{R}_G$ the adjacency matrix of G, by $\mathbf{D}_G$ the diagonal matrix whose diagonal entries are equal to the degrees of the vertices in G and by $\boldsymbol{\Lambda}_G$ the matrix $\mathbf{D}_G + \mathbf{R}_G$. We will also be interested in the case where the graph G is weighted. We assume that the weights on the edges are positive and are given by the diagonal entries of a square diagonal matrix $\mathbf{W}_G$. We denote by $\mathbf{R}^w_G, \mathbf{D}^w_G, \boldsymbol{\Lambda}^w_G$ the weighted versions of the matrices $\mathbf{R}_G, \mathbf{D}_G, \boldsymbol{\Lambda}_G$. It is well known and easy to check that

$$\mathbf{A}_G \mathbf{A}_G^\intercal = \boldsymbol{\Lambda}_G, \;\; \mathbf{A}_G \mathbf{W}_G \mathbf{A}_G^\intercal = \boldsymbol{\Lambda}^w_G. \tag{126}$$

Before defining graphs to which (126) will be applied, it will be useful to state the unconstrained Hessians. From Lemmas 44 and 45, the unconstrained Hessians of interest are: (i) $\mathbf{H}_{1,\mathbf{x}}$, the Hessian of Υ_1/Δ with respect to $\mathbf{x}$ when $\boldsymbol{\alpha}, \boldsymbol{\beta}$ are fixed, (ii) $\mathbf{H}_{2,\mathbf{y}}$, the Hessian of Υ_2/Δ with respect to $\mathbf{y}$ when $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\delta}$ are fixed, (iii) $\mathbf{H}_2$, the Hessian of Υ_2/Δ with respect to $\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y}$ when $\boldsymbol{\alpha}, \boldsymbol{\beta}$ are fixed. These three matrices are all diagonal and it is straightforward to see

$$\mathrm{Det}(\mathbf{H}_{1,\mathbf{x}})^{-1} = (-1)^{q^2-|P_1|} \prod_{i,j\in P_1} x_{ij}, \qquad \mathrm{Det}(\mathbf{H}_{2,\mathbf{y}})^{-1} = (-1)^{q^4-|P_2|} \prod_{(i,k,j,l)\in P_2} y_{ikjl},$$
$$\mathrm{Det}(\mathbf{H}_2)^{-1} = \mathrm{Det}(\mathbf{H}_{2,\mathbf{y}})^{-1} \Big(\frac{\Delta-1}{\Delta}\Big)^{2q^2} \prod_{i,k\in[q]} \gamma_{ik} \prod_{j,l\in[q]} \delta_{jl}, \tag{127}$$
$$(\mathbf{H}_{1,\mathbf{x}})^{-1} = -\mathbf{x}^D,\ (\mathbf{H}_{2,\mathbf{y}})^{-1} = -\mathbf{y}^D,\ (\mathbf{H}_2)^{-1} = \frac{\Delta}{\Delta-1}\boldsymbol{\gamma}^D \oplus \frac{\Delta}{\Delta-1}\boldsymbol{\delta}^D \oplus (\mathbf{H}_{2,\mathbf{y}})^{-1}.$$

We next proceed to the definition of the bipartite graphs for each of the above matrices. For $\mathbf{A}_{1,\mathbf{x}}$, the related graph $G_\mathbf{x}$ is a bipartite graph with vertex bipartition $([q],[q])$ and an edge (i,j) is present iff $(i,j)\in P_1$, that is, $B_{ij}>0$. Since $\mathbf{B}$ is symmetric and irreducible, $G_\mathbf{x}$ is undirected and connected. For $\mathbf{A}_{2,\mathbf{y}}$, the related graph $G_\mathbf{y}$ is a bipartite graph with vertex bipartition $([q]^2,[q]^2)$ and an edge $\big((i,k),(j,l)\big)$ is present iff $(i,k,j,l)\in P_2$, that is, $B_{ikjl}>0$. For $\mathbf{A}_{2,\boldsymbol{\gamma\delta}}$, the graph has two connected components indexed by $\boldsymbol{\gamma},\boldsymbol{\delta}$, respectively. Each is isomorphic to the complete bipartite graph $K_{q,q}$ with vertex bipartition $([q],[q])$. We denote by $\mathbf{A}_{2q,q^2}$ the incidence matrix of $K_{q,q}$, so that $\mathbf{A}_{2,\boldsymbol{\gamma\delta}} = \oplus_2 \mathbf{A}_{2q,q^2}$.

The weights on the edges of the graphs $G_\mathbf{x}, G_\mathbf{y}, G_{\boldsymbol{\gamma\delta}}$ are given in principle by (127). Explicitly, an edge (i,j) in $G_\mathbf{x}$ has weight x_{ij}, an edge $\big((i,k),(j,l)\big)$ in $G_\mathbf{y}$ has weight y_{ikjl} and an edge (i,j) in $G_{\boldsymbol{\gamma\delta}}$ is γ_{ij} if it belongs to the connected component indexed by $\boldsymbol{\gamma}$ and δ_{ij} if it belongs to the connected component indexed by $\boldsymbol{\delta}$. The factor $\Delta/(\Delta-1)$ in $(\mathbf{H}_2)^{-1}$ will be accounted otherwise. In the languange of (126),

$$\mathbf{W}_{G_\mathbf{x}} = \mathbf{x}^D,\ \mathbf{W}_{G_\mathbf{y}} = \mathbf{y}^D,\ \mathbf{W}_{G_{\boldsymbol{\gamma\delta}}} = \boldsymbol{\gamma}^D \oplus \boldsymbol{\delta}^D. \tag{128}$$

The otherwise straightforward application of (126) on each of the graphs $G_\mathbf{x}, G_\mathbf{y}, G_{\boldsymbol{\gamma\delta}}$ is useful to do explicitly in order to decompose the resulting matrices. In particular, since these graphs are undirected and bipartite, we have

$$\boldsymbol{\Lambda}^w(G_\mathbf{x}) = \begin{bmatrix} \boldsymbol{\alpha}^D & \mathbf{S}_\mathbf{x} \\ \mathbf{S}_\mathbf{x}^\intercal & \boldsymbol{\beta}^D \end{bmatrix},\ \boldsymbol{\Lambda}^w(G_\mathbf{y}) = \begin{bmatrix} \boldsymbol{\gamma}^D & \mathbf{S}_\mathbf{y} \\ \mathbf{S}_\mathbf{y}^\intercal & \boldsymbol{\delta}^D \end{bmatrix}, \tag{129}$$
$$\boldsymbol{\Lambda}^w(G_{\boldsymbol{\gamma\delta}}) = \begin{bmatrix} \boldsymbol{\alpha}^D & \mathbf{S}_\gamma \\ \mathbf{S}_\gamma & \boldsymbol{\alpha}^D \end{bmatrix} \bigoplus \begin{bmatrix} \boldsymbol{\beta}^D & \mathbf{S}_{\boldsymbol{\delta}} \\ \mathbf{S}_{\boldsymbol{\delta}} & \boldsymbol{\beta}^D \end{bmatrix}, \tag{130}$$

where $\mathbf{S}_{\mathbf{x}}, \mathbf{S}_{\boldsymbol{\gamma}}, \mathbf{S}_{\boldsymbol{\delta}}$ are the $q \times q$ matrices whose (i,j) entries are x_{ij}, γ_{ij}, δ_{ij}, respectively, and $\mathbf{S}_{\mathbf{y}}$ is the $q^2 \times q^2$ matrix whose $((i,k),(j,l))$ entry is y_{ikjl}. Note that the matrices $\mathbf{D}^w$ were substituted using (110) and (112).

We are now ready to evaluate these matrices at a global maximum $(\boldsymbol{\alpha}^*, \boldsymbol{\beta}^*, \mathbf{x}^*)$ of Υ_1 and $(\boldsymbol{\gamma}^*, \boldsymbol{\delta}^*, \mathbf{y}^*)$ of Υ_2. From now on, we will not explicitly use asterisks in the notation with the understanding that the values of all the variables are fixed to their optimal values. Note by Lemma 31, we have

$$\boldsymbol{\gamma} = \boldsymbol{\alpha} \otimes \boldsymbol{\alpha},\ \boldsymbol{\delta} = \boldsymbol{\beta} \otimes \boldsymbol{\beta},\ \mathbf{y} = \mathbf{x} \otimes \mathbf{x}. \tag{131}$$

We will apply Lemma 49 to the matrices $\mathbf{H}_{1,\mathbf{x}}, \mathbf{H}_{2,\mathbf{y}}, \mathbf{H}_2$ using the matrices

$$\mathbf{T}_1 = \boldsymbol{\alpha}^D \oplus \boldsymbol{\beta}^D,\ \mathbf{T}_{2,\mathbf{y}} = \boldsymbol{\gamma}^D \oplus \boldsymbol{\delta}^D,\ \mathbf{T}_2 = \boldsymbol{\alpha}^D \oplus \boldsymbol{\alpha}^D \oplus \boldsymbol{\beta}^D \oplus \boldsymbol{\beta}^D \oplus \mathbf{0}_{2q^2}, \tag{132}$$

respectively. We first compute the determinants of $\mathbf{M}_1 = \varepsilon \mathbf{T}_1 - \mathbf{A}_1(\mathbf{H}_{1,\mathbf{x}})^{-1}\mathbf{A}_1^{\intercal}$, $\mathbf{M}_{2,\mathbf{y}} = \varepsilon \mathbf{T}_{2,\mathbf{y}} - \mathbf{A}_{2,\mathbf{y}}(\mathbf{H}_{2,\mathbf{y}})^{-1}\mathbf{A}_{2,\mathbf{y}}^{\intercal}$, $\mathbf{M}_2 = \varepsilon \mathbf{T}_2 - \mathbf{A}_2(\mathbf{H}_2)^{-1}\mathbf{A}_2^{\intercal}$, which contribute the most interesting factors in Lemma 49.

We begin with the simplest of these matrices, $\mathbf{M}_1$. Note that $\mathbf{A}_1$ has rank $2q-1$, so by Lemma 49 we want to compute $[\varepsilon]\operatorname{Det}(\mathbf{M}_1)$. Using (126), (127), (128), (132), it is straightforward to check that $\mathbf{M}_1$ has the following form

$$\mathbf{M}_1 = \begin{bmatrix} \boldsymbol{\alpha}^D(\varepsilon \mathbf{I}_q + \mathbf{I}_q) & \mathbf{S}_{\mathbf{x}} \\ \mathbf{S}_{\mathbf{x}}^{\intercal} & \boldsymbol{\beta}^D(\varepsilon \mathbf{I}_q + \mathbf{I}_q) \end{bmatrix},\ \text{so } \operatorname{Det}(\mathbf{M}_1) = \Big(\prod_{i\in[q]} \alpha_i \prod_{j\in[q]} \beta_j\Big) \operatorname{Det}\big(\varepsilon \mathbf{I}_q + \mathbf{I}_q + \mathbf{J}\big), \tag{133}$$

where $\mathbf{J}$ is the matrix in Lemma 38. Note that in Equation (133), to get the second equality, we did the following operations on $\mathbf{M}_1$: for $i = 1, \ldots, q$, we divided the i-th row of by $\sqrt{\alpha_i}$, the i-th column by $\sqrt{\alpha_i}$, the $(i+q)$-th row by $\sqrt{\beta_i}$, the $(i+q)$-th column by $\sqrt{\beta_i}$. The eigenvalues of the matrix $\varepsilon \mathbf{I}_q + \mathbf{I}_q + \mathbf{J}$ are appropriate shifts of the eigenvalues of $\mathbf{J}$ and are given by

$$\varepsilon, \varepsilon + 2, \varepsilon + 1 \pm \lambda_1, \ldots, \varepsilon + 1 \pm \lambda_{q-1},$$

c.f., Lemma 38 for the definition of the λ_i and their properties. We thus obtain

$$[\varepsilon]\operatorname{Det}\big(\varepsilon \mathbf{T}_1 - \mathbf{A}_1(\mathbf{H}_{1,\mathbf{x}})^{-1}\mathbf{A}_1^{\intercal}\big) = 2 \prod_{i\in[q]} \alpha_i \prod_{j\in[q]} \beta_j \prod_{i\in[q-1]} \big(1 - \lambda_i^2\big). \tag{134}$$

In a completely analogous manner, we can also obtain

$$[\varepsilon]\,\mathrm{Det}\big(\varepsilon \mathbf{T}_{2,\mathbf{y}} - \mathbf{A}_{2,\mathbf{y}}(\mathbf{H}_{2,\mathbf{y}})^{-1}\mathbf{A}_{2,\mathbf{y}}^{\intercal}\big) = 2\prod_{i,k\in[q]}\gamma_{ik}\prod_{j,l\in[q]}\delta_{jl}\prod_{i\in[q-1]}\big(1-\lambda_i^2\big)^2\prod_{i,j\in[q-1]}\big(1-\lambda_i^2\lambda_j^2\big). \tag{135}$$

Note here that the λ_i's are again the same as in Lemma 38. The more complicated product in (135) is due to the eigenvalues of a Kronecker product. We omit the details since these can be inferred from Lemma 51 which will be stated and proved afterwards. There, the shifting of the eigenvalues is different, but otherwise everything else is the same.

The determinant of the matrix $\mathbf{M}_2$ is quite more complicated to compute due to its more intricate complicate block structure, which requires using Schur's complement formula to handle. The following is proved in Section 4.4.3.1.

$$\frac{[\varepsilon^3]\mathrm{Det}\big(\varepsilon \mathbf{T}_2 - \mathbf{A}_2(\mathbf{H}_2')^{-1}\mathbf{A}_2^{\intercal}\big)}{\prod_{i\in[q]}\alpha_i\prod_{j\in[q]}\beta_j\prod_{i,k\in[q]}\gamma_{ik}\prod_{j,l\in[q]}\delta_{jl}} = \frac{4\Delta^{4q-2}}{(\Delta-1)^{2q^2}}\prod_{i,j\in[q-1]}\big(1-(\Delta-1)^2\lambda_i^2\lambda_j^2\big)\prod_{i\in[q-1]}\big(1-\lambda_i^2\big)^2. \tag{136}$$

Equations (127), (134), (135), (136) deal with the factors in Lemma 49 which are tied to the Hessians of the functions. While these contribute the most interesting factors, some care is needed to deal with the remaining factors. This is accomplished in the following lemma, which is proved in Section 4.4.3.2.

Lemma 50. *Let* $\big((\mathbf{A}_1)_f, (\mathbf{A}_1)_{fs}\big)$, $\big((\mathbf{A}_{2,\mathbf{y}})_f, (\mathbf{A}_{2,\mathbf{y}})_{fs}\big)$, $\big((\mathbf{A}_2)_f, (\mathbf{A}_2)_{fs}\big)$ *specify arbitrary full dimensional representations of the spaces* $\mathbf{A}_1\mathbf{x} = \mathbf{0}$, $\mathbf{A}_{2,\mathbf{y}}\mathbf{y} = \mathbf{0}$, $\mathbf{A}_2\,[\boldsymbol{\gamma}, \boldsymbol{\delta}, \mathbf{y}]^{\intercal} = \mathbf{0}$. *Then:*

$$\mathrm{Det}\big((\mathbf{A}_1)_{fs}\big)^2 = \mathrm{Det}\big((\mathbf{A}_{2,\mathbf{y}})_{fs}\big)^2 = \mathrm{Det}\big((\mathbf{A}_2)_{fs}\big)^2 = 1, \tag{137}$$

$$L\big((\mathbf{A}_1)_f, \mathbf{A}_1, \mathbf{T}_1\big) = L\big((\mathbf{A}_{2,\mathbf{y}})_f, \mathbf{A}_{2,\mathbf{y}}, \mathbf{T}_{2,\mathbf{y}}\big) = 1/2, \quad L\big((\mathbf{A}_2)_f, \mathbf{A}_2, \mathbf{T}_2\big) = 1/4, \tag{138}$$

where $\mathbf{T}_1, \mathbf{T}_{2,\mathbf{y}}, \mathbf{T}_2$ *are given by* (132) *and the quantities in* (138) *are defined in Lemma 49.*

We are now ready to finish the proof of Lemma 43.

Proof of Lemma 43. Apply Lemma 49 three times to unravel the determinants appearing in Lemmas 44 and 45. Each of the resulting quantities has been computed and appears in one of (127), (134), (135), (136) or Lemma 50. The proof of the lemma is completed with careful but otherwise straightforward substitutions. □

4.4.3 Remaining proofs

4.4.3.1 Details for the Second Moment Determinant

We give here the details of the proof of (136), which gives the determinant of the more intricate matrix $\mathbf{M}_2 := \varepsilon \mathbf{T}_2 - \mathbf{A}_2(\mathbf{H}_2)^{-1}\mathbf{A}_2^\intercal$. The first step of the computation is the same as in the previous arguments and consists of writing out its block structure and then appropriately normalizing the resulting matrix. Here the normalization is slightly more intricate. The analog of (133) is

$$\mathrm{Det}\big(\mathbf{M}_2\big) = \mathrm{Det}\big(\mathbf{H}_2'\big) \prod_{i\in[q]} \alpha_i \prod_{j\in[q]} \beta_j \prod_{i,k\in[q]} \gamma_{ik} \prod_{j,l\in[q]} \delta_{jl}, \tag{139}$$

where

$$\mathbf{H}_2' = \frac{\Delta}{\Delta-1}\begin{bmatrix} \varepsilon\frac{\Delta-1}{\Delta}\mathbf{I}_{4q} - \mathbf{V}\mathbf{V}^\intercal & \mathbf{V} \\ \mathbf{V}^\intercal & -\frac{\Delta-1}{\Delta}\mathbf{W} \end{bmatrix}, \tag{140}$$

and the matrices $\mathbf{W}, \mathbf{V}$ are given by

$$\mathbf{W} = \frac{1}{\Delta-1}\mathbf{I}_{2q^2} - \begin{bmatrix} \mathbf{0} & \mathbf{L}\otimes\mathbf{L} \\ (\mathbf{L}\otimes\mathbf{L})^\intercal & \mathbf{0} \end{bmatrix},$$
$$\mathbf{V} = (\oplus_2\boldsymbol{\alpha}^D)^{-1/2}\mathbf{A}_{2q,q^2}\,(\boldsymbol{\gamma}^D)^{1/2}\bigoplus(\oplus_2\boldsymbol{\beta}^D)^{-1/2}\mathbf{A}_{2q,q^2}\,(\boldsymbol{\delta}^D)^{1/2}.$$

Equation (139) can be obtained by performing the following operations on $\mathbf{M}_2$: for $i = 1,\dots,q$, divide the $i, i+q$ rows by $\sqrt{\alpha_i}$, the $i+2q, i+3q$ rows by $\sqrt{\beta_i}$; for $i,j = 1,\dots,q$, divide the $4q+q(i-1)+j$ row by $\sqrt{\gamma_{ij}}$ and the $4q+q^2+q(i-1)+j$ row by $\sqrt{\delta_{ij}}$; and the same operations on columns. These operations are captured by the matrices $(\oplus_2\boldsymbol{\alpha}^D)^{-1/2}$, $(\oplus_2\boldsymbol{\beta}^D)^{-1/2}$, $(\boldsymbol{\gamma}^D)^{1/2}$, $(\boldsymbol{\delta}^D)^{-1/2}$ which appear in the matrix $\mathbf{V}$. The matrix $\mathbf{W}$ is the normalized form of $-\mathbf{A}_{2,\mathbf{y}}(\mathbf{H}_{2,\mathbf{y}})^{-1}\mathbf{A}_{2,\mathbf{y}}^\intercal$ but with some corrections in the diagonal which stem from the coupled form of the matrix $\mathbf{A}_2$ and the factor $\Delta/(\Delta-1)$ which appears in $\mathbf{H}_2$, c.f., Equations (125) and (127).

In light of (139), it suffices to compute $\mathrm{Det}(\mathbf{H}_2')$. To do this, we proceed by taking the Schur complement of the matrix $\mathbf{W}$. For this, we need that $\mathbf{W}$ is invertible, as the following lemma guarantees.

Lemma 51. *Let $t = 1/(\Delta - 1)$. In the notation and setting of Lemma 38, the spectrum of* $\mathbf{W}$ *is given by*

$t\pm 1, t\pm\lambda_1, t\pm\lambda_1, \ldots, t\pm\lambda_{q-1}, t\pm\lambda_{q-1}, t\pm\lambda_1^2, t\pm\lambda_1\lambda_2, \ldots, t\pm\lambda_1\lambda_{q-1}, t\pm\lambda_2\lambda_1, \ldots, t\pm\lambda_{q-1}^2.$

Recall that the λ_i are non-negative and $\max\lambda_i < \frac{1}{\Delta-1}$. Hence, $\mathbf{W}$ is invertible. We also have

$$\mathrm{Det}(\mathbf{W}) = -\frac{\Delta(\Delta-2)}{(\Delta-1)^{2q^2}} \prod_{i\in[q-1]} \left(1-(\Delta-1)^2\lambda_i^2\right)^2 \prod_{i,j\in[q-1]} \left(1-(\Delta-1)^2\lambda_i^2\lambda_j^2\right). \tag{141}$$

Proof of Lemma 51. The matrix $t\mathbf{I}_{2q^2}$ shifts the eigenvalues of $\begin{bmatrix} \mathbf{0} & \mathbf{L}\otimes\mathbf{L} \\ (\mathbf{L}\otimes\mathbf{L})^\intercal & \mathbf{0} \end{bmatrix}$ by t, so it suffices to find the eigenvalues of the latter matrix.

These come in pairs $(\sigma_i^2, -\sigma_i^2)$, $i \in [q^2]$, where $\{\sigma_i^2\}_{i\in[q^2]}$ are the singular values of $\mathbf{L}\otimes\mathbf{L}$. By properties of the Kronecker Product, these are equal to $\{(\sigma_i')^2(\sigma_j')^2\}_{i,j\in[q]}$, where $\{(\sigma_i')^2\}_{i\in[q]}$ are the singular values of $\mathbf{L}$. The latter are the non-negative eigenvalues of $\begin{bmatrix} \mathbf{0} & \mathbf{L} \\ \mathbf{L}^\intercal & \mathbf{0} \end{bmatrix}$ and are precisely $\{1, \lambda_1, \ldots, \lambda_{q-1}\}$. The lemma follows. $\square$

After taking the Schur complement of the matrix $\mathbf{W}$, we obtain

$$\mathrm{Det}(\mathbf{H}_2') = \left(\frac{\Delta}{\Delta-1}\right)^{4q} \mathrm{Det}\big(\mathbf{W}\big)\,\mathrm{Det}\Big(\varepsilon\frac{\Delta-1}{\Delta}\mathbf{I}_{4q} + \mathbf{Z}\Big), \text{ with } \mathbf{Z} = \frac{\Delta}{\Delta-1}\mathbf{V}\mathbf{W}^{-1}\mathbf{V}^\intercal - \mathbf{V}\mathbf{V}^\intercal. \tag{142}$$

We are left with the evaluation of $\mathrm{Det}\big(\varepsilon\frac{\Delta-1}{\Delta}\mathbf{I}_{4q} + \mathbf{Z}\big)$. This is possible once we know the eigenvalues of $\mathbf{Z}$, since the identity matrix just shifts its eigenvalues. The complication here is the nontrivial inverse of $\mathbf{W}$ appearing in the formulation of $\mathbf{Z}$. The next couple of lemmas circumvent the computation of $\mathbf{W}^{-1}$ and reduce it to the inverse of a much simpler matrix.

Lemma 52. *In the notation and setting of Lemma 38, it holds that*

$$\mathbf{V}\mathbf{W} = \Big(\frac{1}{\Delta-1}\mathbf{I}_{4q} - \mathbf{L}'\Big)\mathbf{V}, \textit{ where } \mathbf{L}' = \begin{bmatrix} \mathbf{0} & \mathbf{L}\oplus\mathbf{L} \\ (\mathbf{L}\oplus\mathbf{L})^\intercal & \mathbf{0} \end{bmatrix}.$$

Lemma 52 is useful only if the matrix $\frac{1}{\Delta-1}\mathbf{I}_{4q} - \mathbf{L}'$ is invertible. This is an immediate corollary of the following.

Lemma 53. *The spectrum of the matrix* $\mathbf{L}'$ *defined in Lemma 52 is given by*

$$\pm 1, \pm 1, \pm\lambda_1, \pm\lambda_1, \ldots, \pm\lambda_{q-1}, \pm\lambda_{q-1}.$$

Recall that the λ_i *are non-negative and* $\max_i \lambda_i < \frac{1}{\Delta-1}$. *Hence,* $\frac{1}{\Delta-1}\mathbf{I}_{4q} - \mathbf{L}'$ *is invertible.*

Proof of Lemma 53. The matrix $\mathbf{L}'$ is similar by a permutation matrix to the direct sum of two copies of the matrix $\mathbf{J}$ of Lemma 38. Thus, the eigenvalues of $\mathbf{L}'$ can be easily derived by properties of the direct sum. The matrix $\frac{1}{\Delta-1}\mathbf{I}_{4q}$ just shifts the eigenvalues of $-\mathbf{L}'$ by $1/(\Delta-1)$. The lemma follows. $\square$

Combining Lemmas 52 and 53, we obtain

$$\mathbf{Z} = \Big[-\mathbf{I}_{4q} + \frac{\Delta}{\Delta-1}\Big(\frac{1}{\Delta-1}\mathbf{I}_{4q} - \mathbf{L}'\Big)^{-1}\Big]\mathbf{V}\mathbf{V}^{\intercal} = \big(\mathbf{I}_{4q} + \mathbf{L}'\big)\big(\frac{1}{\Delta-1}\mathbf{I}_{4q} - \mathbf{L}'\big)^{-1}\mathbf{V}\mathbf{V}^{\intercal}, \tag{143}$$

By (142), $\mathbf{Z}$ is trivially symmetric. As a consequence of Equation (143), we are now in position to study the eigenvalues of $\mathbf{Z}$.

Lemma 54. *The spectrum of* $\mathbf{Z}$ *is given by*

$$0, 0, 0, 2f(1), f(\pm\lambda_1), f(\pm\lambda_1), \ldots, f(\pm\lambda_{q-1}), f(\pm\lambda_{q-1}),$$

where $f(x) = (1+x)(\frac{1}{\Delta-1} - x)^{-1}$.

To simplify slightly the expressions, set $r = (\Delta-1)/\Delta$. The matrix $\varepsilon r\mathbf{I}_{4q}$ shifts the eigenvalues of $\mathbf{Z}$ by εr. Thus, Lemma 54 yields

$$\mathrm{Det}\big(\varepsilon r\mathbf{I}_{4q} + \mathbf{Z}\big) = \varepsilon^3 r^3\big(\varepsilon r + 2f(1)\big)\prod_{i\in[q-1]}\big(\varepsilon r + f(\lambda_i)\big)^2\prod_{i\in[q-1]}\big(\varepsilon r + f(-\lambda_i)\big)^2.$$

By Lemma 54, we have $f(1), f(\pm\lambda_i) \neq 0$ for every $i \in [q-1]$, so that

$$[\varepsilon^3]\mathrm{Det}\big(\varepsilon r\mathbf{I}_{4q} + \mathbf{Z}\big) = 2r^3 f(1)\prod_{i\in[q-1]}\big(f(-\lambda_i)f(\lambda_i)\big)^2 = -\frac{4(\Delta-1)^{4q}}{\Delta^3(\Delta-2)}\prod_{i\in[q-1]}\left(\frac{1-\lambda_i^2}{1-(\Delta-1)^2\lambda_i^2}\right)^2. \tag{144}$$

Plugging (141) and (144) in (142), we obtain

$$[\varepsilon^3]\mathrm{Det}(\mathbf{H}_2') = \frac{4\Delta^{4q-2}}{(\Delta-1)^{2q^2}} \prod_{i,j\in[q-1]} \left(1-(\Delta-1)^2\lambda_i^2\lambda_j^2\right) \prod_{i\in[q-1]} \left(1-\lambda_i^2\right)^2.$$

Using this and (139), we obtain (136).

Proof of Lemma 54. Define

$$\begin{aligned} \mathbf{u}_1 &= \left[\sqrt{\boldsymbol{\alpha}}, \sqrt{\boldsymbol{\alpha}}, \sqrt{\boldsymbol{\beta}}, \sqrt{\boldsymbol{\beta}}\right]^\intercal, & \mathbf{u}_3 &= \left[\sqrt{\boldsymbol{\alpha}}, \mathbf{0}, -\sqrt{\boldsymbol{\beta}}, \mathbf{0}\right]^\intercal, \\ \mathbf{u}_2 &= \left[\sqrt{\boldsymbol{\alpha}}, -\sqrt{\boldsymbol{\alpha}}, \sqrt{\boldsymbol{\beta}}, -\sqrt{\boldsymbol{\beta}}\right]^\intercal, & \mathbf{u}_4 &= \left[\mathbf{0}, \sqrt{\boldsymbol{\alpha}}, \mathbf{0}, -\sqrt{\boldsymbol{\beta}}\right]^\intercal, \end{aligned}$$

and set $S_1 = \{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \mathbf{u}_4\}$ and $S_2 = \{\mathbf{v} : \mathbf{v} \text{ eigenvector of } \mathbf{L}' \text{ with eigenvalue } \mu,\ |\mu| \neq 1\}$. We claim that $S_1 \cup S_2$ is a linearly independent set of eigenvectors of $\mathbf{Z}$.

Let $\mathbf{v} \in S_2$ and let μ be the eigenvalue of $\mathbf{L}'$ associated to $\mathbf{v}$. Then it is easy to check that $\mathbf{V}\mathbf{V}^\intercal\mathbf{v} = \mathbf{v}$, so that $\mathbf{Z}\mathbf{v} = f(\mu)\mathbf{v}$. This gives us $4(q-1)$ eigenvectors of $\mathbf{V}\mathbf{V}^\intercal$ with eigenvalue 1.

Moreover, it is easy to see that

$$[\sqrt{\boldsymbol{\alpha}}, \sqrt{\boldsymbol{\alpha}}, \mathbf{0}, \mathbf{0}]^\intercal, [\mathbf{0}, \mathbf{0}, \sqrt{\boldsymbol{\beta}}, \sqrt{\boldsymbol{\beta}}]^\intercal, [\sqrt{\boldsymbol{\alpha}}, -\sqrt{\boldsymbol{\alpha}}, \mathbf{0}, \mathbf{0}]^\intercal, [\mathbf{0}, \mathbf{0}, \sqrt{\boldsymbol{\beta}}, -\sqrt{\boldsymbol{\beta}}]^\intercal$$

are linearly independent eigenvectors of $\mathbf{V}\mathbf{V}^\intercal$ with eigenvalues $2, 2, 0, 0$ respectively.

Using these or by straightforward calculation, we obtain $\mathbf{V}\mathbf{V}^\intercal\mathbf{u}_1 = 2\mathbf{u}_1$, $\mathbf{V}\mathbf{V}^\intercal\mathbf{u}_2 = 0$, $\mathbf{V}\mathbf{V}^\intercal\mathbf{u}_3 = \mathbf{u}_2$, $\mathbf{V}\mathbf{V}^\intercal\mathbf{u}_4 = \mathbf{u}_2$. Noting that $\mathbf{u}_1, \mathbf{u}_2$ are eigenvectors of $\mathbf{L}'$ with eigenvalues 1 and -1 respectively, we obtain that the vectors in S_1 are eigenvectors of $\mathbf{Z}$ with eigenvalues $2f(1), 0, 0, 0$ respectively. □

Proof of Lemma 52. For notational convenience, set

$$\mathbf{R} = (\oplus_2\boldsymbol{\alpha}^D)^{-1/2}\mathbf{A}_{2q,q^2}\,(\boldsymbol{\gamma}^D)^{1/2}, \quad \mathbf{S} = (\oplus_2\boldsymbol{\beta}^D)^{-1/2}\mathbf{A}_{2q,q^2}(\boldsymbol{\delta}^D)^{1/2}, \quad \mathbf{N} = \mathbf{L}\otimes\mathbf{L}.$$

Note that $\mathbf{V} = \mathbf{R}\oplus\mathbf{S}$. The lemma clearly reduces to proving

$$\mathbf{R}\,\mathbf{N} = \begin{bmatrix} \mathbf{L} & \mathbf{0} \\ \mathbf{0} & \mathbf{L} \end{bmatrix} \mathbf{S}, \qquad \mathbf{S}\,\mathbf{N}^\intercal = \begin{bmatrix} \mathbf{L}^\intercal & \mathbf{0} \\ \mathbf{0} & \mathbf{L}^\intercal \end{bmatrix} \mathbf{R}.$$

We prove the first of these equalities, since the proof of the latter is identical. Let $\mathbf{A}, \mathbf{B}$ be the matrices in the lhs and rhs of the equality. $\mathbf{A}, \mathbf{B}$ clearly have the same dimensions, since $\mathbf{R}, \mathbf{S}$ have dimensions $2q \times q^2$ and $\mathbf{N}$ has dimensions $q^2 \times q^2$. So, it remains to check that the entries of $\mathbf{A}, \mathbf{B}$ are equal. To do this, we first explore the entries of $\mathbf{R}, \mathbf{S}, \mathbf{N}$. It is easy to see that their entries are given by

$$R_{t,(i,k)} = \begin{cases} \sqrt{\frac{\gamma_{ik}}{\alpha_i}}\mathbf{1}\{i = t\}, & 1 \le t \le q \\ \sqrt{\frac{\gamma_{ik}}{\alpha_k}}\mathbf{1}\{k = t - q\}, & q+1 \le t \le 2q \end{cases}, S_{i,(j,l)} = \begin{cases} \sqrt{\frac{\delta_{jl}}{\beta_j}}\mathbf{1}\{j = i\}, & 1 \le i \le q \\ \sqrt{\frac{\delta_{jl}}{\beta_l}}\mathbf{1}\{l = i - q\}, & q+1 \le i \le 2q \end{cases},$$

and $N_{(i,k),(j,l)} = x_{ij}x_{kl}/\sqrt{\gamma_{ik}\delta_{jl}}$. We are now ready to prove that $\mathbf{A} = \mathbf{B}$. Consider the $\big(i, (j,l)\big)$ entries of these matrices. Wlog we may assume $i \le q$. We have

$$A_{i,(j,l)} = \sum_{i',k} R_{i,(i',k)} N_{(i',k),(j,l)} = \sum_{i',k} \sqrt{\frac{\gamma_{i'k}}{\alpha_{i'}}}\mathbf{1}\{i' = i\} \frac{x_{i'j}x_{kl}}{\sqrt{\gamma_{i'k}\delta_{jl}}} = \frac{x_{ij}}{\sqrt{\alpha_i}\sqrt{\delta_{jl}}} \sum_k x_{kl} = \frac{\beta_l x_{ij}}{\sqrt{\alpha_i}\sqrt{\delta_{jl}}}.$$

$$B_{i,(j,l)} = \sum_{j'} L_{i,j'} S_{j',(j,l)} = \sum_{j'} L_{i,j'} \sqrt{\frac{\delta_{jl}}{\beta_j}}\mathbf{1}\{j = j'\} = L_{i,j}\sqrt{\frac{\delta_{jl}}{\beta_j}} = \frac{x_{ij}\sqrt{\delta_{jl}}}{\beta_j\sqrt{\alpha_i}}.$$

Thus $A_{i,(j,l)} = B_{i,(j,l)}$ for every i, j, l and we are done. □

4.4.3.2 Proof of Lemma 50

We first prove (137). Since $\mathbf{A}_1$, $\mathbf{A}_{2,\mathbf{y}}$ are incidence matrices of the bipartite graphs $G_{\mathbf{x}}, G_{\mathbf{y}}$, they are totally unimodular matrices. By the way full dimensional representations are chosen, the matrices $(\mathbf{A}_1)_{fs}$, $(\mathbf{A}_{2,\mathbf{y}})_{fs}$ are invertible and hence their determinants squared equal 1. For $(\mathbf{A}_2)_{fs}$, the same logic applies: by (125), it is immediate to verify that $(\mathbf{A}_2)_{fs}$ has the block decomposition

$$(\mathbf{A}_2)_{fs} = \begin{bmatrix} (\mathbf{A}_{2,\boldsymbol{\gamma\delta}})_{fs} & \mathbf{0} \\ -\mathbf{I} & (\mathbf{A}_{2,\mathbf{y}})_{fs} \end{bmatrix}, \text{ so that } \mathrm{Det}\big((\mathbf{A}_2)_{fs}\big) = \mathrm{Det}\big((\mathbf{A}_{2,\boldsymbol{\gamma\delta}})_{fs}\big)\,\mathrm{Det}\big((\mathbf{A}_{2,\mathbf{y}})_{fs}\big).$$

Since $\mathbf{A}_{2,\boldsymbol{\gamma\delta}}$, $\mathbf{A}_{2,\mathbf{y}}$ are totally unimodular, any invertible submatrix of them has determinant ± 1. This concludes the proof of (137).

We next turn to (138). We begin with $L\big((\mathbf{A}_1)_f, \mathbf{A}_1, \mathbf{T}_1\big)$. The argument is closely related to the proof of Kirchoff's Matrix-Tree Theorem, but is written in a way that it easily extends to the more complicated $L\big((\mathbf{A}_2)_f, \mathbf{A}_2, \mathbf{T}_2\big)$.

Denote by $\mu_1, \dots, \mu_{2q-1}$ the non-zero eigenvalues of $\mathbf{A}_1\mathbf{A}_1^\intercal$; there are exactly $2q-1$ of those since $G_{\mathbf{x}}$ is a connected bipartite graph. Moreover, $\mathbf{v}_0^\intercal = \frac{1}{\sqrt{2q}}[-\mathbf{1}_q \quad \mathbf{1}_q]$ is the unit eigenvector of $\mathbf{A}_1\mathbf{A}_1^\intercal$ with eigenvalue 0. We claim that

$$[\varepsilon]\mathrm{Det}\big(\varepsilon\mathbf{T}_{1,\mathbf{x}} - \mathbf{A}_1\mathbf{A}_1^\intercal\big) = -\frac{\prod_{i\in[2q-1]}\mu_i}{q}, \qquad \mathrm{Det}\big((\mathbf{A}_1)_f\,(\mathbf{A}_1)_f^\intercal\big) = -\frac{\prod_{i\in[2q-1]}\mu_i}{2q}, \tag{145}$$

which clearly proves $L\big((\mathbf{A}_1)_f, \mathbf{A}_1, \mathbf{T}_1\big) = 1/2$. The first equality is a direct application of Lemma 47, after observing that $\mathbf{v}_0^\intercal\mathbf{T}\mathbf{v}_0 = 1/q$. The second can be proved as follows. The matrix $(\mathbf{A}_1)_f\,(\mathbf{A}_1)_f^\intercal$ is a principal minor of $\mathbf{A}_1\mathbf{A}_1^\intercal$, the specific principal minor is clearly determined by which row of $\mathbf{A}_1$ we chose to delete to obtain $(\mathbf{A}_1)_f$. Since $\mathbf{A}_1\mathbf{A}_1^\intercal$ has exactly one zero eigenvalue, we have

$$-\prod_{i\in[2q-1]}\mu_i = \sum_{W\in\binom{[2q]}{2q-1}}\mathrm{Det}\big((\mathbf{A}_1)_W\,(\mathbf{A}_1)_W^\intercal\big), \tag{146}$$

where $(\mathbf{A}_1)_W$ is the submatrix of $\mathbf{A}_1$ induced by the rows indexed with W. It is easily checked that for any $W, W' \in \binom{[2q]}{2q-1}$, there exists a unitary matrix $\mathbf{P}$ such that $(\mathbf{A}_1)_W = \mathbf{P}(\mathbf{A}_1)_{W'}$, so that all summands in (146) are equal. Indeed, since $\mathbf{A}$ corresponds to the incidence matrix of a bipartite graph or otherwise, the sum of the first q rows (as vectors) equals the sum of the last q rows. It follows that any row of $\mathbf{A}_1$ can be expressed as a $\{1,-1\}$ linear combination of the remaining rows, which easily yields the existence of $\mathbf{P}$ with the desired properties. Hence, for any $(\mathbf{A}_1)_f$ as in the statement of the lemma, the second equality in (145) holds as well.

The argument for $L\big((\mathbf{A}_{2,\mathbf{y}})_f, \mathbf{A}_{2,\mathbf{y}}, \mathbf{T}_{2,\mathbf{y}}\big) = 1/2$ is completely analogous. We give a proof sketch for $L\big((\mathbf{A}_2)_f, \mathbf{A}_2, \mathbf{T}_2\big) = 1/4$. The matrix $\mathbf{A}_2\mathbf{A}_2^\intercal$ has zero as an eigenvalue by multiplicity three. Denote by $\sigma_1, \dots, \sigma_{2q^2+4q-3}$ the non-zero eigenvalues of $\mathbf{A}_2\mathbf{A}_2^\intercal$. By looking at the space $\mathbf{z}\mathbf{A}_2 = \mathbf{0}$, it is easy to derive an orthonormal set of eigenvectors for the eigenvalue 0; these are given by the vectors

$$\mathbf{v}_1 = \frac{1}{\sqrt{2q}}[-\mathbf{1}_q,\, \mathbf{1}_q,\, \mathbf{0}_{2q+2q^2}]^\intercal,\ \mathbf{v}_2 = \frac{1}{\sqrt{2q}}[\mathbf{0}_{2q},\, -\mathbf{1}_q,\, \mathbf{1}_q,\, \mathbf{0}_{2q^2}]^\intercal,$$
$$\mathbf{v}_3 = \frac{1}{2\sqrt{q+2q^2}}[-\mathbf{1}_{2q},\, \mathbf{1}_{2q},\, -\mathbf{2}_{q^2},\, \mathbf{2}_{q^2}]^\intercal.$$

Moreover, the analog of (146) is

$$-\prod_{i\in[2q^2+4q-3]}\sigma_i = \sum_{W\in\binom{[2q^2+4q]}{2q^2+4q-3}}\mathrm{Det}\big((\mathbf{A}_2)_W\,(\mathbf{A}_2)_W^\intercal\big). \tag{147}$$

The difference of (147) from (146) is the existence of $W \in \binom{[2q^2+4q]}{2q^2+4q-3}$ which do not contribute to the sum, i.e., $\mathrm{Det}\big((\mathbf{A}_2)_W\,(\mathbf{A}_2)_W^\intercal\big) = 0$. We thus restrict our attention to W with non-zero contribution. These correspond to deleting three rows of $\mathbf{A}_2$ such that either (i) all three rows are among the first $4q$ rows of $\mathbf{A}_2$, (ii) two of the rows are among the first $4q$ rows of $\mathbf{A}_2$ and the remaining is from the last $2q^2$ rows. Note that the deletions must be such that the rank of $\mathbf{A}_2$ is preserved. Using the block decomposition of $\mathbf{A}_2$, c.f. (125), it can easily be checked that there are $4q^3$ W's of the form (i) and $8q^4$ of the form (ii), for a total of $4q^3(2q+1)$.

Moreover, it can be seen that all non-zero summands in (146) are equal. This can be proved using the following linear dependencies among the rows of $\mathbf{A}_2$: (i) the sum of rows $1,\dots,q$ is equal to the sum of rows $q+1,\dots,2q$, (ii) the sum of rows $2q+1,\dots,3q$ is equal to the sum of rows $3q+1,\dots,4q$, (iii) the sum of rows $1,\dots,q-1,q,4q+1,\dots,4q+q^2-1,4q+q^2$ is equal to the sum of rows $2q+1,\dots,3q-1,3q,4q+q^2+1,\dots,4q+2q^2-1,4q+2q^2$.

Hence, the equality $L\big((\mathbf{A}_2)_f,\mathbf{A}_2,\mathbf{T}_2\big) = 1/4$ is obtained by the following analog of (145)

$$[\varepsilon]\mathrm{Det}\big(\varepsilon\mathbf{T}_2 - \mathbf{A}_2\mathbf{A}_2^\intercal\big) = -\frac{\prod_{i\in[2q^2+4q-3]}\sigma_i}{q^3(2q+1)}, \qquad \mathrm{Det}\big((\mathbf{A}_2)_f\,(\mathbf{A}_2)_f^\intercal\big) = -\frac{\prod_{i\in[2q^2+4q-3]}\sigma_i}{4q^3(2q+1)}.$$

CHAPTER V

NP-HARDNESS FOR COUNTING IN SPIN SYSTEMS

In this chapter, we are going to prove the inapproximability results of this thesis, namely Theorems 11, 12, 13, 14, and 15. Much of our focus will concentrate on the more general (and difficult to prove) Theorem 14, which we recall for convenience here.

Theorem 14. *Let $q \geq 2, \Delta \geq 3$. For an antiferromagnetic q-spin system with interaction matrix $\mathbf{B}$, if the dominant semi-translation invariant Gibbs measures on the tree $\mathbb{T}_\Delta$ are permutation-symmetric and all of them are Hessian dominant and not translation invariant then, unless $\mathsf{NP} = \mathsf{RP}$, there is no FPRAS for approximating the partition function for triangle-free Δ-regular graphs. Moreover, there exists $\varepsilon = \varepsilon(q, \mathbf{B}, \Delta) > 0$ such that, unless $\mathsf{NP} = \mathsf{RP}$, one cannot approximate the partition function within a factor $2^{\varepsilon n}$ for triangle-free Δ-regular graphs (where n is the number of vertices).*

5.1 The phase labeling problem and the reduction scheme

5.1.1 Recap: the properties of the gadget

For convenience, we restate the properties of the graph distribution $\mathcal{G}_n^r$ from Chapter 4. To do this, we briefly recall few definitions from Section 4.1, see therein for more details.

A graph $G \sim \mathcal{G}_n^r$ is bipartite, whose bipartition is labeled as $(+, -)$. Moreover, G has vertices of degree Δ and $\Delta - 1$, and W denotes the set of vertices with degree $\Delta - 1$. Further, $\mathcal{Q}$ is the union of dominant phases $\mathbf{p} = (\boldsymbol{\alpha}, \boldsymbol{\beta})$ (cf. Definition 6) and the product measures $\nu_{\mathbf{p}}^{\otimes}(\cdot)$ on configurations σ on G are defined in terms of the corresponding fixpoints of the tree recursions (cf. equation (92)). Finally, the phase $Y(\sigma)$ of a configuration σ on G is the dominant phase which is closer to the footprint of σ on the vertices of degree Δ in G.

We proved the following theorem in Chapter 4.

Theorem 34. *Let $\Delta \geq 3$ and suppose that the interaction matrix $\mathbf{B}$ satisfies conditions (H1), (H2), (H3). Let r be a fixed constant. Then, for every $\varepsilon > 0$, a random graph*

$G \sim \mathcal{G}^r_{n,\Delta}$ satisfies with probability $1-o(1)$ as $n \to \infty$ all of the following:

1. *For each $\mathbf{p} \in \mathcal{Q}$, $(1-\varepsilon)/|\mathcal{Q}| \le \mu_G(Y(\sigma) = \mathbf{p}) \le (1+\varepsilon)/|\mathcal{Q}|$. That is, the phases in the graph G appear with roughly equal probability.*

2. *For each $\mathbf{p} \in \mathcal{Q}$, for all $\eta : W \to [q]$, $\mu_G(\sigma_W = \eta \,|Y(\sigma) = \mathbf{p})/\nu^{\otimes}_{\mathbf{p}}(\eta) \in [1-\varepsilon, 1+\varepsilon]$. That is, conditioned on the phase $\mathbf{p}$ of the configuration, the spins of the vertices in W are roughly independent and the marginal measure on them can be approximated by the measure $\nu^{\otimes}_{\mathbf{p}}(\cdot)$.*

3. *There is no edge between W^+ and W^-. Moreover, there is no vertex in G which has two neighbors in $W^+ \cup W^-$.*

Moreover, G is simple with asymptotically positive probability and the above continue to hold with probability $1-o(1)$ conditioned on G being simple.

Let $\Delta \ge 3$. In this chapter, we will also need the following condition on the spin system specified by the interaction matrix $\mathbf{B}$:

(H4) There does *not* exist a dominant phase which is translation invariant. In other words, if $\mathbf{p} = (\boldsymbol{\alpha}, \boldsymbol{\beta})$ is a dominant phase (cf. Definition 6) for the spin system with interaction matrix $\mathbf{B}$, it holds that $\boldsymbol{\alpha} \neq \boldsymbol{\beta}$.

5.1.2 The phase labeling problem

In this section, we define the phase labeling problem for a general spin system with interaction matrix $\mathbf{B}$ and degree $\Delta \ge 3$. We have already done this in the specific case of the colorings model in Section 1.4.5.2.

In the remainder of the chapter, we are going to identify a dominant phase $\mathbf{p} = (\boldsymbol{\alpha}, \boldsymbol{\beta})$ with the corresponding fixpoint of the tree recursions $(\mathbf{r}, \mathbf{c})$; this is justified by Theorem 5, equation (11). With a slight abuse of notation, we will write $\mathbf{p} = (\mathbf{r}, \mathbf{c})$ instead of $(\boldsymbol{\alpha}, \boldsymbol{\beta})$; no confusion should arise, we will not need an explicit handle on $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ in this chapter.

From hereon, we are going to assume that we have a list of the dominant phases. For fixed $q, \Delta, \mathbf{B}$, they can be found using existential theory of reals and can be approximated to any desired polynomial accuracy of their values. Let $\mathcal{Q}$ be the list of dominant phases–we will call them ordered phases and view them as ordered pairs of vectors:

$$\mathcal{Q} = \{(\mathbf{r}_1, \mathbf{c}_1), \ldots, (\mathbf{r}_Q, \mathbf{c}_Q)\}.$$

The condition (H4) translates into $\mathbf{r}_i \neq \mathbf{c}_i$ for all $i \in [Q]$. Note that if $(\mathbf{r}, \mathbf{c}) \in \mathcal{Q}$, then by symmetry $(\mathbf{c}, \mathbf{r}) \in \mathcal{Q}$ as well. We will refer to the unordered pair of vectors $\{\mathbf{r}, \mathbf{c}\}$ as an unordered phase. Let $\mathcal{Q}'$ be the set of unordered phases over all dominant phases; namely, $\mathcal{Q}' = \{\{\mathbf{r}, \mathbf{c}\} \mid \mathbf{p} = (\mathbf{r}, \mathbf{c}) \in \mathcal{Q}\}$. Enumerate $\mathcal{Q}'$ as follows:

$$\mathcal{Q}' = \{\{\mathbf{r}_1, \mathbf{c}_1\}, \ldots, \{\mathbf{r}_{Q'}, \mathbf{c}_{Q'}\}\}.$$

Note that $Q = 2Q'$ (as a consequence of condition (H4)).

We will denote unordered phases using $\mathbf{p}$; the two ordered phases corresponding to the unordered phase $\mathbf{p}$ will be denoted by $\mathbf{p}^+$ and $\mathbf{p}^-$. Given a graph H with vertex set V we will assign ordered phases to its vertices—the labeling (also called phase assignment) will be denoted by $\mathcal{Y} : V \to \mathcal{Q}$. The corresponding labeling by unordered phases (where the ordering is removed) will be denoted by $\mathcal{Y}'$.

Now we define the weight of a phase assignment. We will have two types of edges in H: parallel or symmetric; the type of an edge will only impact the weight of a phase assignment. In particular, a parallel edge whose endpoints have labels $(\mathbf{r}_1, \mathbf{c}_1)$ and $(\mathbf{r}_2, \mathbf{c}_2)$ incurs weight

$$w_p((\mathbf{r}_1, \mathbf{c}_1), (\mathbf{r}_2, \mathbf{c}_2)) = \ln(\mathbf{r}_1^\intercal \mathbf{B} \mathbf{r}_2) + \ln(\mathbf{c}_1^\intercal \mathbf{B} \mathbf{c}_2),$$

while a symmetric edge incurs weight

$$w_s((\mathbf{r}_1, \mathbf{c}_1), (\mathbf{r}_2, \mathbf{c}_2)) = w_p((\mathbf{r}_1, \mathbf{c}_1), (\mathbf{r}_2, \mathbf{c}_2)) + w_p((\mathbf{r}_1, \mathbf{c}_1), (\mathbf{c}_2, \mathbf{r}_2)).$$

Note that if we flip $(\mathbf{r}_1, \mathbf{c}_1)$, that is, replace it by $(\mathbf{c}_1, \mathbf{r}_1)$, the weight of the symmetric edge does not change.

We will use the following problem in our reduction.

Phase Labeling Problem($\mathbf{B}, \mathcal{Q}$):

INPUT: undirected edge-weighted multigraph $H = (V, E)$ and a partition of the edges $\{E_p, E_s\}$.

OUTPUT: $\textsc{MaxLwt}(H) := \max_{\mathcal{Y}} \textsc{Lwt}_H(\mathcal{Y})$, where the maximization is over all possible phase labelings $\mathcal{Y} : V \to \mathcal{Q}$ and

$$\textsc{Lwt}_H(\mathcal{Y}) = \sum_{\{u,v\}\in E_s} w_s(\mathcal{Y}(u), \mathcal{Y}(v)) + \sum_{\{u,v\}\in E_p} w_p(\mathcal{Y}(u), \mathcal{Y}(v)).$$

The motivation for the Phase Labeling problem is the following lemma. The proof roughly follows the lines of [64] and is given in Section 5.2.

Lemma 55. *Let $\Delta \geq 3$ and $\mathbf{B}$ specify the interaction matrix of a spin system. Assume further that conditions (H1), (H2), (H3) are satisfied. Then, a (randomized) algorithm that approximates the partition function on triangle free Δ-regular graphs within an arbitrarily small exponential factor yields a* PRAS *for the phase labeling problem with parameters $\mathbf{B}, \mathcal{Q}$ on bounded degree graphs.*

The most difficult part of our arguments is the following lemma which is given in Section 5.5.

Lemma 56. *Let $\Delta \geq 3$ and $\mathbf{B}$ specify the interaction matrix of a spin system. Assume further that condition (H4) is satisfied. A* PRAS *for the phase labeling problem with parameters $\mathbf{B}, \mathcal{Q}$ on bounded degree graphs yields a* PRAS *for* MaxCut *on 3-regular graphs.*

Using Lemmas 55 and 56, we obtain Theorem 14.

Proof of Theorem 14. Suppose that there exists a (randomized) algorithm to approximate the partition function on Δ-regular graphs with interaction matrix $\mathbf{B}$ up to an arbitrarily small exponential factor. Then, combinining Lemmas 55 and 56, we obtain a (randomized) algorithm to approximate MaxCut on 3-regular graphs within a factor of $1 - o(1)$. This contradicts the result of [3]. □

5.2 Connection between approximating the partition function and the phase labeling problem

The purpose of this section is to prove Lemma 55.

We begin by clarifying what degree means in our setting. Let $H = (V, E)$ be an instance of the phase labeling problem, and $\{E_p, E_s\}$ be a partition of the edges and $|V| = m$. The degree of a vertex $v \in V$ will be defined as $2d_s + d_p + 4l_s + 2l_p$, where d_s, d_p are the numbers of symmetric and parallel edges joining v to a distinct vertex u and l_s, l_p are the numbers of symmetric and parallel loops from v to itself. The bounded degree assumption means there is an absolute constant D (not depending on m) which bounds the degree of any $v \in V$.

To approximate the phase labeling problem on H with parameters $\mathbf{B}, \mathcal{Q}$, we will replace each vertex in the graph H by a suitable graph in a family of gadgets $\mathcal{F}$. The construction has a parameter k which roughly controls the accuracy of the approximation we want to achieve. The family $\mathcal{F}$ will be of the form $\{G^d\}_{d\in[D]}$ and the gadget for a vertex v will be G^d where d is the degree of v. Note that the cardinality of $\mathcal{F}$ is bounded by the absolute constant D. The gadgets G^d are selected from the graph distribution $\mathcal{G}_n^{kd}$ for some n which is sufficiently large.

An immediate consequence of Theorem 34 is the following.

Corollary 57. *Let k be an arbitrarily large constant. For $d \in [D]$, let $G^d \sim \mathcal{G}_n^{kd}$ and set $\mathcal{F} = \{G^d\}_{d\in[D]}$. Then, for all sufficiently large n, G^d is simple and satisfies Items 1, 2 and 3 of Theorem 34 with positive probability for every $d \in [D]$.*

Corollary 57 also yields a trivial randomized algorithm to construct the family $\mathcal{F}$ for an arbitrary constant k. In fact, since all the parameters are constants, one can construct the family $\mathcal{F}$ by brute force search. With the family $\mathcal{F}$ in our hands, we can now give the details of the construction.

The first step consists of replacing each vertex $v \in H$ with degree d with a distinct copy of the gadget $G^d \in \mathcal{F}$. We will refer to the gadget corresponding to vertex v as G_v and the respective sets in G_v as $W_v, W_v^{\pm}, U_v^{\pm}$. Denote by $\widehat{H}$ the graph obtained by the disconnected copies of the gadgets.

The second step consists of placing the edges of H in $\widehat{H}$, that is, making connections between the gadgets. The final graph will be denoted as $H_{\mathcal{F}}$. The edges we are going to place will form a perfect matching on $\cup_{v\in H} W_v$ and as a result $H_{\mathcal{F}}$ will be Δ-regular. Every parallel edge of H corresponds to $2k$ edges in $H_{\mathcal{F}}$, while every symmetric to $4k$. Roughly, parallel and symmetric indicate which parts of two gadgets get connected (recall that the gadgets are bipartite). Loops are treated as if they were connecting distinct vertices.

In detail, let (u, v) be an edge e of H. Suppose first that $u \neq v$. If e is parallel, place k edges between W_u^s and W_v^s for $s \in \{+,-\}$. If e is symmetric, place k edges between W_u^s and W_v^s and k edges between W_u^s and W_v^{-s} for $s \in \{+,-\}$. Suppose now that $u = v$. If e is parallel, place k edges between distinct vertices in W_v^+ and k edges between distinct vertices in W_v^-. If e is symmetric, place $2k$ edges between W_v^+ and W_v^-, k edges between distinct vertices in W_v^+ and k edges between distinct vertices in W_v^-.

The first step of the construction guarantees that the second step can be done in a (deterministic) way so that $H_{\mathcal{F}}$ is Δ-regular. Moreover, by Corollary 57 and item 3 of Lemma 34, $H_{\mathcal{F}}$ is a simple, triangle-free graph.

Proof of Lemma 55. The argument in [64, Lemma 4.3] almost verbatim gives

$$\frac{(1-\varepsilon)^{2m}}{|\mathcal{Q}|^m} \leq \frac{Z_{H_{\mathcal{F}}}/Z_{\widehat{H}}}{\exp(k \cdot \textsc{MaxLwt}(H))} \leq (1+\varepsilon)^m.$$

This can be rearranged into

$$\frac{1}{k}\log\left(\frac{Z_{H_{\mathcal{F}}}}{Z_{\widehat{H}}}\right) - \frac{m}{k}\log(1+\varepsilon) \leq \textsc{MaxLwt}(H) \leq \frac{1}{k}\log\left(\frac{Z_{H_{\mathcal{F}}}}{Z_{\widehat{H}}}\right) - \frac{m}{k}\big[2\log(1-\varepsilon) - \log|\mathcal{Q}|\big].$$

The argument in [64, Proof of Theorems 1 and 2] gives the desired result. We give the short details. The graph $\widehat{H}$ consists of m disconnected subgraphs, each of constant size. Hence, we can compute $Z_{\widehat{H}}$ exactly in polynomial time. Assume now that $Z_{H_{\mathcal{F}}}$ can be approximated within a factor of $\exp\left(c|\widehat{H}|\right)$ in polynomial time for any $c > 0$. Since $\log\left(Z_{H_{\mathcal{F}}}\right)$ is bounded above by $O(|\widehat{H}|)$, the ratio $\log\left(Z_{H_{\mathcal{F}}}/Z_{\widehat{H}}\right)$ can be approximated within an additive $O(c|\widehat{H}|) = O[cm(n+kD)]$. Thus we obtain upper and lower bounds for $\textsc{MaxLwt}(H)$ which differ by $O[(cn+1)m/k]$. A random phase labeling yields the lower bound $\textsc{MaxLwt}(H) \geq \Omega(m)$. Thus, the final approximation is within a multiplicative factor $1 + O[(cn+1)/k]$ of $\textsc{MaxLwt}(H)$. To make this multiplicative factor arbitrarily close to 1, first take k large and make cn small with a sufficiently small value of c. $\square$

5.3 Reducing MaxCut to the phase labeling problem for spin systems with a unique dominant phase

Proof of Theorem 15. The assumptions of the Theorem imply that $\mathcal{Q} = \{\mathbf{p}^+, \mathbf{p}^-\}$, where $\mathbf{p} = \{\mathbf{r}, \mathbf{c}\}$. Using Lemma 55, the hardness will follow from the APX-hardness of MaxCut [3] and

$$w_p(\mathbf{p}^+, \mathbf{p}^+) = w_p(\mathbf{p}^-, \mathbf{p}^-) < w_p(\mathbf{p}^+, \mathbf{p}^-).$$

The equality follows from the definition of $w_p(\cdot, \cdot)$ and $\mathbf{p}^+, \mathbf{p}^-$, while the inequality is equivalent to

$$(\mathbf{r}^\intercal \mathbf{B} \mathbf{r})(\mathbf{c}^\intercal \mathbf{B} \mathbf{c}) \leq (\mathbf{r}^\intercal \mathbf{B} \mathbf{c})^2.$$

To see the inequality, recall that $(\mathbf{r}, \mathbf{c})$ are maximizers of $\Phi(\mathbf{r}, \mathbf{c}) = \dfrac{\mathbf{r}^\intercal \mathbf{B} \mathbf{c}}{\|\mathbf{r}\|_p \|\mathbf{c}\|_p}$, where $p = \frac{\Delta}{\Delta - 1}$. By cancelling the denominators in $\Phi(\mathbf{r}, \mathbf{r})\Phi(\mathbf{c}, \mathbf{c}) \leq (\Phi(\mathbf{r}, \mathbf{c}))^2$, we obtain the desired inequality, as wanted. □

5.4 Reducing MaxCut to the colorings phase labeling problem

In this section, as a preparation to the proof of the more general Lemma 56, we first prove the following lemma.

Lemma 58. *A* PRAS *for* Colorings Phase Labeling Problem($\mathcal{Q}$) *yields a* PRAS *for* Max-Cut *on 3-regular graphs.*

The reduction in the proof of Lemma 58 relies on the following gadget which "prefers" the unordered phase of two distinguished vertices u and v to agree. Recall that for a phase assignment $\mathcal{Y}$ with ordered phases, we denote by $\mathcal{Y}'$ the respective phase assignment with unordered phases.

Lemma 59. *A constant sized gadget J_1 with two distinguished vertices u, v can be constructed with the following property: all edges of J_1 are symmetric and the following is true,*

$$\max_{\mathcal{Y};\, \mathcal{Y}'(u) = \mathcal{Y}'(v)} \textsc{Lwt}_{J_1}(\mathcal{Y}) > \varepsilon_1 + \max_{\mathcal{Y};\, \mathcal{Y}'(u) \neq \mathcal{Y}'(v)} \textsc{Lwt}_{J_1}(\mathcal{Y}), \tag{148}$$

where $\varepsilon_1 > 0$ is a constant depending only on k and Δ.

With Lemma 59 at hand, we can derive Lemma 58 fairly easily.

Proof of Lemma 58. Let ε_1 be as in Lemma 59 and

$$t := 2\left\lceil \frac{\max_{\mathbf{p}_1,\mathbf{p}_2} w_p(\mathbf{p}_1,\mathbf{p}_2) - \min_{\mathbf{p}_1,\mathbf{p}_2} w_p(\mathbf{p}_1,\mathbf{p}_2)}{\varepsilon_1}\right\rceil.$$

Given a 3-regular instance $H = (V, E)$ of Max-Cut, we first declare all edges of H to be parallel. Moreover, for every edge (u', v') of H, take t copies of gadget J_1 from Lemma 59, identify (merge) their u vertices with u', and identify (merge) their v vertices with v'. Let H' be the final graph.

To find the optimal phase labeling of H', we may focus on the phase assignment restricted to vertices in H, since each gadget J_1 can be independently set to its optimal value conditioned on the phases for its distinguished vertices u and v. We claim that

$$\textsc{MaxLwt}(H') = C_1\textsc{MaxCut}(H) + (C_2 + C_3 t)|E|, \tag{149}$$

for constants C_1, C_2, C_3 to be specified later (depending only on k, Δ). Using the trivial bound $\textsc{MaxCut}(H) \geq |E|/2 = 3|V|/4$, the lemma follows easily from (149). We thus focus on proving (149).

The key idea is that for any phase labeling $\mathcal{Y} : V \to \mathcal{Q}$, changing the unordered phases of vertices in H to the same unordered phase $\mathbf{p} \in \mathcal{Q}'$, while keeping the spins, can only increase the weight of the labeling. Indeed, for $(u, v) \in E$ such that $\mathcal{Y}'(u) = \mathcal{Y}'(v)$, no change in the weight of the labeling occurs, using (148). For $(u, v) \in E$ such that $\mathcal{Y}'(u) \neq \mathcal{Y}'(v)$, the potential (weight) loss from the parallel edge (u, v) is compensated by the gain on the t copies of J_1 by (148) and the choice of t.

For phase labelings which assign vertices of H the same unordered phase $\mathbf{p}$, to attain the maximum weight for a phase labeling, we only need to choose the spins, in order to maximize the contribution from parallel edges (the edges of H). The same argument we discussed for the hard-core model, (19) yields that the optimal choice of spins to the phases induces a maximum-cut partition of H. For such a spin assignment, the contribution from parallel edges is $C_1\textsc{MaxCut}(H) + C_2|E|$, where

$$C_1 := w_p(\mathbf{p}^+, \mathbf{p}^-) - w_p(\mathbf{p}^-, \mathbf{p}^-) \text{ and } C_2 := w_p(\mathbf{p}^-, \mathbf{p}^-).$$

Finally, if we let $C_3 := \max_{\mathcal{Y};\, \mathcal{Y}'(u)=\mathcal{Y}'(v)=\mathbf{p}} \textsc{Lwt}_{J_1}(\mathcal{Y})$, it is simple to see that the contribution from symmetric edges is $C_3 t|E|$. This proves (149). $\square$

We next give the proof of the critical Lemma 59.

Proof of Lemma 59. Let $\mathcal{Q}' := \{\mathbf{p}_1, \dots, \mathbf{p}_{Q'}\}$ and $\mathbf{p}_i := \{\mathbf{r}_i, \mathbf{c}_i\}$ for $i \in [Q']$. Denote by K the multigraph on Q' vertices $b_1, b_2, \dots, b_{Q'}$ with the following symmetric edges: self-loop on b_i for $i \in [Q']$ and two edges between b_i and b_j for every $i, j \in [Q']$ with $i \neq j$. We first prove that the optimal phase assignments $\mathcal{Y}$ of K are those which assign each vertex b_i a distinct phase from $\mathcal{Q}'$ (note that the spin of the phase does not matter since all edges of K are symmetric). The desired gadget J_1 will be constructed afterwards.

Let $\mathcal{Y}$ be a phase labeling of K and s_i be the number of vertices assigned phase $\mathbf{p}_i$. Let $\mathbf{s}$ be the vector $(s_1, \dots, s_{Q'})^\intercal$. Note that $\mathbf{1}^\intercal \mathbf{s} = Q'$, where $\mathbf{1}$ is the all one vector with dimension Q'. Then

$$\textsc{Lwt}_K(\mathcal{Y}) = \sum_{i,j \in [Q']} s_i s_j w_s(\mathbf{p}_i, \mathbf{p}_j) = \mathbf{s}^\intercal \mathbf{A} \mathbf{s},$$

where $\mathbf{A}$ is the $Q' \times Q'$ matrix whose (i,j) entry equals $w_s(\mathbf{p}_i, \mathbf{p}_j)$. Note that $\mathbf{A}$ is symmetric and $\mathbf{1}$ is an eigenvector of $\mathbf{A}$ (because of the transitive symmetry of phases). Moreover, if we let $\mathbf{s}' = \mathbf{s} - \mathbf{1}$, then $\mathbf{1}^\intercal \mathbf{s}' = 0$. It follows that

$$\mathbf{s}^\intercal \mathbf{A} \mathbf{s} = \mathbf{1}^\intercal \mathbf{A} \mathbf{1} + (\mathbf{s}')^\intercal \mathbf{A} \mathbf{s}'. \tag{150}$$

If $\mathbf{A}$ is negative definite, equation (150) shows that the all ones labeling is better than any other labeling. Hence the result will follow if we prove that $\mathbf{A}$ is negative definite.

Let $\mathbf{z}_1, \dots, \mathbf{z}_Q := \mathbf{r}_1, \dots, \mathbf{r}_{Q'}, \mathbf{c}_1, \dots, \mathbf{c}_{Q'}$ and let $\hat{\mathbf{A}}$ be the $Q \times Q$ matrix whose ij-entry is $\ln(1 - \mathbf{z}_i^\intercal \mathbf{z}_j)$. Using the definition of the weights $w_s(\cdot,\cdot)$, it is easy to check that for any vector $\mathbf{s}$ it holds that

$$\mathbf{s}^\intercal \mathbf{A} \mathbf{s} = (\mathbf{s}, \mathbf{s})^\intercal \hat{\mathbf{A}} (\mathbf{s}, \mathbf{s}),$$

so it suffices to prove that $\hat{\mathbf{A}}$ is negative definite. We will show here that $\hat{\mathbf{A}}$ is negative semi-definite; the proof that $\hat{\mathbf{A}}$ is regular (and hence negative definite) is trickier and is given in the proof of the more general Lemma 62. Note that the entries of $\hat{\mathbf{A}}$ are obtained by applying $z \mapsto \ln(1-z)$ to each entry of the Gram matrix of the vectors $\mathbf{z}_1, \dots, \mathbf{z}_Q$. Since for $|z| < 1$ we have $\ln(1-z) = -z - z^2/2 - z^3/3 - \dots$, by Schur's product theorem (see Corollary 7.5.9 in [33]) we obtain that $\hat{\mathbf{A}}$ is negative semi-definite, as desired.

To construct the gadget J_1, we overlay two copies of K as follows. Let K_u (resp. K_v) be a copy of K, where the image of $b_{Q'}$ is renamed to u (resp. v). Overlay K_u, K_v by identifying the images of $b_1, \dots, b_{Q'-1}$ in the two copies. Thus, the resulting graph J_1 has two self loops on b_i for $i \in [Q'-1]$, four edges between b_i and b_j for every $i, j \in [Q'-1]$ with $i \neq j$, two edges between u and b_i for $i \in [Q'-1]$, two edges between v and b_i for $i \in [Q'-1]$ and a self loop on u, v.

Note that for every phase labeling $\mathcal{Y}$ of J_1, we have

$$\textsc{Lwt}_{J_1}(\mathcal{Y}) = \textsc{Lwt}_{K_u}(\mathcal{Y}) + \textsc{Lwt}_{K_v}(\mathcal{Y})$$

and hence

$$\textsc{MaxLwt}(J_1) \leq 2\textsc{MaxLwt}(K).$$

Using that the optimal phase labelings for K are those which assign each vertex a distinct phase from $\mathcal{Q}'$, we obtain that the inequality holds at equality for those (and only those) phase labelings which assign u, v a common phase $\mathbf{p} \in \mathcal{Q}'$ and vertices $b_1, \dots, b_{Q'-1}$ a distinct phase from $\mathcal{Q}' - \{\mathbf{p}\}$. This yields the ε_1 in the statement of the lemma. Note that ε_1 depends only on $\mathcal{Q}'$, which in turn is completely determined by k, Δ. $\square$

5.5 *Reducing MaxCut to the phase labeling problem for general antiferromagnetic spin systems*

5.5.1 Antiferromagnetic spin systems and their properties

5.5.1.1 Definition of antiferromagnetic spin systems

In this section, we define antiferromagnetic models in terms of the eigenvalues of the interaction matrix $\mathbf{B}$. Using this definition, we then state and prove properties of antiferromagnetic models relevant to this work.

The interaction matrix $\mathbf{B}$ is assumed to be symmetric and have non-negative entries. These are standard assumptions since we are interested in undirected graphs and the Gibbs distribution should be a probability distribution. We will also assume that $\mathbf{B}$ is primitive, i.e., irreducible and aperiodic. As we demonstrate next, this is also a natural restriction which does not affect generality.

If $\mathbf{B}$ is reducible, by a suitable permutation of the labels of colors, $\mathbf{B}$ can be put into block diagonal form (which coincides with the normal form of the reducible $\mathbf{B}$) where each of the blocks is either irreducible or zero. Intuitively, this says that the original spin model can be studied by considering the induced sub-models of each block which correspond to irreducible symmetric matrices (where our results apply). Indeed, these sub-models are "non-interacting" for connected graphs G, that is, the partition function for the original model is simply the sum of the partition functions of each sub-model.

We are now ready to give the definition of antiferromagnetism we use.

Definition 11. *Let* $\mathbf{B}$ *be the interaction matrix of a* q*-state spin system. Since* $\mathbf{B}$ *is symmetric all of its eigenvalues are real. Also note that it has non-negative entries and by irreducibility, the Perron-Frobenius theorem implies that one of the eigenvalues of* $\mathbf{B}$ *with the largest magnitude is positive and simple, i.e., the associated eigenspace is one-dimensional. The model is called* antiferromagnetic *if all the other eigenvalues are negative. Note that no eigenvalue is allowed to be zero and hence* $\mathbf{B}$ *is regular.*

The above definition generalizes antiferromagnetism for 2-spin systems, and captures colorings as well as the antiferromagnetic region for the Potts models. Moreover, the above definition seems natural in that it implies that neighboring vertices prefer to have different spin assignments (see Corollary 61 below). Another nice feature of Definition 11 is that it does not depend on the presence of external fields. Specifically, for Δ-regular graphs, any external field can be pushed into the interaction matrix $\mathbf{B}$ with a congruence transformation of the matrix $\mathbf{B}$. The resulting interaction matrix, by Sylvester's law of inertia, has the same number of positive, zero and negative eigenvalues and in particular remains antiferromagnetic.

We conclude this discussion by pointing out that some of our results for general models are more easily stated when $\mathbf{B}$ is further assumed to be aperiodic. We shall refer to such matrices $\mathbf{B}$ (irreducible and aperiodic) as ergodic. Note that if $\mathbf{B}$ is periodic, its period must be two, since $\mathbf{B}$ is symmetric. Such a model is only interesting on bipartite graphs (otherwise the partition function is zero). It can be verified that Definition 11 of antiferromagnetism implies that the matrix $\mathbf{B}$ is ergodic whenever $q \geq 3$ (note that it is trivial to compute the partition function on periodic models with $q = 2$).

5.5.1.2 Properties

As a consequence of the Perron-Frobenius theorem and the antiferromagnetism definition, we may decompose the interaction matrix $\mathbf{B}$ of an antiferromagnetic model as

$$\mathbf{B} = \mathbf{u}\mathbf{u}^{\intercal} - \mathbf{P}^{\intercal}\mathbf{P}, \tag{151}$$

where the vector $\mathbf{u}$ has positive entries and $\mathbf{P}$ is a square matrix. Using the decomposition (151), we prove the following two lemmas which are used in the reduction.

Lemma 60. *For antiferromagnetic* $\mathbf{B}$*, and vectors* $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^q_{\geq 0}$ *with* $\|\mathbf{z}_1\|_1 = \|\mathbf{z}_2\|_1 = 1$*, we have*

$$(\mathbf{z}_1^{\intercal}\mathbf{B}\mathbf{z}_1)(\mathbf{z}_2^{\intercal}\mathbf{B}\mathbf{z}_2) \leq (\mathbf{z}_1^{\intercal}\mathbf{B}\mathbf{z}_2)^2.$$

Equality holds iff $\mathbf{z}_1 = \mathbf{z}_2$*.*

Proof. Set $\mathbf{w}_1 = \mathbf{P}\mathbf{z}_1$, $\mathbf{w}_2 = \mathbf{P}\mathbf{z}_2$, $a_1 = \mathbf{u}^{\intercal}\mathbf{z}_1$, $a_2 = \mathbf{u}^{\intercal}\mathbf{z}_2$. Then

$$\mathbf{z}_1^{\intercal}\mathbf{B}\mathbf{z}_1 = a_1^2 - \mathbf{w}_1^{\intercal}\mathbf{w}_1,\ \mathbf{z}_2^{\intercal}\mathbf{B}\mathbf{z}_2 = a_2^2 - \mathbf{w}_2^{\intercal}\mathbf{w}_2,\ \mathbf{z}_1^{\intercal}\mathbf{B}\mathbf{z}_2 = a_1 a_2 - \mathbf{w}_1^{\intercal}\mathbf{w}_2.$$

Since $\mathbf{B}, \mathbf{z}_1, \mathbf{z}_2$ have non-negative entries, the above equalities imply that

$$a_1^2 - \mathbf{w}_1^{\intercal}\mathbf{w}_1, a_2^2 - \mathbf{w}_2^{\intercal}\mathbf{w}_2, a_1 a_2 - \mathbf{w}_1^{\intercal}\mathbf{w}_2 \geq 0.$$

The desired inequality reduces to

$$\left(a_1^2 - \mathbf{w}_1^{\intercal}\mathbf{w}_1\right)\left(a_2^2 - \mathbf{w}_2^{\intercal}\mathbf{w}_2\right) \leq \left(a_1 a_2 - \mathbf{w}_1^{\intercal}\mathbf{w}_2\right)^2.$$

This is known as Aczél's inequality. The fastest proof goes as follows: set $b_1^2 = a_1^2 - \mathbf{w}_1^{\intercal}\mathbf{w}_1$ and $b_2^2 = a_2^2 - \mathbf{w}_2^{\intercal}\mathbf{w}_2$, so that by Cauchy-Schwarz $a_1 a_2 \geq b_1 b_2 + \mathbf{w}_1^{\intercal}\mathbf{w}_2$, implying the inequality.

Equality can only hold if $a_1 = \lambda a_2$ and $\mathbf{w}_1 = \lambda \mathbf{w}_2$, yielding $\mathbf{u}^{\intercal}(\mathbf{z}_1 - \lambda\mathbf{z}_2) = 0$ and $\mathbf{P}(\mathbf{z}_1 - \lambda\mathbf{z}_2) = 0$. We easily obtain $\mathbf{B}(\mathbf{z}_1 - \lambda\mathbf{z}_2) = 0$ and since $\mathbf{B}$ is invertible, $\mathbf{z}_1 = \lambda\mathbf{z}_2$. The assumption $\|\mathbf{z}_1\|_1 = \|\mathbf{z}_2\|_1 = 1$ implies $\lambda = 1$, as wanted. $\square$

Corollary 61. *By plugging in the inequality of Lemma 60 the vectors with a single 1 in the positions i and j respectively, we obtain that any two spins i, j induce an antiferromagnetic two-spin system.*

Lemma 62. *Let $\mathbf{z}_1, \ldots, \mathbf{z}_n \in \mathbb{R}^d$ be a collection of distinct non-negative vectors such that $\|\mathbf{z}_i\|_1 = 1$ for $i \in [n]$. Let $a_i = \mathbf{z}_i^\intercal \mathbf{u}$, where $\mathbf{u}$ is as in (151). Let $\mathbf{A}'$ be the $n \times n$ matrix whose ij-th entry is $\ln(\mathbf{z}_i^\intercal \mathbf{B} \mathbf{z}_j) - \ln(a_i) - \ln(a_j)$. Then $\mathbf{A}'$ is negative definite.*

Proof. Let $\mathbf{w}_i = \frac{1}{a_i} \mathbf{P} \mathbf{z}_i$ and let $\mathbf{W}$ be the $q \times n$ matrix whose columns are $\mathbf{w}_1, \ldots, \mathbf{w}_n$. We first argue $\mathbf{w}_i \neq \mathbf{w}_j$ for $i \neq j$. Suppose $\mathbf{w}_i = \mathbf{w}_j$. Let $\mathbf{z} = \frac{1}{a_i}\mathbf{z}_i - \frac{1}{a_j}\mathbf{z}_j$. We have $\mathbf{P}\mathbf{z} = \mathbf{w}_i - \mathbf{w}_j = 0$ and $\mathbf{u}^\intercal \mathbf{z} = 1 - 1 = 0$ and hence $\mathbf{B}\mathbf{z} = 0$. Since $\mathbf{B}$ is regular we have $\mathbf{z} = 0$. Thus $0 = \mathbf{z}^\intercal \mathbf{1} = \frac{1}{a_i} - \frac{1}{a_j}$ which implies $a_i = a_j$ which in turn implies $\mathbf{z}_i = \mathbf{z}_j$, a contradiction. Thus $\mathbf{w}_i \neq \mathbf{w}_j$ for $i \neq j$.

Note that we have

$$\ln(1 - \mathbf{w}_i^\intercal \mathbf{w}_j) = \ln(a_i a_j - \mathbf{z}_i^\intercal \mathbf{P}^\intercal \mathbf{P} \mathbf{z}_j) - \ln(a_i a_j) = A'_{ij}.$$

Thus the ij-th entry in $\mathbf{A}'$ is obtained by applying $z \mapsto \ln(1-z)$ to each entry of the Gramm matrix $\mathbf{W}^\intercal \mathbf{W}$. Note that for $|z| < 1$ we have $\ln(1-z) = -z - z^2/2 - z^3/3 - \ldots$ and hence by Schur product theorem $\mathbf{A}'$ is negative semi-definite (see Corollary 7.5.9 in [33]).

Now we argue that $\mathbf{A}'$ is regular (and hence negative definite). We have

$$-\mathbf{A}' = \sum_{k=1}^{\infty} \frac{1}{k} \mathbf{W}_k^\intercal \mathbf{W}_k, \tag{152}$$

where $\mathbf{W}_k$ is the $q^k \times n$ matrix whose columns are $w_1^{\otimes k}, \ldots, w_n^{\otimes k}$. Note that if $\mathbf{A}'$ is singular then there exists a non-zero vector v such that $\mathbf{v}^\intercal \mathbf{A}' \mathbf{v} = 0$ and for this to happen we would have to have

$$\mathbf{W}_k \mathbf{v} = 0 \tag{153}$$

for all $k \geq 1$ (the terms on the right-hand side of (152) are non-negative and if even one of them is positive then $\mathbf{v}^\intercal \mathbf{A}' \mathbf{v} < 0$).

There exists a vector $\mathbf{h} \in \mathbb{R}^q$ such that $\alpha_i = \mathbf{h}^\intercal \mathbf{w}_i, i = 1, \ldots, n$ are distinct real numbers (the $\mathbf{w}_i$'s are distinct and hence for any $i \neq j$ the measure of $\mathbf{h} \in [0,1]^q$ such that $\mathbf{h}^\intercal \mathbf{w}_i =$

$\mathbf{h}^\intercal \mathbf{w}_j$ is zero). Note that $(\mathbf{h}^{\otimes k})^\intercal \mathbf{W}_k$ is $(\alpha_1^k, \dots, \alpha_n^k)$. From (153) we obtain that for every integer $k \geq 1$ we have $(\alpha_1^k, \dots, \alpha_n^k)\mathbf{v} = 0$ and hence $\mathbf{v} = 0$ (by considering the Vandermonde matrix $\{\alpha_i^k\}$), a contradiction. Hence $\mathbf{A}'$ is regular and negative definite. □

5.5.2 The reduction

The remainder of this section is devoted to the proof of Lemma 56. We will use the existence of the following gadget which "prefers" the unordered phase of two vertices to agree. We postpone the proof to Section 5.5.3.

Lemma 63. *A constant sized gadget J_1 with two special vertices u, v can be constructed with the following property: all edges of J_1 are symmetric and the following is true,*

$$\max_{\mathcal{Y};\, \mathcal{Y}'(u)=\mathcal{Y}'(v)} \textsc{Lwt}_{J_1}(\mathcal{Y}) > \varepsilon_1 + \max_{\mathcal{Y};\, \mathcal{Y}'(u)\neq\mathcal{Y}'(v)} \textsc{Lwt}_{J_1}(\mathcal{Y}), \tag{154}$$

where $\varepsilon_1 > 0$ is a constant depending only on the spin model and Δ.

In Lemma 60 we proved that for a parallel edge and any phase $\mathbf{p}$ we have $w(\mathbf{p}^+, \mathbf{p}^+) = w_p(\mathbf{p}^-, \mathbf{p}^-) < w_p(\mathbf{p}^+, \mathbf{p}^-)$ and hence there exists a constant $\varepsilon_2 > 0$ depending only on the model and Δ such that for every phase $\mathbf{p} \in \mathcal{Q}$ we have

$$w_p(\mathbf{p}^+, \mathbf{p}^+) = w_p(\mathbf{p}^-, \mathbf{p}^-) < w_p(\mathbf{p}^+, \mathbf{p}^-) - \varepsilon_2. \tag{155}$$

Combining Lemma 63 with equation (155) we can construct a gadget that "prefers" the unordered phase of two vertices to agree and also "prefers" the spin assignment to disagree.

Lemma 64. *A constant sized gadget J_2 can be constructed with two special vertices u, v and the following property: there exists a phase $\mathbf{p} \in \mathcal{Q}'$ satisfying simultaneously all of the following:*

1. *$A_1(\mathbf{p}) = \textsc{MaxLwt}(J_2)$, where*

$$A_1(\mathbf{p}) := \max_{\mathcal{Y};\, \mathcal{Y}(u)=\mathbf{p}^+, \mathcal{Y}(v)=\mathbf{p}^-} \textsc{Lwt}_{J_2}(\mathcal{Y}) = \max_{\mathcal{Y};\, \mathcal{Y}(u)=\mathbf{p}^-, \mathcal{Y}(v)=\mathbf{p}^+} \textsc{Lwt}_{J_2}(\mathcal{Y}). \tag{156}$$

2. *Among $\mathbf{p}$ that satisfy Item 1, $\mathbf{p}$ maximizes*

$$A_2(\mathbf{p}) := \max_{\mathcal{Y};\, \mathcal{Y}(u)=\mathbf{p}^+, \mathcal{Y}(v)=\mathbf{p}^+} \textsc{Lwt}_{J_2}(\mathcal{Y}) = \max_{\mathcal{Y};\mathcal{Y}(u)=\mathbf{p}^-, \mathcal{Y}(v)=\mathbf{p}^-} \textsc{Lwt}_{J_2}(\mathcal{Y}). \tag{157}$$

3. *The following inequalities hold*

$$A_1(\mathbf{p}) > A_2(\mathbf{p}) + \varepsilon_3 \quad \textit{and} \quad A_2(\mathbf{p}) > \varepsilon_3 + \max_{\mathcal{Y};\mathcal{Y}'(u)\neq\mathcal{Y}'(v)} \mathrm{L{\small WT}}_{J_2}(\mathcal{Y}), \tag{158}$$

where $\varepsilon_3 > 0$ is a constant (depending only on the model and Δ).

Proof. To construct J_2 we take $t := 3\lceil(\max_{\mathbf{p_1},\mathbf{p_2}} w_p(\mathbf{p}_1,\mathbf{p}_2) - \min_{\mathbf{p_1},\mathbf{p_2}} w_p(\mathbf{p}_1,\mathbf{p}_2))/\varepsilon_1\rceil$ copies of gadget J_1 from Lemma 63, identify (merge) their u vertices, and identify (merge) their v vertices. Finally we add a parallel edge between u and v.

Let $\mathbf{p}$ be the unordered phase that is the common value of $\mathcal{Y}'(u)$ and $\mathcal{Y}'(v)$ for which the maximum on the left-hand side of (154) is achieved (note that $\mathbf{p}$ is not unique; we just take one such $\mathbf{p}$). Let

$$A_4 := \max_{\mathcal{Y};\, \mathcal{Y}'(u)=\mathbf{p},\mathcal{Y}'(v)=\mathbf{p}} \mathrm{L{\small WT}}_{J_2}(\mathcal{Y}) \quad \text{and} \quad A_5 := \max_{\mathcal{Y};\, \mathcal{Y}'(u)\neq\mathcal{Y}'(v)} \mathrm{L{\small WT}}_{J_2}(\mathcal{Y}).$$

Then applying (154) on each copy of J_1 in J_2 we obtain

$$A_4 > A_5 + 2(\max_{\mathbf{p_1},\mathbf{p_2}} w_p(\mathbf{p}_1,\mathbf{p}_2) - \min_{\mathbf{p_1},\mathbf{p_2}} w_p(\mathbf{p}_1,\mathbf{p}_2)). \tag{159}$$

Thus the maximizer of $\max_{\mathcal{Y}} \mathrm{L{\small WT}}_{J_2}(\mathcal{Y})$ happens for $\mathcal{Y}$ with $\mathcal{Y}'(u) = \mathcal{Y}'(v)$. Only the parallel edge is influenced by the spin and hence, by (155), we have

$$\max_{\mathcal{Y}} \mathrm{L{\small WT}}_{J_2}(\mathcal{Y}) = \max_{\mathbf{p}} \max_{\mathcal{Y}\colon \mathcal{Y}(u)=\mathbf{p}^+,\mathcal{Y}(v)=\mathbf{p}^-} \mathrm{L{\small WT}}_{J_2}(\mathcal{Y}). \tag{160}$$

Let $\mathbf{p}$ be the maximizer on the right-hand side of (160) that (secondarily) maximizes the second expression in (157). Note that $\mathbf{p}$ satisfies the first and second condition of the lemma. The first part of the third condition is satisfied for any $\varepsilon_3 \leq \varepsilon_2$ (using (155)). Recall that $\varepsilon_2 > 0$. The second part of the third condition is satisfied for $\varepsilon_3 \leq \max_{\mathbf{p_1},\mathbf{p_2}} w_p(\mathbf{p}_1,\mathbf{p}_2) - \min_{\mathbf{p_1},\mathbf{p_2}} w_p(\mathbf{p}_1,\mathbf{p}_2)$. Recall that $\max_{\mathbf{p_1},\mathbf{p_2}} w_p(\mathbf{p}_1,\mathbf{p}_2) - \min_{\mathbf{p_1},\mathbf{p_2}} w_p(\mathbf{p}_1,\mathbf{p}_2) > 0$. Thus we can take $\varepsilon_3 > 0$ to be the smaller of the two upper bounds (each of which is a constant depending on the model and Δ only). $\square$

Lemma 65. *Let $\mathbf{B}$ be the interaction matrix of an antiferromagnetic spin system. Let A_1, A_2 be the constants defined in Lemma 64. There exist constants D_1, D_2, D_3 depending* *only on the model and Δ such that the following is true. Given a cubic graph H we can, in polynomial-time, construct a max-degree-D_1 graph G with $|V(G)| \leq D_2|V(H)|$ such that*

$$\textsc{MaxLwt}(G) = (A_1 - A_2)\textsc{MaxCut}(H) + A_2|E(H)| + A_1 D_3|V(H)|.$$

We can now go back and prove the inapproximability result for the phase labeling problem.

Proof of Lemma 56. Since A_1, A_2, D_3 are constants depending only on the model and Δ, the trivial algorithm gives the bound $\textsc{MaxCut}(H) \geq 1/2|E(H)| = 3/2|V(H)|$. Together with Lemma 65 we obtain the result. $\square$

Proof of Lemma 65. Replace each edge of H by gadget J_2 and for each vertex $w \in V(H)$ add D_3 new vertices $w_1, \dots, w_{D_3}$ and add a gadgets J_2 between w and w_i (for $i \in [D_3]$), where D_3 will be determined shortly.

The purpose of the D_3 copies of J_2 is to force phase $\mathbf{p}$ (from Lemma 64) to be used on the special vertices in a labeling of G with maximum weight. A phase $\mathbf{s} \neq \mathbf{p}$ can have

$$\ell_1(\mathbf{s}) := \max_{\mathcal{Y};\, \mathcal{Y}(u)=\mathbf{s}^+, \mathcal{Y}(v)=\mathbf{s}^+} \textsc{Lwt}_{J_2}(\mathcal{Y}) - \max_{\mathcal{Y};\, \mathcal{Y}(u)=\mathbf{p}^+, \mathcal{Y}(v)=\mathbf{p}^+} \textsc{Lwt}_{J_2}(\mathcal{Y}) > 0, \tag{161}$$

but then by the choice of $\mathbf{p}$

$$\ell_2(\mathbf{s}) := \max_{\mathcal{Y};\, \mathcal{Y}(u)=\mathbf{p}^+, \mathcal{Y}(v)=\mathbf{p}^-} \textsc{Lwt}_{J_2}(\mathcal{Y}) - \max_{\mathcal{Y};\, \mathcal{Y}(u)=\mathbf{s}^+, \mathcal{Y}(v)=\mathbf{s}^-} \textsc{Lwt}_{J_2}(\mathcal{Y}) > 0. \tag{162}$$

Let

$$D_3 = 4 + 3\left\lceil \max_{\mathbf{s}} \frac{\ell_1(\mathbf{s})}{\ell_2(\mathbf{s})} \right\rceil,$$

where the maximum is taken over $\mathbf{s}$ such that (161) is satisfied (if no such $\mathbf{s}$ exists we can take $D_3 = 0$). Note that D_3 is a constant depending on the model and Δ only.

Now we want to find the maximum weight labeling of G. We are only going to focus on labeling of the special vertices (u's and v's in the J_2 gadgets), since once those are fixed one just finds the optimal labeling in each gadget (conditioned on the labels of special vertices). Let $\mathcal{U}$ be a labeling of the special vertices that leads to the maximum weight labeling of G. Let $\widehat{\mathcal{U}}$ be the labeling obtained from $\mathcal{U}$ by changing the phase of each special vertex to $\mathbf{p}$ while (1) keeping the original spin on the vertices of H, and (2) making the spin of $w_1, \ldots, w_{D_3}$ the opposite of the spin of $\mathcal{U}$ (for each $w \in V(H)$). Now we compare $\mathcal{U}$ and $\widehat{\mathcal{U}}$ for each J_2 gadget corresponding to edge of H:

- if in $\mathcal{U}$ the phase of u and v were different then $\widehat{\mathcal{U}}$ has higher weight than $\mathcal{U}$ on the gadget, using (158);
- if in $\mathcal{U}$ the phase of u and v is the same but the spin is different then $\widehat{\mathcal{U}}$ has greater or equal weight than $\mathcal{U}$ on the gadget, using (156);
- if in $\mathcal{U}$ the phase of u and v is the same and the spin is the same that the loss of $\widehat{\mathcal{U}}$ on the gadget (compared to $\mathcal{U}$) is $\ell_1(\mathbf{s})$ (where $\mathbf{s}$ is the phase of u, v in $\mathcal{U}$).

For the J_2 gadgets connecting $\mathcal{U}$ to $w_1, \ldots, w_{D_3}$ we have

- if the phase of w in $\mathcal{U}$ was $\mathbf{s}$ such that $\ell_1(\mathbf{s}) > 0$ then the gain of $\widehat{\mathcal{U}}$ on each gadget (compared to $\mathcal{U}$) is at least $\ell_2(\mathbf{s})$;
- otherwise, by (156) $\widehat{\mathcal{U}}$ has greater or equal weight than $\mathcal{U}$ on the gadget.

For each vertex whose phase in $\mathcal{U}$ was $\mathbf{s}$ such that $\ell_1(\mathbf{s}) > 0$ there are 3 edges where $\widehat{\mathcal{U}}$ can lose $\ell_1(\mathbf{s})$ (compared to $\mathcal{U}$) but there are D_3 edges where $\widehat{\mathcal{U}}$ gains $\ell_2(\mathbf{s})$ (compared to $\mathcal{U}$). Since $D_3\ell_2(\mathbf{s}) > 3\ell_1(\mathbf{s})$ we have that $\widehat{\mathcal{U}}$ has at least as large weight as $\mathcal{U}$ (and hence is also optimal).

Now we just argue how the spins should be assigned. The largest number of J_2 gadgets with opposite spins on the special vertices arises when we take the max-cut of H and assign the spin according to the cut. □

5.5.3 The gadget

Proof of Lemma 63. Let $\mathbf{z}_1, \ldots, \mathbf{z}_Q := \mathbf{r}_1, \ldots, \mathbf{r}_{Q'}, \mathbf{c}_1, \ldots, \mathbf{c}_{Q'}$. Let $\mathbf{u}$ be defined as in Equation (151). In Section 5.5.1, Lemma 62 it is proved that the $Q \times Q$ matrix $\hat{\mathbf{A}}$ whose ij-th entry is $\ln(\mathbf{z}_i^\intercal \mathbf{B} \mathbf{z}_j) - \ln(\mathbf{z}_i^\intercal \mathbf{u}) - \ln(\mathbf{z}_j^\intercal \mathbf{u})$ is negative definite. Let $\mathbf{A}'$ be the $Q' \times Q'$ matrix obtained by the following "folding" of $\hat{\mathbf{A}}$:

$$\mathbf{A}'_{ij} = \hat{\mathbf{A}}_{i,j} + \hat{\mathbf{A}}_{i+Q',j} + \hat{\mathbf{A}}_{i,j+Q'} + \hat{\mathbf{A}}_{i+Q',j+Q'}.$$

We have that $\mathbf{A}'$ is also negative definite (since $\mathbf{x}^\intercal\mathbf{A}'\mathbf{x} = \mathbf{y}^\intercal\hat{\mathbf{A}}\mathbf{y}'$, where $\mathbf{y}^\intercal = (\mathbf{x}^\intercal, \mathbf{x}^\intercal)$). Note that

$$\mathbf{A}'_{ij} = w_s((\mathbf{x}_i, \mathbf{y}_i), (\mathbf{x}_j, \mathbf{y}_j)) - a'_i - a'_j,$$

where $a'_i := 2\ln(\mathbf{x}_i^\intercal\mathbf{u}) + 2\ln(\mathbf{y}_i^\intercal\mathbf{u})$.

Let λ_1 be largest eigenvalue of $-\mathbf{A}'$ and let λ_2 be the smallest eigenvalue of $-\mathbf{A}'$. Note that $0 < \lambda_2 \leq \lambda_1$. Define $\mathbf{A}$ to be the $Q' \times Q'$ matrix with $A_{ij} = A'_{ij} + a'_i + a'_j$ and consider the following maximization problem

$$\max_{\mathbf{x};\mathbf{x}^\intercal\mathbf{1}=1,\mathbf{x}\geq 0} \mathbf{x}^\intercal\mathbf{A}\mathbf{x}. \tag{163}$$

Note that for $\mathbf{x}$ with $\mathbf{x}^\intercal\mathbf{1} = 1$ we have

$$\mathbf{x}^\intercal\mathbf{A}\mathbf{x} = 2\mathbf{a}'^\intercal\mathbf{x} + \mathbf{x}^\intercal\mathbf{A}'\mathbf{x}, \tag{164}$$

where $\mathbf{A}'$ is negative definite. Note that if $\mathbf{x}$ and $\mathbf{y}$ are distinct optimal solutions of (163) then $(x + y)/2$ satisfies all the constraints, and from (164) and negative definiteness of A' we have

$$((\mathbf{x}+\mathbf{y})/2)^\intercal\mathbf{A}((\mathbf{x}+\mathbf{y})/2) > \left(\mathbf{x}^\intercal\mathbf{A}\mathbf{x} + \mathbf{y}^\intercal\mathbf{A}\mathbf{y}\right)/2,$$

a contradiction (with optimality of both x and y). Thus (163) has a unique maximum; let $\mathbf{x}^*$ be the value of $\mathbf{x}$ achieving it. Let O^* be $(\mathbf{x}^*)^\intercal\mathbf{A}\mathbf{x}^*$. Let S be the set of non-zero coordinates in $\mathbf{x}^*$.

Let $\mathbf{y} \in \mathbb{R}^{Q'}$ be such that $\mathbf{y}^\intercal\mathbf{1} = 0$ and $\mathbf{y}$ is zero on coordinates outside S. Then from (local) optimality of $\mathbf{x}^*$ we have

$$(\mathbf{x}^*+\mathbf{y})^\intercal\mathbf{A}(\mathbf{x}^*+\mathbf{y}) = O^*+2(\mathbf{a}'^\intercal+(\mathbf{x}^*)^\intercal\mathbf{A}')\mathbf{y}+\mathbf{y}^\intercal\mathbf{A}'\mathbf{y} = O^*+\mathbf{y}^\intercal\mathbf{A}'\mathbf{y} \geq O^*-\lambda_1\|\mathbf{y}\|_2^2. \tag{165}$$

Equation (165) tells us that moving slightly from the optimum the objective decreases at most quadratically in the length of $\mathbf{y}$.

Let $\mathbf{y} \in \mathbb{R}^{Q'}$ be such that $\mathbf{y}^\intercal\mathbf{1} = 0$ and $\mathbf{y}$ is non-negative on coordinates outside S. Then from (local) optimality of $\mathbf{x}^*$ we have

$$(\mathbf{x}^*+\mathbf{y})^\intercal\mathbf{A}(\mathbf{x}^*+\mathbf{y}) = O^*+2(\mathbf{a}'^\intercal+(\mathbf{x}^*)^\intercal\mathbf{A}')\mathbf{y}+\mathbf{y}^\intercal\mathbf{A}'\mathbf{y} = O^*+\mathbf{y}^\intercal\mathbf{A}'\mathbf{y} \geq O^*-\lambda_2\|\mathbf{y}\|_2^2. \tag{166}$$

Equation (165) tells us that moving slightly from the optimum the objective decreases at least quadratically in the length of $\mathbf{y}$.

Let $Z \geq (4Q'\lambda_1/\lambda_2)^{Q'}$. Note that Z is a constant depending only on the model and Δ. Let $z_1/z, \ldots, z_{Q'}/z$ be the optimal simultaneous Diophantine approximation of $x_1^*, \ldots, x_{Q'}^*$ with $z_1, \ldots, z_{Q'}, z \in \mathbb{Z}$ and $1 \leq z \leq Z$. By Dirichlet's theorem we have

$$|zx_i^* - z_i| \leq Z^{-1/Q'} < 1. \tag{167}$$

Note that (167) implies

$$\text{if } x_i^* = 0 \text{ then } z_i = 0. \tag{168}$$

Also note that

$$\left| \sum_{i=1}^{Q'} zx_i^* - \sum_{i=1}^{Q'} z_i \right| \leq \sum_{i=1}^{Q'} |zx_i^* - z_i| \leq Q'Z^{-1/Q'} < 1,$$

and since z and z_i's are integers and $(\mathbf{x}^*)^\intercal \mathbf{1} = 1$ we have

$$\sum_{i=1}^{Q'} \frac{z_i}{z} = 1. \tag{169}$$

From (168) and (169) we have that for $\mathbf{y} := (z_1/z, \ldots, z_{Q'}/z) - \mathbf{x}^*$ we can apply (165) and hence

$$(z_1/z, \ldots, z_{Q'}/z)\mathbf{A}(z_1/z, \ldots, z_{Q'}/z)^\intercal \geq O^* - \lambda_1 Q' Z^{-2/Q'} z^{-2}. \tag{170}$$

Now we are ready to construct the gadget J_1. First, let K be the multigraph on z vertices $b_1, b_2, \ldots, b_z$ with the following symmetric edges: self-loop on b_i for $i \in [z]$ and two edges between b_i and b_j for every $i, j \in [z]$ with $i \neq j$. To obtain J_1, we overlay two copies of K as follows. Let K_u (resp. K_v) be a copy of K, where the image of b_z is renamed to u (resp. v). Overlay K_u, K_v by identifying the images of $b_1, \ldots, b_{z-1}$ in the two copies. Thus, the resulting graph J_1 has $z+1$ vertices and the following edges: two self loops on b_i for $i \in [z-1]$, four edges between b_i and b_j for every $i, j \in [Q'-1]$ with $i \neq j$, two edges between u and b_i for $i \in [z-1]$, two edges between v and b_i for $i \in [z-1]$ and a self loop on u, v.

Note that the weight of a phase assignment on J_1 is the sum of the induced phase assignments on K_u and K_v. Consider an assignment of phases $\mathcal{Y}_o$ such that in each complete graph z_i vertices get phase i (note that this forces the phases of u and v to be the same). The weight of the phase assignment $\mathcal{Y}_o$ is

$$\textsc{Lwt}_{J_1}(\mathcal{Y}_o) = \mathrm{S}_1 := 2(z_1, \ldots, z_{Q'})\mathbf{A}(z_1, \ldots, z_{Q'})^\intercal \geq 2z^2 O^* - 2\lambda_1 Q' Z^{-2/Q'}. \tag{171}$$

Now suppose that we have a phase assignment $\mathcal{Y}$ for J_1 where the phases of u and v are different. Let $\hat{\mathbf{u}}$ be the vector with $\hat{u}_i$ counting the number of vertices with phase i in K_u and define similarly $\hat{\mathbf{v}}$.

Note that $\|\hat{\mathbf{u}} - \hat{\mathbf{v}}\|_2^2 = 2$ (since $\hat{\mathbf{u}}$ and $\hat{\mathbf{v}}$ differ in two coordinates—the phases of u and v in the assignment). By triangle inequality we have $\|\hat{\mathbf{u}}/z - \mathbf{x}^*\|_2 \geq 1/(z\sqrt{2})$ or $\|\hat{\mathbf{v}}/z - \mathbf{x}^*\|_2 \geq 1/(z\sqrt{2})$ (otherwise we would have $\|\hat{\mathbf{u}}/z - \hat{\mathbf{v}}/z\|_2 < \sqrt{2}/z$). W.l.o.g. assume that $\hat{\mathbf{u}}/z$ has the greater distance from $\mathbf{x}^*$. We have

$$\textsc{Lwt}_{J_1}(\mathcal{Y}_o) = \mathrm{S}_2 := \hat{\mathbf{u}}^\intercal \mathbf{A}\hat{\mathbf{u}} + \hat{\mathbf{v}}^\intercal \mathbf{A}\hat{\mathbf{v}} \leq z^2(2O^* - \lambda_2/(2z^2)) = 2z^2 O^* - \lambda_2/2. \tag{172}$$

By our choice of Z we have $S_1 > S_2$ and hence in an optimal phase assignment for J_1 we have that u and v get the same phase. Note that we did not show which phase assignment is optimal; we only found a phase assignment in which u, v have the same phase that is better than any assignment in which u, v have different phases. $\square$

CHAPTER VI

PHASE DIAGRAMS

In this chapter, we study the phase diagrams for spin systems of particular interest. We will be working on random Δ-regular bipartite graphs, and hence phase diagrams should be understood as the maximizers $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ of the limit $\frac{1}{n} \log \mathbf{E}_G[Z_G^{\boldsymbol{\alpha},\boldsymbol{\beta}}]$, see also Definition 6 in Section 3.3. We do this for the following spin systems: antiferromagnetic 2-spin systems with external field, the ferromagnetic Potts model, the antiferromagnetic Potts and the colorings model.

Our general approach will be to use Theorem 2 as far as possible, that is, we will identify whenever possible the local maxima of the function Ψ_1. By Theorem 2, this is equivalent to identifying the stable fixpoints of the tree recursions. Of course, this is trivial when the parameters of the model lie in the uniqueness regime since there only one fixpoint exists, which by Theorem 5 corresponds to the (unique) global maximum of Ψ_1. Thus, the more interesting case is the non-uniqueness regime.

For antiferromagnetic 2-spin systems, identifying the stable fixpoints of the tree recursions will be sufficient to find the phases since they will be the *only local* maximizers of the function Ψ_1. In contrast, for the q-spin systems we investigate, more work is needed, essentially because the number of non-permutation symmetric fixpoints is a function of q.

For the ferromagnetic Potts model, the situation remains relatively simple: the non-permutation symmetric stable fixpoints turn out to not vary with q and Δ, and hence we can find the dominant phases by comparing the value of Ψ_1 at the local maximizers. An interesting phenomenon arises here, the so called phase coexistence. In particular, at a certain temperature $B = \mathfrak{B}_o$ (where $\mathfrak{B}_o$ is a function of q, Δ), the global maximizers are no longer permutation symmetric and each has total measure a constant fraction of the Gibbs distribution. The phase coexistence phenomenon for the ferromagnetic Potts model has been observed and rigorously proved on the complete graph (this setting should be considered simpler than ours) and on the grid $\mathbb{Z}^2$ (this setting should be considered harder than ours).

The analysis of dominant phases becomes much more complicated for the antiferromagnetic Potts model and the colorings model. First, the local maximizers now turn out to depend on both q, Δ. Second, the structure of the local maximizers needs several variables to be captured, which makes analytical penetration much more intricate.

6.1 Antiferromagnetic 2-spin systems

A 2-spin system is specified by parameters $B_1, B_2 \geq 0$ and $\lambda > 0$. To avoid trivial models we will assume that at least one of B_1, B_2 is bigger than 0. The edge interaction matrix for a 2-spin system with parameters B_1, B_2 is given by

$$\mathbf{B} = \begin{bmatrix} B_1 & 1 \\ 1 & B_2 \end{bmatrix}.$$

The system is antiferromagnetic if $B_1 B_2 < 1$ and ferromagnetic otherwise. For the purposes of this section only, we will index the rows and columns of $\mathbf{B}$ by $\{0,1\}$ (instead of $\{1, 2 \ldots, q\}$ used for q-spin systems).

For a graph $G = (V, E)$ and a configuration $\sigma : V \to \{0, 1\}$, the weight of the configuration σ is given by

$$w_G(\sigma) = \lambda^{|\sigma^{-1}(0)|} \prod_{(u,v) \in E} B_{\sigma(u), \sigma(v)},$$

and the partition function is given by $Z_G = \sum_\sigma w_G(\sigma)$.

We note that for Δ-regular graphs, the external field λ may be pushed into the interaction matrix $\mathbf{B}$, giving a 2-spin system specified by the interaction matrix

$$\mathbf{B}_\lambda = \begin{bmatrix} \lambda^{2/\Delta} B_1 & \lambda^{1/\Delta} \\ \lambda^{1/\Delta} & B_2 \end{bmatrix}.$$

Thus, Theorem 14 applies, provided that its assumptions are satisfied. As a first step, it is important to characterize the semi-translation invariant non-uniqueness regime for $\mathbf{B}_\lambda$. For two-spin systems, this turns out to coincide with the non-uniqueness regime. As a starting point, we study the fixpoints of the tree recursions 8 for $\mathbf{B}_\lambda$, which can be written as

$$\begin{aligned} R_0 &\propto (B_1 C_0 + \lambda^{1/\Delta} C_1)^{\Delta-1} \text{ and } R_1 \propto (\lambda^{1/\Delta} C_0 + \lambda^{2/\Delta} B_2 C_1)^{\Delta-1}, \\ C_0 &\propto (B_1 R_0 + \lambda^{1/\Delta} R_1)^{\Delta-1} \text{ and } C_1 \propto (\lambda^{1/\Delta} R_0 + \lambda^{2/\Delta} B_2 R_1)^{\Delta-1}. \end{aligned} \tag{173}$$

We can easily transform the system (173) into a system with two variables, by substituting $x' = R_0/R_1$ and $y' = C_0/C_1$. This gives the equivalent system of equations

$$x' = \left(\frac{\lambda^{2/\Delta} B_1 y' + \lambda^{1/\Delta}}{\lambda^{1/\Delta} y' + B_2}\right)^{\Delta-1} \text{ and } y' = \left(\frac{\lambda^{2/\Delta} B_1 x' + \lambda^{1/\Delta}}{\lambda^{1/\Delta} x' + B_2}\right)^{\Delta-1} \tag{174}$$

A further simplification is also possible giving a direct correspondence with the standard form of the tree recursions for the original spin system (with the external field λ). Namely, if we set $x = \lambda^{1/\Delta} x'$ and $y = \lambda^{1/\Delta} y'$, we obtain

$$x = \lambda\left(\frac{B_1 y + 1}{y + B_2}\right)^{\Delta-1} \text{ and } y = \lambda\left(\frac{B_1 y + 1}{y + B_2}\right)^{\Delta-1} \tag{175}$$

The system (175) always has a positive solution with $x = y =: p$. The quantity p captures the uniqueness regime on the infinite tree. Precisely, the 2-spin system with parameters B_1, B_2, λ is in the uniqueness regime of $\hat{\mathbb{T}}_\Delta$ iff

$$(\Delta - 1)\omega \le 1, \text{ where } \omega := \frac{(1 - B_1 B_2)p}{(B_2 + p)(1 + B_1 p)}. \tag{176}$$

To gain some intuition, let us look at the stability matrix of the tree recursions. By Lemma 33, for a general interaction matrix $\mathbf{B}$, the condition for Jacobian stability of a fixpoint of the tree recursions is related to the spectrum of $\mathbf{L} = \left[\begin{smallmatrix} \mathbf{0} & \mathbf{A} \\ \mathbf{A}^\intercal & \mathbf{0} \end{smallmatrix}\right]$, where $\mathbf{A}$ is the $q \times q$ matrix whose ij-entry is given by $A_{ij} = B_{ij} R_i C_j / \sqrt{\alpha_i \beta_j}$ and α_i, β_j are given by (43). Recall that ± 1 are eigenvalues of $\mathbf{L}$ and the condition for Jacobian stability is that all the other eigenvalues have absolute value less than $1/(\Delta - 1)$ (see for details the proof of Theorem 2 in Section 3.4.3). We next show that $\pm\omega$ in (176) are the nontrivial eigenvalues of $\mathbf{L}$ and hence, when $(\Delta - 1)\omega < 1$ the fixpoint $x = y = p$ is Jacobian stable.

To do this, we trace back a solution to (173), which is given by $R_0 \propto p/\lambda^{1/\Delta}, R_1 \propto 1$ and $C_0 \propto p/\lambda^{1/\Delta}, C_1 \propto 1$. The matrix $\mathbf{A}$ for the spin model $\mathbf{B}_\lambda$ is thus given by

$$\begin{bmatrix} \frac{B_1 p}{B_1 p + 1} & \frac{p^{1/2}}{(B_1 p+1)^{1/2}(p+B_2)^{1/2}} \\ \frac{p^{1/2}}{(B_1 p+1)^{1/2}(p+B_2)^{1/2}} & \frac{B_2}{p+B_2} \end{bmatrix},$$

whose eigenvalues are easily checked to be $1, -\omega$. Thus, the eigenvalues of $\mathbf{J}$ are $\pm 1, \pm\omega$, as claimed.

The above arguments show that the translation invariant fixpoint of (173) is unstable in the non-uniqueness regime of $\mathbb{T}_\Delta$. By Theorem 2, it does not correspond to a Hessian local maximum of $\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$. Thus, in the non-uniqueness regime of $\mathbb{T}_\Delta$, the dominant phases correspond to positive solutions of (175) with $x \neq y$. It is a standard fact that there are exactly two such pairs $(p^+, p^-), (p^-, p^+)$. Both of these should correspond to the global maxima of $\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$. All we are left to do now is verify that they are Hessian local maxima, or equivalently (by Theorem 2) that the respective fixpoint of (175) is Jacobian stable.

Similar arguments as above yield that Jacobian stability of the two fixpoints is captured by the following inequality. For convenience, let $x = p^+, y = p^-$.

$$(\Delta - 1)^2 \omega^* < 1 \text{ where } \omega^* = \frac{(1 - B_1 B_2)^2 xy}{(B_1 x + 1)(B_1 y + 1)(B_2 + x)(B_2 + y)}. \tag{177}$$

We prove the following technical lemma.

Lemma 66. *Every positive solution of* (175) *with $x \neq y$ satisfies* (177)*.*

We prove Lemma 66 shortly. First, we give the proof of Theorem 11.

Proof of Theorem 11. Recall that for 2-spin systems, non-uniqueness coincides with the existence of multiple semi-translation invariant Gibbs measures on the tree. Moreover, in non-uniqueness, the maximizers of Ψ_1 are exactly two pairs $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ and $(\boldsymbol{\beta}, \boldsymbol{\alpha})$ with $\boldsymbol{\alpha} \neq \boldsymbol{\beta}$. Thus, they do not correspond to translation invariant measures and satisfy the permutation symmetric property. Finally, as discussed above, Lemma 66 establishes that they are Hessian dominant and hence the hypotheses of Theorem 14 are satisfied. $\square$

For the proof of Lemma 66, we will use the following intermediate lemma.

Lemma 67. *For $\Delta \geq 3$, every positive solution of* (175) *with $x \neq y$ satisfies*

$$B_1 xy + B_1 B_2 (x + y) + B_2 \geq (\Delta - 2)(1 - B_1 B_2)\sqrt{xy}.$$

With Lemma 67, it is easy to establish Lemma 177.

Proof of Lemma 66. Let $W := B_1xy + B_1B_2(x+y) + B_2$. Observe that

$$W = (B_1x+1)(y+B_2) - (1-B_1B_2)y = (B_1y+1)(x+B_2) - (1-B_1B_2)x.$$

Using the two expressions of W and the AM-GM inequality, we obtain

$$\begin{aligned} W^2 &= (B_1x+1)(y+B_2)(B_1y+1)(x+B_2) + (1-B_1B_2)^2xy \\ &\quad - (1-B_1B_2)\big(y(B_1x+1)(B_2+y) + x(B_1y+1)(B_2+x)\big) \\ &\le \left(\sqrt{(B_1x+1)(B_1y+1)(B_2+x)(B_2+y)} - (1-B_1B_2)\sqrt{xy}\right)^2. \end{aligned} \tag{178}$$

Equality in (178) can only hold if $y(B_1x+1)(B_2+y) = x(B_1y+1)(B_2+x)$, which in conjuction with (175) gives $x = y$. Thus, the inequality (178) is in fact strict. To take square roots in (178), note that trivially

$$(B_1x+1)(B_1y+1)(B_2+x)(B_2+y) > xy > (1-B_1B_2)^2xy,$$

so (178) and Lemma 67 give

$$\sqrt{(B_1x+1)(y+B_2)(B_1y+1)(x+B_2)} - (1-B_1B_2)\sqrt{xy} > (\Delta-2)(1-B_1B_2).$$

which after massaging gives (177). □

Proof of Lemma 67. For convenience, set $d = \Delta - 1$. The x, y satisfy

$$x = \lambda\left(\frac{B_1y+1}{y+B_2}\right)^d \text{ and } y = \lambda\left(\frac{B_1x+1}{x+B_2}\right)^d.$$

It follows that

$$x\left(\frac{B_1x+1}{x+B_2}\right)^d = y\left(\frac{B_1y+1}{y+B_2}\right)^d \Rightarrow x\Big((B_1x+1)(y+B_2)\Big)^d = y\Big((B_1y+1)(x+B_2)\Big)^d. \tag{179}$$

Let $W := B_1xy + B_1B_2x + B_1B_2y + B_2$. Observe that

$$\begin{aligned} (B_1x+1)(y+B_2) &= B_1xy + B_1B_2x + y + B_2 = W + (1-B_1B_2)y, \\ (B_1y+1)(x+B_2) &= B_1xy + B_1B_2y + x + B_2 = W + (1-B_1B_2)x. \end{aligned}$$

Hence, (179) gives

$$x\left(W + (1-B_1B_2)y\right)^d = y\left(W + (1-B_1B_2)x\right)^d.$$

Expanding using Newton's formula gives

$$x\sum_{k=0}^{d}\binom{d}{k}W^{d-k}\big((1-B_1B_2)y\big)^k = y\sum_{k=0}^{d}\binom{d}{k}W^{d-k}\big((1-B_1B_2)x\big)^k,$$

which is equivalent to

$$W^d(x-y) = xy\sum_{k=2}^{d}\binom{d}{k}W^{d-k}(1-B_1B_2)^k\left(x^{k-1}-y^{k-1}\right).$$

Since $x \neq y$, this can be rewritten as

$$W^d = xy\sum_{k=2}^{d}\binom{d}{k}W^{d-k}(1-B_1B_2)^k\left(\frac{x^{k-1}-y^{k-1}}{x-y}\right). \tag{180}$$

Claim 68. *For $k \geq 2$ and $x, y > 0$ with $x \neq y$, it holds that $\frac{x^{k-1}-y^{k-1}}{x-y} \geq (k-1)(xy)^{(k-2)/2}$.*

The simple proof of Claim 68 is given at the end. Using Claim 68, (180) gives

$$W^d \geq \sum_{k=2}^{d}\binom{d}{k}(k-1)W^{d-k}(1-B_1B_2)^k(xy)^{k/2},$$

or equivalently

$$\underbrace{W^d + \sum_{k=2}^{d}\binom{d}{k}W^{d-k}(1-B_1B_2)^k(xy)^{k/2}}_{C} \geq \underbrace{\sum_{k=2}^{d}\binom{d}{k}kW^{d-k}(1-B_1B_2)^k(xy)^{k/2}}_{D}. \tag{181}$$

Using again Newton's formula and the identity $\binom{d}{k} = \frac{d}{k}\binom{d-1}{k-1}$, we have

$$\begin{aligned}
C &= \big(W + (1-B_1B_2)\sqrt{xy}\big)^d - dW^{d-1}(1-B_1B_2)\sqrt{xy},\\
D &= d(1-B_1B_2)\sqrt{xy}\sum_{k=2}^{d}\binom{d-1}{k-1}W^{d-k}(1-B_1B_2)^{k-1}(\sqrt{xy})^{k-1}\\
&= d(1-B_1B_2)\sqrt{xy}\Big(\big(W + (1-B_1B_2)\sqrt{xy}\big)^{d-1} - W^{d-1}\Big).
\end{aligned}$$

Thus, (181) gives

$$W \geq (d-1)(1-B_1B_2)\sqrt{xy},$$

which is exactly the inequality we wanted. Finally, we give the proof of Claim 68.

Proof of Claim 68. Since $k \geq 2$, observe that

$$\frac{x^{k-1} - y^{k-1}}{x - y} = x^{k-2} + x^{k-3}y + \ldots + y^{k-2} \geq (k-1)\Big((xy)^{(k-1)(k-2)/2}\Big)^{1/(k-1)} = (k-1)(xy)^{(k-2)/2}.$$

The inequality is an application of the AM-GM inequality to $x^{k-2}, \ldots, y^{k-2}$. Equality holds iff $x = y$. □

This completes the proof of Lemma 67. □

6.2 Ferromagnetic Potts model

Häggström [32] established that the uniqueness threshold $\mathfrak{B}_u$ for the q-state ferromagnetic Potts model with parameter B is the unique value for which the following polynomial has a double root in $(0, 1)$:

$$(q-1)x^{\Delta} + (2 - B - q)\, x^{\Delta-1} + Bx - 1. \tag{182}$$

We prove that the ferromagnetic Potts model on random Δ-regular graphs undergoes a first order phase transition at a parameter $\mathfrak{B}_o > \mathfrak{B}_u$,(which was considered by Peruggi et al. [59]):

$$\mathfrak{B}_o := \frac{q-2}{(q-1)^{(1-2/\Delta)} - 1}.$$

Finally, Häggström [32] considers the following activity $\mathfrak{B}_{rc}$, which he conjectures is a (second) threshold for uniqueness of the random-cluster model, defined as:

$$\mathfrak{B}_{rc} := 1 + \frac{q}{\Delta - 2}.$$

Note, $\mathfrak{B}_u < \mathfrak{B}_o < \mathfrak{B}_{rc}$.

In this section we prove the following Theorem 69 detailing the phase diagram for the Potts model. This consists of $q + 1$ phases; the disordered phase and q permutation symmetric ordered phases. The disordered phase corresponds the uniform fixpoint with $\boldsymbol{\alpha} = (1/q, \ldots, 1/q)$. When we refer to an ordered phase $\boldsymbol{\alpha}$ this refers specifically to a phase with one color dominating in the following sense: one coordinate is equal to $a > 1/q$ and the other $q - 1$ coordinates are equal to $(1 - a)/(q - 1)$.

Theorem 69. *For the ferromagnetic Potts model the following holds at activity B:*

$B < \mathfrak{B}_u$: *There is a unique infinite-volume Gibbs measure on $\mathbb{T}_\Delta$. The disordered phase is Hessian dominant phase, and there are no other local maxima of Ψ_1.*

$\mathfrak{B}_u < B < \mathfrak{B}_{rc}$: *The local maxima of Ψ_1 are the disordered phase $\mathbf{u}$ and the q ordered phases (the ordered phases are permutations of each other). All of these $q+1$ phases are Hessian local maxima. Moreover:*

$\mathfrak{B}_u < B < \mathfrak{B}_o$: *The disordered phase is Hessian dominant.*

$B = \mathfrak{B}_o$: *Both the disordered phase and the ordered phases are Hessian dominant.*

$\mathfrak{B}_o < B < \mathfrak{B}_{rc}$: *The ordered phases are Hessian dominant.*

$B \geq \mathfrak{B}_{rc}$: *The q ordered phases (which are permutations of each other) are Hessian dominant. For $B > \mathfrak{B}_{rc}$ there are no other local maxima of Ψ_1.*

To prove Theorem 69, in light of Theorem 2, in order to determine the local maxima for the Potts model, we need to compute the spectral radius of the map $L : (r_1, \ldots, r_q) \mapsto (\widehat{r}_1, \ldots, \widehat{r}_q)$, given by

$$\widehat{r}_i = \sum_{j=1}^{q} \frac{B_{ij} R_i R_j}{\sqrt{\alpha_i \alpha_j}} r_j \tag{183}$$

in the subspace

$$\sum_{i=1}^{q} \sqrt{\alpha_i} r_i = 0, \tag{184}$$

where the R_i's are fixed points of the tree recursions and the α_i's are given by

$$\alpha_i = R_i \sum_{j=1}^{q} B_{ij} R_j \text{ for } i = 1, \ldots, q.$$

Our goal is to determine the local maxima by verifying when the spectal radius of this map (in the subspace) is less than $1/(\Delta - 1)$.

First we argue that the solutions of the tree recurrences for the Potts model are simple—they have only two values.

Lemma 70. *Let $(R_1, \ldots, R_q)$ be a solution of the tree recursion of ferromagnetic Potts model. Then the R_i's have at most two distinct values.*

Proof. Let $r_i = R_i^{1/d}$ and $r = \sum_{i=1}^{q} r_i^d$, where $d := \Delta - 1$. We have

$$r_i = r + (B-1)r_i^d.$$

The polynomial $f(x) = (B-1)x^d - x + r$ has at most 2 positive roots (counted with their multiplicities; by the Descartes' rule of signs) and hence there are at most 2 different values of the r_i's. $\square$

Let

$$\mathbf{M} = \left\{ \frac{B_{ij} R_i R_j}{\sqrt{\alpha_i \alpha_j}} \right\}_{i,j=1}^{q}$$

be the matrix of the linear map L. Note that $\mathbf{M}$ is symmetric and has an eigenvalue equal to 1 with eigenvector $e = \left[\sqrt{\alpha_1}, \ldots, \sqrt{\alpha_q}\right]^{\mathrm{T}}$.

Lemma 71. *The fixed points of the tree recursions, assuming $R_1 \geq R_2 \geq \ldots R_q$, satisfy $R_1 = R_2 = \ldots = R_t$ and $R_{t+1} = \ldots R_q$ for some $1 \leq t \leq q$. It follows that $\alpha_1 = \alpha_2 = \ldots = \alpha_t$ and $\alpha_{t+1} = \ldots \alpha_q$.*

Proof. This follows from Lemma 70. $\square$

Remark 15. *Two settings for t in the setting of Lemma 71 will be of particular interest, namely $t = 1$ and $t = q$. We shall refer to the latter as the uniform fixpoint, and this corresponds to the disordered phase. We shall refer to fixpoints with $t = 1$ as the "majority" fixpoints. This class includes either one or two (depending on the value of B, c.f. Lemma 75) fixpoints where color 1 dominates and the remaining appear with equal probability. The ordered phases correspond to the majority fixpoint for which the ratio R_1/R_q is maximum.*

Lemma 71 implies that $\mathbf{M}$ has a very simple structure. The following simple lemma describes the eigenvalues of $\mathbf{M}$.

Lemma 72. *Assume $1 \leq t < q$. Then $\mathbf{M}$ has the following eigenvalues:*

- 1 *with multiplicity* 1,
- $(B-1)R_1^2/\alpha_1$ *with multiplicity* $t-1$ *(assuming* $t > 1$*),*
- $(B-1)R_q^2/\alpha_q$ *with multiplicity* $q-t-1$ *(assuming* $t < q-1$*), and*

- $(B+t-1)R_1^2/\alpha_1 + (B+q-t-1)R_q^2/\alpha_q - 1$ *with multiplicity* 1.

For $t = q$ *the eigenvalues of* $\mathbf{M}$ *are*

- 1 *with multiplicity* 1,

- $(B-1)R_1^2/\alpha_1$ *with multiplicity* $q-1$.

Proof. We already described the eigenvector for 1. A vector with 1 at position 1 and -1 at a position i for $2 \le i \le t$ (and zeros elsewhere) yields eigenvalue $(B-1)R_1^2/\alpha_1$. Similarly, a vector with 1 at position q and -1 at a position i for $t+1 \le i < q$ (and zeros elsewhere) yields eigenvalue $(B-1)R_q^2/\alpha_q$. Note that in the case $t = q$ this accounts for all the eigenvalues. In the case $t < q$ we deduce the remaining eigenvalue by considering the trace of $\mathbf{M}$:

$$t\frac{BR_1^2}{\alpha_1} + (q-t)t\frac{BR_q^2}{\alpha_q} - (t-1)\frac{(B-1)R_1^2}{\alpha_1} - (q-t-1)\frac{(B-1)R_q^2}{\alpha_q} - 1.$$

□

Lemma 73. *The uniform fixed point is stable if* $(\Delta-2)(B-1) < q$. *The uniform fixed point is unstable if* $(\Delta-2)(B-1) > q$.

Proof. The solution of the tree recurrences considered is $R_1 = \cdots = R_q$ and hence $\alpha_1 = \cdots = \alpha_q = (B+q-1)R_1^2$. The only relevant eigenvalue is $(B-1)/(B+q-1)$ (with multiplicity $q-1$), which we compare with $1/(\Delta-1)$ to obtain the result. □

Lemma 73 allows us to restrict our focus on $q-1 \ge t \ge 1$. In this setting, the tree equations give, with $x := y^d := \frac{R_1}{R_q}$ and $d = \Delta - 1$,

$$x = \left(\frac{(B+t-1)x + (q-t)}{tx + (B+q-t-1)}\right)^d \text{ or } B - 1 = \frac{(y-1)\left(ty^d + q - t\right)}{y^d - y}. \tag{185}$$

The following lemma implies that all fixed points with $q-1 \ge t \ge 2$ are unstable in the whole non-uniqueness regime, since the respective matrices have an eigenvalue greater than $1/(\Delta-1)$.

Lemma 74. *When* $\frac{R_1}{R_q} > 1$, *it holds that* $(B-1)\frac{R_1^2}{\alpha_1} > \frac{1}{\Delta-1}$.

Proof. The desired inequality is equivalent to

$$(\Delta-1)(B-1)R_1 > (B+t-1)R_1 + (q-t)R_q,$$

which after simple manipulations reduces into

$$\big((\Delta-2)(B-1)-t\big)\frac{R_1}{R_q} > q-t.$$

Substituting $\frac{R_1}{R_q} = y^d$ and $B-1$ from equation (185), the inequality becomes

$$\left(\frac{(d-1)(y-1)\left(ty^d+q-t\right)}{y^d-y} - t\right) y^d - (q-t) > 0.$$

Doing the necessary simplifications, we obtain the following equivalent inequality

$$\frac{\left((d-1)y^{1+d}-dy^d+y\right)\left(q+t\left(y^d-1\right)\right)}{y^d-y} > 0.$$

Since $y>1$, the only non-trivial factor to prove positivity is $p(y) := (d-1)y^{d+1}-dy^d+y$. By Descartes' rule of signs, p can have at most two positive roots. It holds that $p(1)=p'(1)=0$, so that $p(y)$ is always positive for $y>1$. □

In light of Lemma 74, we need to classify fixpoints with $t=1$, the majority fixpoints. The following lemma gives the number of such majotity fixpoints in the regimes of interest.

Lemma 75. *When $\mathfrak{B}_u \le B < \mathfrak{B}_{rc}$, there are exactly two majority fixpoints. When $B \ge \mathfrak{B}_{rc}$, there is exactly one majority fixpoint.*

Proof. We need to look at (185) for $t=1$ and check how many values of $y>1$ satisfy the equation in the two different regimes. For $t=1$, the equation reads as

$$B-1 = f(y) := \frac{(y-1)(y^d+q-1)}{y^d-y}, \text{ so that } f'(y) = \frac{p(y)}{(y^d-y)^2}, \tag{186}$$

where $p(y)$ is the polynomial

$$p(y) := y^{2d} - dy^{d+1} - (d-1)(q-2)y^d + d(q-1)y^{d-1} - (q-1). \tag{187}$$

Employing the Descartes' rule of signs we see that p has one or three positivite roots counted by multiplicities. It is easy to check that $p(1)=p'(1)=0$, so that p has in fact 3 positive roots (since 1 is a double root), let ρ denote the other positive root. We next prove that $\rho > 1$ so that $p(y) \geq 0$ if $1 \leq y \leq \rho$ and $p(y) \geq 0$ if $y \geq \rho$. It follows that for positive y we have $p(y) > 0$ iff $y > \rho$.

To prove that $\rho > 1$, for the sake of contradiction assume that $0 < \rho \leq 1$. If $\rho = 1$, then 1 is a root with multiplicity 3 of the polynomial $p(y)$ and hence $p''(1) = 0$. By straightforward calculations we see that $p''(1) = (q-2)(d-d^2)$ which is clearly non-zero for $q \geq 3$ and $d \geq 3$. Thus, we may assume that $0 < \rho < 1$. Since $p(1) = p(\rho) = 0$, by Rolle's theorem there is a root $\rho' \in (\rho, 1)$ of the polynomial $p'(y) = dy^{d-2}g(y)$ where

$$g(y) := 2y^{d+1} - (d+1)y^2 - (d-1)(q-2)y + (d-1)(q-1).$$

Since $g(1) = g(\rho') = 0$, by the same token there is a root $\rho'' \in (\rho', 1)$ of

$$g'(y) = 2(d+1)y^d - 2(d+1)y - (d-1)(q-2).$$

We thus obtain the desired contradiction since, for $q \geq 3$ and $d \geq 2$, $g'(y) < 0$ for all $y \in [0,1]$.

Now observe that $f(y) \to \infty$ as $y \to \infty$, while $f(y) \to \frac{q}{d-1}$ as $y \to 1^+$. Thus, when $y \to 1^+$, B, as given by (186), goes to $\mathfrak{B}_{rc}$ from below.

To obtain the lemma, it thus suffices to show that at $y = \rho$, we have $B = \mathfrak{B}_u$. Recall, that $\mathfrak{B}_u$ is the unique value of B for which the polynomial $(q-1)z^{d+1} + (2-B-q)z^d + Bz - 1$ has a double root in $(0,1)$. We reparameterize $z \to 1/z$, so that $\mathfrak{B}_u$ is the unique value of B for which the following polynomial has a double root in $(1, \infty)$:

$$r(z) = z^{d+1} - Bz^d - (2-B-q)z - (q-1).$$

Let z_c be the double root of this polynomial when $B = \mathfrak{B}_u$. Solving each of $r(z_c) = 0$ and $r(z_c) = 0$ with respect to B and equating the expressions, we obtain that $p(z_c) = 0$. It follows that $\rho = z_c$, as wanted. □

We can now classify the stability of fixpoints with $t = 1$.

Lemma 76. *Exactly one majority fixed point is stable. More precisely, the only stable fixed point with $t = 1$ is the one maximizing the ratio x.*

Proof. In the setting of Lemma 71 we have $t = 1$ and thus the interesting eigenvalues of M are R_2^q/α_q and $BR_1^2/\alpha_1 + (B+q-2)R_q^2/\alpha_q - 1$. One can easily check that the former is always larger, so it suffices to check when the following inequality holds

$$Q := BR_1^2/\alpha_1 + (B+q-2)R_q^2/\alpha_q - \frac{\Delta}{\Delta-1} < 0. \tag{188}$$

Expanding everything out, we have

$$\begin{aligned} Q &= \frac{BR_1}{BR_1+(q-1)R_q} + \frac{(B+q-2)R_q}{R_1+(q-2+B)R_q} - \frac{\Delta}{\Delta-1} \\ &= \frac{R_1R_q((\Delta-2)B(q-2+B)-\Delta(q-1)) - BR_1^2 - (q-1)(q-2+B)R_q^2)}{(\Delta-1)(BR_1+(q-1)R_q)(R_1+(q-2+B)R_q)}. \end{aligned}$$

Thus it suffices to check, with $x = \frac{R_1}{R_q}$, when

$$((\Delta-2)B(q-2+B)-\Delta(q-1))x < Bx^2 + (q-1)(q-2+B). \tag{189}$$

Substituting $x = y^d$ and $B-1$ from (185), we obtain the equivalent inequality

$$0 < \frac{y\left(y^d-1\right)\left(y^d+q-1\right)p(y)}{\left(y^d-y\right)^2}.$$

where $p(y)$ is the polynomial given in (187). By the proof of Lemma 75, $p(y) > 0$ iff $y > \rho$. The latter inequality, throughout the regime $B \geq \mathfrak{B}_u$, is only satisfied by the majority fixpoint with x maximum, concluding the proof. $\square$

Having classified the fixpoints which are Jacobian attractive, we now need to see when these are dominant. This entails comparing the values of Ψ_1 for the respective phases. Rather than doing this directly, we use Lemma 5. In particular, it is equivalent to compare the values of Φ_1 at the fixpoints. Moreover, note that the expression (249) is invariant upon scaling R_i's by the same factor and hence we only need to compare $\Phi_1(x, 1, \dots, 1)$ and $\Phi_1(1, \dots, 1)$, where x is a solution of (185) for $t = 1$.

Lemma 77. *Let $t = 1$ and x be the solution of* (185) *with x maximum. Then*

$$\Phi_1(x, 1, \dots, 1) \geq \Phi_1(1, 1, \dots, 1) \textit{ iff } B \geq \mathfrak{B}_o.$$

Equality holds iff $B = \mathfrak{B}_o$.

Proof. By a direct calculation

$$\Phi_1(x,1,\ldots,1) = \frac{\Delta}{2}\log\big((x+q-1)^2+(B-1)(x^2+q-1)\big) - (\Delta-1)\log\big(x^{\Delta/\Delta-1}+q-1\big),$$
$$\Phi_1(1,1,\ldots,1) = \frac{\Delta}{2}\log\big(q^2+(B-1)q\big) - (\Delta-1)\log\big(q\big).$$

Using the substitutions $d=\Delta-1$, $x=y^d$ and the second equation in (185), after careful manipulations we obtain

$$DIF := \Phi_1(x,1,\ldots,1) - \Phi_1(1,1,\ldots,1) = \log\left(\frac{\big(y^d+q-1\big)^{d+1}}{\big(y^{d+1}+q-1\big)^{d-1}}/\frac{(q+y-1)^{d+1}}{q^{d-1}}\right).$$

It is straightforward to check that for $y=(q-1)^{2/(d+1)}$, $DIF=0$. The respective value of B for this value of y is given by the second equation in (185) and equals $\mathfrak{B}_o$. Thus, it suffices to show that y is an increasing function of B and DIF increases as y increases. This is indeed true. By (185), one calculates

$$\frac{\partial y}{\partial B}\cdot\frac{p(y)}{y(y^d-y)^2}=1, \quad \text{and} \quad \frac{\partial DIF}{\partial y} = \frac{(d+1)(q-1)p(y)}{y(y+q-1)(y^d+q-1)(y^{d+1}+q-1)},$$

where $p(y)$ is the polynomial defined in Lemma 74, whose positivity has already been established. The claim follows. □

Proof of Theorem 69. We first argue about the local maxima of the function Ψ_1. To use the results proved in this section, we apply Theorem 2: we just need to check the stability of the corresponding fixpoints. As proved in Lemma 74, only the disordered phase and the q ordered phases can be local maxima. The disordered phase, by Lemma 73, is Jacobian stable when $1<B\le\mathfrak{B}_{rc}$. The q ordered phases (which are permutations of each other), by Lemmas 74 and 76, are Jacobian stable when $B\ge\mathfrak{B}_u$.

To argue about the dominant phases, we need to find the regimes where the disordered/ordered phases are dominant. Lemma 77 says that the disordered phase is dominant iff $B\le\mathfrak{B}_o$, whereas the q ordered phases iff $B\ge\mathfrak{B}_o$. To finish the proof of Theorem 69, we only need to argue that the phases are also Hesian dominant in the respective regimes. This follows immediately by Theorem 2, since Jacobian stability is equivalent to a Hessian phase. Thus, since we classified local maxima of Ψ_1 using the Jacobian stability of the corresponding fixpoints, the statements about the Hessian phases follow from the preceding discussion. □

6.3 *Antiferromagnetic Potts model and colorings*

In this section, we establish the dominant phases for the antiferromagnetic Potts model and colorings and hence prove Theorem 16. We also verify the hypotheses of Theorem 14 for these two models and hence complete the proofs of Theorems 12 and 13.

6.3.1 Proof outline

To obtain Theorems 12 and 13, we verify the hypotheses of Theorem 14 for the dominant phases of the antiferromagnetic Potts and colorings models on random Δ-regular bipartite graphs. Recall, the interaction matrix $\mathbf{B}$ for the Potts model is completely determined by a parameter B, which is equal to $\exp(-\beta)$ where β is the inverse temperature in the standard notation for the Potts model. The antiferromagnetic regime corresponds to $0 < B < 1$. The coloring model is the zero temperature limit of the Potts model and corresponds to the particular case $B = 0$ in what follows. We should note that in Statistical Physics terms, the arguments of this section are closely related to the phase diagrams of the models.

The most crucial component is to obtain characterizations of the global maxima of Ψ_1. To be able to apply Theorem 14, our goal is to prove that the global maxima of Ψ_1 in the regime $0 \leq B < \frac{\Delta-q}{\Delta}$ are (i) not translation-invariant, (ii) Hessian maxima, and (iii) permutation symmetric, see Section 1.4.3 for the definitions of these terms. We will in fact show the stronger statement that condition (i) is met only in the regime $0 \leq B < \frac{\Delta-q}{\Delta}$, which shows that the inapproximability results of Theorems 12 and 13 are best possible for any reduction which uses bipartite graphs as gadgets. We next outline our methods.

By Theorem 5 specified to the antiferromagnetic Potts and colorings models, studying the global maxima of Ψ_1 is equivalent to studying the global maxima of Φ. Moreover, the global maxima of Φ and Ψ_1 occur at their critical points. Since there is a one-to-one correspondence between the critical points of Φ and the critical points of Ψ_1 (given by (11)), we will freely interchange our focus between critical points of Φ and Ψ_1.

The critical points of Φ, by the first part of Theorem 5, are given by fixpoints of the tree recursions (8), which for the Potts model are positive solutions to (18) stated in the

Introduction. For convenience, we restate the equations (18):

$$R_i \propto \Big(BC_i + \textstyle\sum_{j\neq i} C_j\Big)^d, \quad C_j \propto \Big(BR_j + \textstyle\sum_{i\neq j} R_i\Big)^d, \tag{18}$$

where $i, j = 1, \dots, q$ and d is the notational convenient substitution $d := \Delta - 1 \geq 2$. Given a fixpoint of the tree recursions (18), we will classify whether it is a Hessian local maximum of Ψ_1 using Theorem 2.

Once we find the global maxima of Ψ_1, it will be simple to prove that they are Hessian and permutation symmetric. Finding however the global maxima of Ψ_1 is going to be more intricate, mainly because the number of local maxima varies according to the value of B. We will thus have to compare the values of Ψ_1 at the critical points. Rather than doing this directly (which seems as a difficult task), we solve a relaxed optimisation problem, which for q even can be tied to the maximization of Ψ_1. We next give the details.

We begin our considerations by examining when a fixpoint (18) is translation invariant, i.e., satisfies $R_i \propto C_i$ for every $i \in [q]$.

Lemma 78. *Let $0 \leq B < 1$ and $\Delta \geq 3$. Then a solution of* (18) *satisfies $R_i \propto C_i$ for every $i \in [q]$ iff $R_1 = \ldots = R_q$ and $C_1 = \ldots = C_q$.*

Proof of Lemma 78. We first prove the forward direction. By the symmetries of the model, we may assume an arbitrary ordering of the R_i's. Since $0 \leq B < 1$, equations (18) easily imply the reverse ordering of the C_i's. Thus, $R_i \propto C_i$ for every $i \in [q]$ implies that the ordering must be trivial, i.e, $R_1 = \ldots = R_q$ and $C_1 = \ldots = C_q$. The backward direction is trivial. □

Corollary 79. *Translation invariant fixpoints* (18) *always exist and are unique up to scaling.*

We next explore in which regimes of B, the critical points of Φ consist solely of translation invariant fixpoints. In this regime, we immediately obtain by Theorem 5 that the global maximum of Ψ_1 (and hence the global maximum of Φ as well) is achieved at a translation invariant fixpoint.

Lemma 80. *Let $0 \leq B < 1$ and $\Delta \geq 3$. When $q \geq \Delta$ and $0 \leq B < 1$ or $q < \Delta$ and $\frac{\Delta - q}{\Delta} \leq B < 1$, the system of equations* (18) *admits a positive solution iff $R_1 = \ldots = R_q$ and $C_1 = \ldots = C_q$.*

The proof of Lemma 80 is an extension of an argument in [11] for colorings and is given in Section 6.3.5.1. The next lemma states that in the complementary regime of Lemma 80, the translation invariant fixpoint does not correspond to a local maximum of Ψ_1 and hence, by Theorem 5, the global maximum of Ψ_1 occurs at a fixpoint of (18) which is not translation invariant. In particular, in this regime we have semi-translational non-uniqueness.

Lemma 81. *For $q < \Delta$ and $0 \leq B < \frac{\Delta - q}{\Delta}$, the global maximum of Ψ_1 is not achieved at the translation invariant fixpoint.*

Proof of Lemma 81. We apply Theorem 2 by showing that the translation invariant fixpoint is Jacobian unstable and hence not a local maximum of Ψ_1. By Lemma 33, for a general interaction matrix $\mathbf{B}$, the condition for Jacobian stability of a fixpoint of the tree recursions is related to the spectrum of $\mathbf{L} = \left[\begin{smallmatrix} \mathbf{0} & \mathbf{A} \\ \mathbf{A}^{\intercal} & \mathbf{0} \end{smallmatrix}\right]$, where $\mathbf{A}$ is the $q \times q$ matrix whose ij-entry is given by $A_{ij} = B_{ij} R_i C_j / \sqrt{\alpha_i \beta_j}$ and α_i, β_j are given by (43). Recall that ± 1 are eigenvalues of $\mathbf{L}$ and the condition for Jacobian stability is that all the other eigenvalues have absolute value less than $1/(\Delta - 1)$ (see for details the proof of Theorem 2 in Section 3.4.3).

In the setting of the lemma, the matrix $\mathbf{A}$ for the translation invariant fixpoint has off-diagonal entries equal to $1/(B + q - 1)$ and diagonal entries equal to $B/(B + q - 1)$. It follows that the eigenvalues of $\mathbf{L}$ are ± 1 by multiplicity 1 and $\pm(1 - B)/(B + q - 1)$ by multiplicity $q - 1$. The absolute value of the latter is greater than $\frac{1}{\Delta - 1}$ for $0 \leq B < \frac{\Delta - q}{\Delta}$, as claimed. □

We summarize the above results into the following corollary.

Corollary 82. *When $q \geq \Delta$ and $0 \leq B < 1$ or $q < \Delta$ and $\frac{\Delta - q}{\Delta} \leq B < 1$, Ψ_1 has a unique global maximum for $\alpha_1 = \ldots = \alpha_q = \beta_1 = \ldots = \beta_q = 1/q$ or, in other words, the global maximum of Ψ_1 is achieved by the fixpoint which corresponds to the (unique) translation invariant Gibbs measure. In the complementary regime $q < \Delta$ and $0 \leq B < \frac{\Delta - q}{\Delta}$, the* *maximum of Ψ_1 is not achieved at the translation invariant fixpoint, and hence it is achieved at a semi-translation invariant fixpoint which is not translation invariant.*

Corollary 82 is not sufficient to obtain Theorems 12 and 13, since we need to verify that the global maxima of Ψ_1 in semi-translational non-uniqueness are Hessian and permutation symmetric. We do this by identifying the critical points which are maxima of Ψ_1.

To state the result, we first need the following structural statement for the solutions of equations (18), namely that solutions of (18) are supported on at most 3 values for the R_i's and similarly for the C_i's.

Lemma 83. *Let* $(R_1, \ldots, R_q, C_1, \ldots, C_q)$ *be a positive solution of the system* (18). *Let* t_R *be the number of values on which the* R_i*'s are supported and define similarly* t_C. *Then* $t_R, t_C \leq 3$ *and* $t_R = t_C =: t$.

The proof of Lemma 83 is given in Section 6.3.5. Lemma 83 motivates the following definition.

Definition 12. *From Lemma 83, the* R_i*'s and* C_j*'s of a fixpoint of* (18) *attain at most* $t \leq 3$ *different values. Let* $\tilde{R}_1, \ldots, \tilde{R}_t$ *and* $\tilde{C}_1, \ldots, \tilde{C}_t$ *be their values and let* $q_1, \ldots, q_t \geq 1$ *be their multiplicities. When* $t = 1$, *define* $q_2 = q_3 = 0$; *when* $t = 2$, *define* $q_3 = 0$; *when* $q_i = 0$, *define the values of* $\tilde{R}_i, \tilde{C}_i$ *to be zero. The corresponding solution of* (18) *or equivalently the fixpoint of the tree recursions is then defined to be of type* (q_1, q_2, q_3). *Note that* $q_1 + q_2 + q_3 = q$ *and the* q_i*'s are non-negative integers. Call a* (q_1, q_2, q_3)*-type fixpoint to be* t*-supported if the number of* q_i*'s which are non-zero equals* t.

Finding the types of fixpoints which correspond (via (11)) to global maxima of Ψ_1 is not a trivial task. While 2-supported fixpoints are simple to handle, this is not the case for 3-supported fixpoints. The main lemma we prove is the following, which identifies the type of fixpoints which maximize Ψ_1.

Lemma 84. *For* $0 \leq B < \frac{\Delta - q}{\Delta}$ *and even* $q \geq 3$, *the maximum of* Ψ_1 *over* (q_1, q_2, q_3)*-type solutions of* (18) *is attained at fixpoints of type* $(q/2, q/2, 0)$.

The final piece is to show that fixpoints of type $(q/2, q/2, 0)$ are Hessian maxima of Ψ_1 and permutation symmetric. This is the scope of the next lemma, whose proof is given in Section 6.3.5.

Lemma 85. *For $0 \leq B < \frac{\Delta - q}{\Delta}$ and even $q \geq 3$, fixpoints of type $(q/2, q/2, 0)$ are Jacobian stable and hence correspond to Hessian maxima of Ψ_1. The values of R_i's and C_j's for fixpoints of type $(q/2, q/2, 0)$ are unique up to scaling and permutations of the colours.*

We are now ready to prove Theorem 16.

Proof of Theorem 16. Items 1 and 2 follow from Corollary 82 (see also Lemmas 80 and 81). Item 3 follows from Lemmas 85, after using the correspondence between fixpoints of the tree recursions (18) and dominant phases of Theorem 5 (equation (11)). □

With Lemmas 84 and 85 at hand, it is also straightforward to obtain Theorems 12 and 13 by applying the inapproximability result of Theorem 14.

Proof of Theorems 12 and 13. By Theorem 14, it suffices to check that the dominant phases of Ψ_1 are permutation symmetric and Hessian. By Theorem 5, the maximum of Ψ_1 happens at a critical point, which correspond (via (11)) to fixpoints (18). By Lemma 84, the fixpoints (18) which correspond to maximizers of Ψ_1 are of type $(q/2, q/2, 0)$. Note that the scaling of R_i's and C_j's does not affect the values of α_i's and β_j's in Theorem 5. It follows by Lemma 85 that the maximizers of Ψ_1 are permutation symmetric and correspond to Hessian maxima of Ψ_1. This verifies the hypotheses of Theorem 14. □

6.3.2 Proof of Lemma 84

In this section, we outline the proof of Lemma 84. We need to find the type(s) of the fixpoints which maximize Ψ_1. Let $\mathbf{q} = (q_1, q_2, q_3)$ specify the type of a fixpoint of (18) and let $\mathbf{r} = (R_1, R_2, R_3)$, $\mathbf{c} = (C_1, C_2, C_3)$ be the respective values of the R_i's and C_j's, see Definition 12. Note that the q_i's are non-negative integers satisfying $q_1 + q_2 + q_3 = q$.

Using Theorem 5, we obtain that the value of $\Psi_1(\boldsymbol{\alpha}, \boldsymbol{\beta})$ corresponding to this fixpoint of (18) is given by the value of the function $\overline{\Phi^S}$, where

$$\begin{aligned}\overline{\Phi^S}(\mathbf{q},\mathbf{r},\mathbf{c}) := (d+1)\ln\Big(\textstyle\sum_{i=1}^3 q_iR_i\sum_{j=1}^3 q_jC_j + (B-1)\sum_i q_iR_iC_i\Big)&\\ - d\ln\Big(\textstyle\sum_{i=1}^3 q_iR_i^{(d+1)/d}\Big) - d\ln\Big(\textstyle\sum_{j=1}^3 q_jC_j^{(d+1)/d}\Big),&\end{aligned} \tag{190}$$

and $d = \Delta - 1$. It is a non-trivial task to directly compare the values of $\overline{\Phi^S}$ over fixpoints of (18). Instead, we will solve a relaxed version of the problem, seeking to maximize $\overline{\Phi^S}$ over non-negative q_i's which satisfy $q_1 + q_2 + q_3 = q$. If this maximum happens to occur for integer $\mathbf{q}$ and the respective values of R_i's and C_j's are solutions of (18), then we have also found the solution to the original maximization problem. It turns out that all of the above are satisfied iff q is even.

To formalize the argument, for non-negative q_i's such that $q_1 + q_2 + q_3 = q$, define

$$\overline{\Phi}(\mathbf{q}) := \max_{\mathbf{r},\mathbf{c}} \overline{\Phi^S}(\mathbf{q},\mathbf{r},\mathbf{c}) \tag{191}$$

where the maximum is over $\mathbf{r} = (R_1, R_2, R_3)^\intercal$, $\mathbf{c} = (C_1, C_2, C_3)^\intercal$ which satisfy

$$\begin{gathered}\textstyle\sum_{i=1}^3 q_iR_i\sum_{j=1}^3 q_jC_j + (B-1)\sum_{i=1}^3 q_iR_iC_i > 0,\\ R_1, R_2, R_3, C_1, C_2, C_3 \geq 0.\end{gathered} \tag{192}$$

It is simple to see that in the region (192), $\overline{\Phi^S}$ is well defined. It is not completely immediate that the maximum in (191) is well defined since the region (192) is not compact. This is a consequence of the following scale-free property of $\overline{\Phi^S}$ with respect to $\mathbf{r}$ and $\mathbf{c}$:

$$\text{for every } c_1, c_2 > 0 \text{ it holds that } \overline{\Phi^S}(\mathbf{q}, c_1\mathbf{r}, c_2\mathbf{c}) = \overline{\Phi^S}(\mathbf{q},\mathbf{r},\mathbf{c}). \tag{193}$$

Using (193), it is simple to obtain the following.

Lemma 86. *Let $B \geq 0$ and $q \geq 2$. For all $q_1, q_2, q_3 \geq 0$ which satisfy $q_1 + q_2 + q_3 = q$, the maximum in* (191) *is well defined. Moreover, the maximum of $\overline{\Phi}(q_1, q_2, q_3)$ over all such q_1, q_2, q_3 is attained.*

We next seek to connect the maximizers of (191) with solutions of (18). To do this, we first need to consider whether the maximum in (191) happens on the boundary of the region (192); it turns out that the maximum can happen at the boundary $R_i = 0$ or $C_i = 0$ if q_i

is close to zero. While the boundary cases are an artifact of allowing q_i's to be non-integer, we will need to treat them explicitly to find the maximum of $\overline{\Phi}$.

Definition 13. *A triple* $\mathbf{q} = (q_1, q_2, q_3)$ *is* good *if the* $\mathbf{r}, \mathbf{c}$ *which achieve the maximum in* (191) *satisfy: for* $i = 1, 2, 3$*,* $q_i > 0$ *implies* $R_i, C_i > 0$*. A triple* $\mathbf{q} = (q_1, q_2, q_3)$ *is* bad *if it is not good.*

To complete the connection, we need to further restrict the set of triples $\mathbf{q}$. To motivate this restriction, note that if we consider the region (192) in the subspace $R_1 = R_2$ and $C_1 = C_2$, we obtain $\overline{\Phi}(q_1 + q_2, q_3, 0) \le \overline{\Phi}(q_1, q_2, q_3)$. To avoid degenerate cases, we consider only triples $\mathbf{q}$ where such simple inequalities do not hold at equality.

Definition 14. *Let* $t = 2$ *or* 3*. A triple* $\mathbf{q} = (q_1, q_2, q_3)$ *is called* t*-maximal if exactly* t *of the* q_i*'s are non-zero and for all distinct* $i, j, k \in \{1, 2, 3\}$ *it holds that* $\overline{\Phi}(q_i + q_j, q_k, 0) < \overline{\Phi}(\mathbf{q})$*.*

Our interest is in maximal good triples $\mathbf{q} = (q_1, q_2, q_3)$. This is justified by the following lemma, whose proof is given in Section 6.3.5.

Lemma 87. *Suppose that* q_1, q_2, q_3 *are non-negative* integers *and the triple* $\mathbf{q} = (q_1, q_2, q_3)$ *is* t*-maximal and good. Then, the* $\mathbf{r}, \mathbf{c}$ *which achieve the maximum in* (18) *specify a* t*-supported fixpoint of* (18) *of type* (q_1, q_2, q_3)*.*

Thus to prove Lemma 84, it suffices to prove that the triple $(q/2, q/2, 0)$ is 2-maximal and good and that the maximum of $\overline{\Phi}(\mathbf{q})$ is achieved at $(q/2, q/2, 0)$. The next lemma examines which maximal good triples can be a maximum of $\overline{\Phi}$.

Lemma 88. *Let* $q \ge 3$ *and* $0 \le B < 1$*. There do not exist 3-maximal good triples* $\mathbf{q}$ *which maximize* $\overline{\Phi}(\mathbf{q})$*. The only 2-maximal good triples* $\mathbf{q}$ *where a maximum of* $\overline{\Phi}(\mathbf{q})$ *can occur are* $(q/2, q/2, 0)$ *or its permutations.*

Lemma 88 is not sufficient to yield Lemma 84 because the maximum of $\overline{\Phi}(\mathbf{q})$ can occur at a bad triple $\mathbf{q}$. This possibility is excluded by the following lemma.

Lemma 89. *Let* $q \ge 3$ *and* $0 \le B < \frac{\Delta - q}{\Delta}$*. There do not exist bad triples* $\mathbf{q}$ *which maximize* $\overline{\Phi}(\mathbf{q})$*.*

Using Lemmas 88 and 89, we can now give the proof of Lemma 84.

Proof of Lemma 84. The maximum of $\overline{\Phi}(\mathbf{q})$ over triples $\mathbf{q}$ is attained by Lemma 86. This maximum can happpen either at a bad or a good triple $\mathbf{q}$. Maxima at bad triples $\mathbf{q}$ are excluded by Lemma 89. Maxima at 3-maximal good triples are excluded by the first part of Lemma 88. Thus, the maximum must happen at a (good) triple of the form $\mathbf{q} = (q_1, q_2, 0)$. The latter can be either 2-maximal or not. If it is not 2-maximal, the maximum must equal $\overline{\Phi}(\mathbf{q})$, which in the regime $0 \leq B < \frac{\Delta - q}{\Delta}$ is excluded by Lemma 81. Thus, the maximum must happen at a 2-maximal good triple, which Lemma 88 asserts that it must be the triple $(q/2, q/2, 0)$. Finally, for q even, by Lemma 87 the $\mathbf{r}, \mathbf{c}$ which achieve the maximum in (191) correspond to a 2-supported fixpoint of (18) of type $(q/2, q/2, 0)$, as wanted. $\square$

For the proofs of Lemmas 88 and 89, we will often perturb the values of q_i's. The following lemma, which is proved in Section 6.3.5 will be very helpful.

Lemma 90. *Let* $\mathbf{q} = (q_1, q_2, q_3)$ *and* $I = \{i \mid q_i > 0\}$*. Suppose that* $\mathbf{r}, \mathbf{c}$ *achieve the maximum in* (191)*. Then, for* $i \in I$ *it holds that*

$$\frac{\partial \overline{\Phi^S}}{\partial q_i}(\mathbf{q}, \mathbf{r}, \mathbf{c}) = \frac{R_i \sum_j q_j C_j + C_i \sum_j q_j R_j + (d-1)(1-B) R_i C_i}{\sum_j q_j R_j \sum_j q_j C_j + (B-1) \sum_j q_j R_j C_j}. \tag{194}$$

Moreover, if there exist $i, j \in I$ *such that* $\frac{\partial \overline{\Phi^S}}{\partial q_i} - \frac{\partial \overline{\Phi^S}}{\partial q_j} \neq 0$*, the maximum of* $\overline{\Phi}$ *is* not *achieved at the triple* $\mathbf{q}$.

6.3.3 Good triples: proof of Lemma 88

We first prove the statement of the lemma for 3-maximal good triples $\mathbf{q} = (q_1, q_2, q_3)$, the proof for 2-maximal good triples will easily be inferred by appropriately modifying the arguments in the special case $q_2 = 0$.

Let $\mathbf{q} = (q_1, q_2, q_3)$ be a 3-maximal good triple. Since $\mathbf{q}$ is 3-maximal all of the q_i's are positive. Moreover, $\mathbf{q}$ is good, and hence the maximum in (191) for $\mathbf{q}$ is attained at positive R_i's and C_j's. Thus, the R_i's and C_j's satisfy $\partial \overline{\Phi^S} / \partial R_i = \partial \overline{\Phi^S} / \partial C_j = 0$ which give

$$R_i^{1/d} \propto q_1 C_1 + q_2 C_2 + q_3 C_3 + (B-1) C_i, \quad C_j^{1/d} \propto q_1 R_1 + q_2 R_2 + q_3 R_3 + (B-1) R_j. \tag{195}$$

Since $\mathbf{q}$ is 3-maximal, we may assume that $\mathbf{r}$ is such that $R_i \neq R_j$ for all $i \neq j$. Otherwise, if for example $R_1 = R_2$, by (195), we have $C_1 = C_2$ as well, so that $\overline{\Phi}(q_1, q_2, q_3) = \overline{\Phi}(q_1 + q_2, q_3, 0)$, contradicting the 3-maximality of $\mathbf{q}$. Thus, we may assume a strict ordering of the R_i's, which by (195) implies the reverse ordering of the C_j's. Wlog, we will use the following ordering:

$$R_1 > R_2 > R_3 > 0 \quad \text{and} \quad 0 < C_1 < C_2 < C_3. \tag{196}$$

The following lemma, together with the second part of Lemma 90, establishes that the maximum of $\overline{\Phi}$ cannot occur at a 3-maximal triple.

Lemma 91. *Suppose that R_i's and C_j's satisfy* (195) *and* (196)*. If $R_1/R_3 \neq C_3/C_1$ then $\frac{\partial \overline{\Phi^S}}{\partial q_1} - \frac{\partial \overline{\Phi^S}}{\partial q_3} \neq 0$. If $R_1/R_3 = C_3/C_1$ then $\frac{\partial \overline{\Phi^S}}{\partial q_1} - \frac{\partial \overline{\Phi^S}}{\partial q_2} \neq 0$.*

We next give the proof of Lemma 91. We will utilize Lemma 90 by specifying a particular scaling of the R_i's and C_j's which will be beneficial. To do this, set

$$r_1^d = R_1/R_3, r_2^d = R_2/R_3, c_2^d = C_2/C_1, c_3^d = C_3/C_1. \tag{197}$$

The R_i's and C_j's may be recovered from r_i's, c_j's using

$$R_1 \propto r_1^d,\ R_2 \propto r_2^d,\ R_3 \propto 1, \quad \text{and} \quad C_1 \propto 1,\ C_2 \propto c_2^d,\ C_3 \propto c_3^d. \tag{198}$$

Translating (196) into r_1, r_2, c_2, c_3 gives

$$r_1 > r_2 > 1 \text{ and } c_3 > c_2 > 1. \tag{199}$$

Moreover, dividing appropriate pairs of (195), we also obtain

$$\begin{aligned} r_1 &= \frac{B + q_1 - 1 + q_2 c_2^d + q_3 c_3^d}{q_1 + q_2 c_2^d + (B + q_3 - 1) c_3^d}, \ c_3 = \frac{B + q_3 - 1 + q_2 r_2^d + q_1 r_1^d}{q_3 + q_2 r_2^d + (B + q_1 - 1) r_1^d}, \\ r_2 &= \frac{q_1 + (B + q_2 - 1) c_2^d + q_3 c_3^d}{q_1 + q_2 c_2^d + (B + q_3 - 1) c_3^d}, \ c_2 = \frac{q_3 + (B + q_2 - 1) r_2^d + q_1 r_1^d}{q_3 + q_2 r_2^d + (B + q_1 - 1) r_1^d}. \end{aligned} \tag{200}$$

It can easily be verified that this system of equations gives

$$q_1 = \frac{(1-B) f(r_1, c_3) + q_2 P \left(c_2^d - c_3^d r_2^d\right)}{P\left(r_1^d c_3^d - 1\right)}, \ q_3 = \frac{(1-B) f(c_3, r_1) + q_2 P \left(r_2^d - r_1^d c_2^d\right)}{P\left(r_1^d c_3^d - 1\right)} \tag{201}$$

$$r_2 = \frac{r_1 c_3^d - 1 - c_2^d (r_1 - 1)}{c_3^d - 1}, \ r_2^d = \frac{r_1^d c_3 - 1 - c_2 (r_1^d - 1)}{c_3 - 1}, \tag{202}$$

$$f(x, y) := x^{d+1} y^{d+1} - x^d y^{d+1} - x y^{d+1} + y^d + y - 1, \ P := (r_1 - 1)(c_3 - 1) > 0.$$

We will need the following lemma.

Lemma 92. *Assume that* $q_1, q_2, q_3, r_1, r_2, c_2, c_3$ *satisfy* (199), (201), (202). *If* $r_1 = c_3$ *then* $r_2 = c_2$ *and* $q_1 = q_3$.

Proof of Lemma 92. We prove that $r_1 = c_3$ implies $r_2 = c_2$. Once this is done, (201) easily gives that $r_1 = c_3$ implies $q_1 = q_3$ as well, thus proving the lemma.

So, suppose that $z = r_1 = c_3$ and for the sake of contradiction assume $r_2 \neq c_2$. By (199) we obtain that $r_2, c_2 \in (1, z)$. Eliminating r_2 from (202) we obtain that c_2 (and by a symmetric argument r_2) satisfies

$$g(s) := \left(\frac{z^{d+1} - 1 - s^d(z-1)}{z^d - 1}\right)^d + \frac{s(z^d - 1) - (z^{d+1} - 1)}{z - 1} = 0.$$

In fact, $g(1) = g(z) = 0$ as well, so that g has at least four distinct roots in $[1, z]$. It follows that $g'(s) = 0$ has at least three distinct solutions in $[1, z]$, say s_i for $i = 1, 2, 3$. As a consequence of $g'(s_i) = 0$, we easily obtain that the s_i's satisfy $h(s_i) = c$ where $h(s) := (z^{d+1} - 1)s - s^{d+1}(z - 1)$ and c is a constant which depends only on z, d. Thus, $h'(s) = 0$ has at least two distinct solutions in $[1, z]$ which is clearly absurd. $\square$

Proof of Lemma 91. Set

$$DIF_{13} := \frac{\partial \overline{\Phi^S}}{\partial q_1} - \frac{\partial \overline{\Phi^S}}{\partial q_3},\ DIF_{12} := \frac{\partial \overline{\Phi^S}}{\partial q_1} - \frac{\partial \overline{\Phi^S}}{\partial q_2},\ S := \textstyle\sum_i q_i R_i \sum_j q_j C_j + (B-1) \sum_i q_i R_i C_i.$$

We use the expressions (194) for the derivatives. The denominators in the expressions are the same, so we may ignore them. Moreover, the expressions therein are scale-free, consequently in order to write the derivatives with respect to r_i's and c_j's we just need to make the substitutions (198).

To prove the first part of the lemma, we eliminate q_1, q_3 from the resulting expression for DIF_{13} using (201). This substitution has the beneficial effect of eliminating q_2, r_2 from the final expression. After straightforward calculations, we obtain the following:

$$\begin{aligned} DIF_{13} &= -\frac{(1-B)\, g(r_1, c_3)}{S(r_1 - 1)(c_3 - 1)}, \quad \text{where} \\ g(r_1, c_3) &:= (r_1 - c_3)(r_1^d - 1)(c_3^d - 1) - d(r_1 - 1)(c_3 - 1)(r_1^d - c_3^d). \end{aligned} \tag{203}$$

It can easily be seen that for $r_1, c_3 > 1$, it holds that $g(r_1, c_3) = 0$ iff $r_1 = c_3$ iff $R_1/R_3 = C_3/C_1$ as desired.

We next prove the second part of the lemma. Since $R_1/R_3 = C_3/C_1$, we have $r_1 = c_3$ and by Lemma 92, $r_2 = c_2$ and $q_1 = q_3$. Using these, (201) and (202) simplify to

$$q_1 = \frac{(1-B)\left(r_1^{d+1}-1\right) - q_2 r_2^d (r_1 - 1)}{(r_1-1)\left(r_1^d+1\right)}, \; r_2 = \frac{r_1^{d+1} - 1 - r_2^d(r_1-1)}{r_1^d - 1}. \tag{204}$$

Moreover, using the substitutions (198) and $q_1 = q_3$, we obtain

$$\begin{aligned} DIF_{12} &= \frac{(q_1 r_1^d + q_2 r_2^d + q_1)(r_1^d - 2r_2^d + 1) + (d-1)(1-B)(r_1^d - r_2^{2d})}{S} \\ &= -\frac{(1-B)\left[(d-1)(r_1-1)r_2^{2d} + 2r_2^d(r_1^{d+1}-1) - (r_1^{2d+1} + dr_1^{d+1} - dr_1^d - 1)\right]}{(r_1-1)S}, \end{aligned}$$

where in the second equality we substituted the value of q_1 from (204). Observe that the numerator is a quadratic polynomial in r_2^d and, by inspection, for $r_1 > 1$, its roots are of opposite sign. Thus, $DIF_{12} = 0$ iff $r_2^d = \rho_1$, where

$$\rho_1(r_1) := \frac{\sqrt{D} - (r_1^{d+1}-1)}{(d-1)(r_1-1)} \text{ and } D := \left(dr_1^{d+1} - (d-1)r_1^d + 1\right)\left(r_1^{d+1} + (d-1)r_1 - d\right).$$

For the sake of contradiction, suppose that $r_2^d = \rho_1$. Then (204) gives that $r_2 = \rho_2$, where

$$\rho_2(r_1) := \frac{d(r_1^{d+1}-1) - \sqrt{D}}{(d-1)(r_1^d-1)}.$$

Thus $\rho_1 = \rho_2^d$. We obtain a contradiction by showing that for every $r_1 > 1$, it holds that $\rho_2^d < \rho_1$ or equivalently $d \ln \rho_2 < \ln \rho_1$. It is easy to see that in the limit $r_1 \downarrow 1$ the inequality is satisfied at equality, thus it suffices to prove that the derivative of the rhs w.r.t r_1 is greater than the respective derivative of the l.h.s. for $r_1 > 1$.

This differentiation is cumbersome but otherwise straightforward. The final result is

$$\frac{1}{\rho_1}\frac{\partial \rho_1}{\partial r_1} - \frac{d}{\rho_2}\frac{\partial \rho_2}{\partial r_1} = \frac{(d+1)g(r_1)h(r_1)}{2(r_1-1)\left(r_1^d-1\right)\left(\sqrt{D} - (r_1^{d+1}-1)\right)\left(d(r_1^{d+1}-1) - \sqrt{D}\right)}, \tag{205}$$

$$\begin{aligned} g(r_1) &:= r_1^{2d} - d^2 r_1^{d+1} + 2(d^2-1)r_1^d - d^2 r_1^{d-1} + 1, \\ h(r_1) &:= (d+1)(r_1^{d+1}-1) - (d-1)(r_1^d - 1) - 2\sqrt{D}. \end{aligned}$$

Note that the denominator in the r.h.s. of (205) is positive for $r_1 > 1$: the terms involving $\sqrt{D}$ are positive since they are the numerators of $\rho_1,\ \rho_2$. The final part of the proof consists of proving that $g(r_1) > 0$ and $h(r_1) > 0$ for $r_1 > 1$.

The polynomial g has 4 sign changes and hence, by the Descartes' rule of signs has at most 4 positive roots. In fact, a tedious calculation shows that $r_1 = 1$ is a root by multiplicity 4, thus proving that $g(r_1) > 0$ for $r_1 > 1$. To prove that $h(r_1) > 0$ for $r_1 > 1$, note the identity

$$\left[(d+1)(r_1^{d+1}-1)-(d-1)(r_1^d-1)\right]^2 - 4D = (d-1)^2(r_1-1)^2(r_1^d-1)^2.$$

This completes the proof. $\square$

To prove the second part of Lemma 88, assume that $\mathbf{q} = (q_1, q_2, q_3)$ is a 2-maximal good triple. Since $\mathbf{q}$ is 2-maximal, w.l.o.g. we may assume that $q_2 = 0$. Note that the values of R_2, C_2 do not affect the value of the derivatives $\partial\overline{\Phi^S}/\partial q_1, \partial\overline{\Phi^S}/\partial q_3$ when $q_2 = 0$. Similarly, (201) continues to hold even when $q_2 = 0$. Thus, the proof of the first part of Lemma 91 carries through verbatim. In particular, if $R_1/R_3 \neq C_3/C_1$, then $\partial\overline{\Phi^S}/\partial q_1 - \partial\overline{\Phi^S}/\partial q_3 \neq 0$. By the second part of Lemma 90, it follows that $\mathbf{q} = (q_1, 0, q_3)$ cannot be a maximum unless $R_1/R_3 = C_3/C_1$. In this case, (201) gives $q_1 = q_3$. Since $q_1 + q_3 = q$, we obtain that the only 2-maximal good triples where the maximum of $\overline{\Phi}$ may occur are $(q/2, q/2, 0)$ or its permutations, as desired.

This concludes the proof of Lemma 88.

6.3.4 Bad triples: proof of Lemma 89

To get a handle on bad triples, we first give necessary conditions so that the maximum in (191) happens at the boundary. The proof of the following lemma is given in Section 6.3.5.

Lemma 93. *Let $0 \leq B < 1$. For a triple $\mathbf{q} = (q_1, q_2, q_3)$, let $\mathbf{r}, \mathbf{c}$ achieve the maximum in* (191)*. Then, if $q_i > 0$, the following implications hold:*

$$R_i = 0 \Rightarrow \textstyle\sum_j q_j C_j \leq (1-B)C_i, \quad C_i = 0 \Rightarrow \textstyle\sum_j q_j R_j \leq (1-B)R_i.$$

In particular, if $q_i > 1 - B$ it holds that $R_i, C_i > 0$ and hence for every $q \geq 3$ there exists i such that $R_i, C_i > 0$.

We next examine bad triples. Note that a bad triple $\mathbf{q} = (q_1, q_2, q_3)$, by the second part of Lemma 90, must have at least two positive entries. We consider cases whether the triple

$\mathbf{q}$ has two or three positive entries. We start with the case where exactly two of the q_i's are positive. We assume throughout the rest of the section that $\mathbf{r}, \mathbf{c}$ achieve the maximum in (191).

Let $\mathbf{q} = (q_1, q_2, 0)$ be a bad triple where $q_1, q_2 > 0$. Since $\mathbf{q}$ is bad, at least one of R_1, R_2, C_1, C_2 is zero. Wlog, we may assume $C_2 = 0$. By the second part of Lemma 93, it follows that $R_1, C_1 > 0$. There are two cases to consider.

$$\text{(I) } R_2 = 0, \text{ (II) } R_2 > 0. \tag{206}$$

Case (I) is straightforward: by the first part of Lemma 90, we trivially have $\frac{\partial \overline{\Phi^S}}{\partial q_1} > 0$ and $\frac{\partial \overline{\Phi^S}}{\partial q_2} = 0$, so that the second part of Lemma 90 yields that $\mathbf{q}$ does not maximize $\overline{\Phi}$.

We next examine case (II). Since $\overline{\Phi^S}$ is scale-free (see (193)), we may assume that $C_1 = 1$. Since R_1, R_2 are positive, it holds that $\partial \overline{\Phi^S}/\partial R_1 = \partial \overline{\Phi^S}/\partial R_2 = 0$, yielding

$$R_1 \propto y^d, \;\; R_2 \propto 1, \;\; \text{where } y = (q_1 + B - 1)/q_1.$$

Expressing q_1, q_2 in terms of y and substituting in $\overline{\Phi^S}$, we obtain the value of $\overline{\Phi}(\mathbf{q})$:

$$\overline{\Phi}(\mathbf{q}) = \log h(y), \;\; \text{where } h(y) := \frac{(1-B)\big(q(1-y) - (1-B)(1-y^{d+1})\big)}{(1-y)^2}.$$

Let I be the interval $[0, (q+B-1)/q]$. Note that for any $y \in I$, there exists a positive $q_1 \in [0, q]$ such that $y = (q_1 + B - 1)/q_1$. Obviously, if $\mathbf{q}$ maximizes $\overline{\Phi}$, it must be the case that y maximizes $h(y)$ in the interval I. We compute $h'(y)$.

$$h'(y) = \frac{(1-B)\, r(y)}{(1-y)^3}, \;\; \text{where } r(y) := q(1-y) - (1-B)\big((d-1)y^{d+1} - (d+1)y^d + 2\big).$$

It is immediate to see that $r(y)$ is convex for $y \in [0,1]$. Since $r(0) = q - 2(1-B) > 0$ and $r(1) = 0$, we obtain that either

$$\text{(i) } r(y) > 0 \text{ for all } y \in I, \;\; \text{or}$$

$$\text{(ii) } \exists\, y_o \in I\colon\; r(y_o) = 0,\; r(y) > 0 \text{ iff } y < y_o.$$

In case (i), $h(y)$ is increasing and hence $h(y)$ is maximized at $y = (q+B-1)/q$. This value of y corresponds to $q_1 = q$ and thus $\Phi(\mathbf{q}) = \Phi(q, 0, 0)$.

In case (ii), we have $h(y) \leq h(y_o)$. The value of q_1 corresponding to y_o is $q_o := (1 - B)/(1 - y_o)$. We will show that the maximum in (191) does not happen at the boundary $C_2 = 0$ when $\mathbf{q} = (q_o, q - q_o, 0)$, implying that $h(y_o)$ does not equal $\overline{\Phi}(\mathbf{q})$ and hence the maximum of $\overline{\Phi}$ as well. To prove the former, we utilize the first part of Lemma 93. In particular, we prove that

$$q_o y_o^d + (q - q_o) > (1 - B). \tag{207}$$

Note that $r(y_o) = 0$ yields $q = (1 - B)\big((d-1)y_o^{d+1} - (d+1)y_o^d + 2\big)/(1 - y_o)$. Plugging this expression into (207), we only need to show that

$$\frac{(d-1)y_o^{d+1} - dy_o^d + 1}{1 - y_o} > 1 \text{ or } (d-1)y_o^d + 1 > dy_o^{d-1}, \tag{208}$$

which holds by the AM-GM inequality for any positive $y_o \neq 1$.

Let $\mathbf{q} = (q_1, q_2, q_3)$ be a bad triple where all of the q_i's are positive. Since $\mathbf{q}$ is bad, at least one of the R_i's and C_j's is zero. W.l.o.g. we may assume $C_2 = 0$. Moreover, by the second part of Lemma 93, we may also assume that $R_1, C_1 > 0$. There are four cases to consider.

$$\text{(I) } R_2 = 0, \;\; \text{(II) } R_2, R_3 > 0, C_3 = 0, \;\; \text{(III) } R_2, R_3, C_3 > 0, \quad \text{(IV) } R_2, C_3 > 0, R_3 = 0.$$

We omitted the case $R_2 > 0$ and $R_3 = C_3 = 0$, which is identical to case (I) after renaming the q_i's.

Case (I) is straightforward: since $R_2 = C_2 = 0$, (194) gives $\partial\overline{\Phi^S}/\partial q_2 = 0$. Since at least one of $\partial\overline{\Phi^S}/\partial q_1$, $\partial\overline{\Phi^S}/\partial q_3$ is positive, the second part of Lemma 90 yields that $\mathbf{q}$ does not maximize $\overline{\Phi}$.

We next examine case (II). Since $\overline{\Phi^S}$ is scale-free (see (193)), we may substitute $C_1 = 1$. Setting the derivatives of $\partial\overline{\Phi^S}/\partial R_1, \partial\overline{\Phi^S}/\partial R_2, \partial\overline{\Phi^S}/\partial R_3$ equal to zero, we obtain

$$R_1 \propto (q_1 + B - 1)^d/q_1^d, \;\; R_2 \propto 1, \;\; R_3 \propto 1.$$

It follows that $\overline{\Phi}(\mathbf{q}) = \overline{\Phi}(q_1, q_2 + q_3, 0)$ and hence the maximum of $\overline{\Phi}$ does not occur at $\mathbf{q}$ by the argument for case (II) in (206).

We next examine case (III). The partial derivatives of $\overline{\Phi^S}$ with respect to R_1, R_2, R_3, C_1, C_3 must vanish so we obtain

$$R_1^{1/d} \propto q_1C_1 + q_2C_3 - (1-B)C_1,\ R_2^{1/d} \propto q_1C_1 + q_3C_3,\ R_3^{1/d} \propto q_1C_1 + q_3C_3 - (1-B)C_3,$$
$$C_1^{1/d} \propto q_1R_1 + q_2R_2 + q_3R_3 - (1-B)R_1,\ C_3^{1/d} \propto q_1R_1 + q_2R_2 + q_3R_3 - (1-B)R_3. \tag{209}$$

If $C_1 = C_3$, then $R_1 = R_3$ and thus we obtain $\overline{\Phi}(q_1, q_2, q_3) = \overline{\Phi}(q_1 + q_3, q_2, 0)$, contradicting the maximality of $\mathbf{q}$ by the argument for case (II) in (206). Thus, wlog we may assume $C_1 < C_3$. By (209), this yields

$$R_2 > R_1 > R_3,\ C_1 < C_3. \tag{210}$$

We have the following analogue of Lemma 91, which proves that the maximum cannot occur at $\mathbf{q}$ by the second part in Lemma 90.

Lemma 94. *Suppose that R_i's and C_j's satisfy* (209) *and* (210)*. If $R_1/R_3 \neq C_3/C_1$ then $\frac{\partial\overline{\Phi^S}}{\partial q_1} - \frac{\partial\overline{\Phi^S}}{\partial q_3} \neq 0$. If $R_1/R_3 = C_3/C_1$ then $\frac{\partial\overline{\Phi^S}}{\partial q_1} - \frac{\partial\overline{\Phi^S}}{\partial q_2} \neq 0$.*

Proof of Lemma 94. The proof is analogous to the proof of Lemma 91, we highlight the main differences. Let $r_1^d = R_1/R_3, r_2^d = R_2/R_3, c_3^d = C_3/C_1$. The R_i's and C_j's may be recovered by the r_i's and c_j's by

$$R_1 \propto r_1^d, R_2 \propto r_2^d, R_3 \propto 1, \text{ and } C_1 \propto 1, C_3 \propto c_3^d. \tag{211}$$

By (210), we have

$$r_2 > r_1 > 1 \text{ and } c_3 > 1.$$

The expressions for r_1, r_2, c_3 in (200) are exactly the same after substituting $c_2 = 0$. The same is true for (201), (202). It follows that the proof for the first part of Lemma 91 holds verbatim in this case as well (note that the ordering of r_1, r_2 is different here but that part of the argument does not use the ordering).

While the proof for the second part of Lemma 91 does not carry through as simply, the changes are minor. We assume that $r_1 = c_3$ and set $DIF_{12} := \frac{\partial\overline{\Phi^S}}{\partial q_1} - \frac{\partial\overline{\Phi^S}}{\partial q_2}$. Plugging $r_1 = c_3$ and $c_2 = 0$ in (201), (202) and then substituting the resulting expressions in DIF_{12}

we obtain

$$DIF_{12} = \frac{(1-B)\,h(r_1)}{r1-1}, \text{ where } h(r_1) := (r_1^{2d+1} + dr_1^{d+1} - dr_1^d - 1) - \frac{(r_1^{d+1}-1)^{d+1}}{(r_1^d-1)^d}.$$

By a first derivative argument, the function

$$g(r_1) := \log\left(\frac{(r_1^{d+1}-1)^{d+1}}{(r_1^d-1)^d\,(r_1^{2d+1} + dr_1^{d+1} - dr_1^d - 1)}\right),$$

is strictly increasing for $r_1 > 1$. Thus, $g(r_1) \geq g(+\infty) = 0$, which gives $h(r_1) > 0$ for all $r_1 > 1$. This proves that $DIF_{12} \neq 0$, as desired. □

Finally, we examine case (IV). The partial derivatives of $\overline{\Phi^S}$ with respect to R_1, R_2, C_1, C_3 must vanish so we obtain

$$\begin{aligned} &R_1^{1/d} \propto q_1C_1 + q_3C_3 - (1-B)C_1,\ R_2^{1/d} \propto q_1C_1 + q_3C_3, \\ &C_1^{1/d} \propto q_1R_1 + q_2R_2 - (1-B)R_1,\ C_3^{1/d} \propto q_1R_1 + q_2R_2. \end{aligned} \tag{212}$$

Note that we have $R_1 < R_2$ and $C_1 < C_3$.

Lemma 95. *If $R_2/R_1 \neq C_3/C_1$ then either $\frac{\partial\overline{\Phi^S}}{\partial q_2} - \frac{\partial\overline{\Phi^S}}{\partial q_3} \neq 0$ or $\frac{\partial\overline{\Phi^S}}{\partial q_1} - \frac{\partial\overline{\Phi^S}}{\partial q_2} \neq 0$. If $R_2/R_1 = C_3/C_1$ and $\frac{\partial\overline{\Phi^S}}{\partial q_1} - \frac{\partial\overline{\Phi^S}}{\partial q_2} = 0$, then the maximum in* (191) *does not happen at the boundary $C_2 = 0$.*

Proof of Lemma 95. The approach for the first part is similar the proof of Lemma 91. Set $r_2^d = R_2/R_1$ and $c_3^d = C_3/C_1$, so that $r_1, c_3 > 1$. Dividing appropriate pairs in (212), we obtain

$$r_2 = \frac{q_1 + q_3c_3^d}{(q_1+B-1) + q_3c_3^d}, \quad c_3 = \frac{q_1 + q_2r_2^d}{(q_1+B-1) + q_2r_2^d}. \tag{213}$$

It follows that

$$q_2 = \frac{q_1 - (q_1+B-1)c_3}{r_2^d(c_3-1)}, \quad q_3 = \frac{q_1 - (q_1+B-1)r_2}{c_3^d(r_2-1)}.$$

Using these, we obtain

$$\frac{\partial\overline{\Phi^S}}{\partial q_2} - \frac{\partial\overline{\Phi^S}}{\partial q_3} = 0 \Rightarrow f(r_2) = f(c_3), \text{ where } f(x) := \frac{x^{d+1}}{x-1}, \tag{214}$$

$$\frac{\partial\overline{\Phi^S}}{\partial q_1} - \frac{\partial\overline{\Phi^S}}{\partial q_3} = 0 \Rightarrow r_2^{d+1}(c_3-1) - (d+1)r_2c_3 + d(r_2+c_3) - (d-1) = 0. \tag{215}$$

It is relatively simple to prove that this cannot be the case unless $r_2 = c_3$. This gives the first part.

For the second part, we have $r_2 = c_3$ and $\frac{\partial \overline{\Phi^S}}{\partial q_1} - \frac{\partial \overline{\Phi^S}}{\partial q_3} = 0$. Thus, (215) yields

$$r_2^{d+1} = (d+1)r_2 - (d-1). \tag{216}$$

To prove that the maximum does not happen at the boundary $C_2 = 0$, we use the first part of Lemma 93. It suffices to prove that

$$q_1 + q_2 r_2^d > (1-B)r_2^d. \tag{217}$$

We have that $q_2 = q_3 = (q - q_1)/2$, so that (213) gives

$$q_1 = \frac{q(r_2^{d+1} - r_2^d) - 2r_2(1-B)}{(r_2 - 1)(r_2^d - 2)}, \quad q_2 = q_3 = \frac{q - r_2(q + B - 1)}{(r_2 - 1)(r_2^d - 2)}. \tag{218}$$

Plugging (218) into (217) gives the equivalent inequality

$$\frac{(1-B)\left(r_2^d + r_2 - r_2^{d+1}\right)}{r_2 - 1} > 0$$

To see the latter, use (216) to obtain

$$r_2^d + r_2 - r_2^{d+1} = r_2^d + (d-1) - dr_2 > 0, \text{ for all } r_2 > 1 \text{ by the AM-GM inequality.}$$

This completes the proof. □

6.3.5 Remaining proofs

Proof of Lemma 83. Let $R_i = r_i^d$, $C_i = c_i^d$, $r = \sum_{i=1}^q r_i^d$, and $c = \sum_{i=1}^q c_i^d$. We have

$$r_i = c - (1-B)c_i^d \quad \text{and} \quad c_i = r - (1-B)r_i^d,$$

It is clear from this equation that $R_i = R_j$ iff $C_i = C_j$ and hence also $t_R = t_C$. We also obtain that for $i = 1, \dots, q$,

$$r_i = c - (1-B)(r - (1-B)r_i^d)^d. \tag{219}$$

Since r is the sum of r_i^d and the r_i are positive, we have $(1-B)r_i^d < r$. Fix the values of r, c and let I be the interval where $(1-B)x^d < r$. Using (219), we shall prove that $t_R \leq 3$

by arguing that $f(x) = c - (1-B)(r-(1-B)x^d)^d - x$ has at most 3 positive roots in the interval I, counted by multiplicities. We have

$$f'(x) = (1-B)^2 d^2 (r-(1-B)x^d)^{d-1} x^{d-1} - 1 = \left(\sum_{i=0}^{d-2} g(x)^i\right)(g(x)-1),$$

where

$$g(x) = ((1-B)d)^{2/(d-1)}(r-(1-B)x^d)x.$$

Note that $g(x) > 0$ in the interval I and hence all roots of $f'(x)$ in this interval come from $g(x)-1$. The polynomial $g(x)-1$ has at most two positive roots by Descartes' rule of signs, hence $f'(x)$ has at most two positive roots in I. Thus, $f(x)$ has at most three positive roots in I, all roots counted by their multiplicities. This concludes the proof. □

Proof of Lemma 85. Let $q' = q/2$. To better align with the results of Section 6.3.3, let us assume that the fixpoint $(q', 0, q')$ maximizes Ψ_1. In Section 6.3.3, we proved that this can be the case only if $R_1/R_3 = C_3/C_1$ or (in the parameterization of Section 6.3.3) $r_1 = c_3 =: x$ where $x > 1$. Equation (200) for $q_2 = 0$, $q_1 = q_3 = q'$ gives that x satisfies

$$x = \frac{B + q' - 1 + q'x^d}{q' + (B + q' - 1)x^d}. \tag{220}$$

It is straightforward to check that (220) has exactly one solution $x > 1$ for all $0 \le B < \frac{\Delta - q}{\Delta}$. The values of R_1, C_1, R_3, C_3 may be recovered by (198), which in the case $q_2 = 0$ give

$$R_1 \propto x^d, R_3 \propto 1 \text{ and } C_1 \propto 1, C_3 \propto x^d.$$

This proves the second part of the lemma. For the first part, to check Jacobian stability, we proceed as in the proof of Lemma 81. The eigenvalues of the matrix $\mathbf{L}$ in this case can be computed easily as well. They are given by ± 1 by multiplicity 1, $\pm\lambda_1$ by multiplicity $q-2$ and $\pm(B+q-1)\lambda_1^2$ by multiplicity 1, where

$$\lambda_1 := \frac{(1-B)x^{d/2}}{\sqrt{(q' + (B+q'-1)x^d)(B+q'-1+q'x^d)}}.$$

To prove that the absolute value of the eigenvalues different from 1 is less than $1/d$, it suffices to prove that $\lambda_1 < 1/d$. Use (220) to solve for q' and plug the value into the expression for λ_1. This yields that λ_1 is equal to $x^{(d-1)/2}(x-1)/(x^d - 1)$, which by the AM-GM inequality is less than $1/d$ for $x > 1$. □

Proof of Lemma 86. For non-negative $\mathbf{q} = (q_1, q_2, q_3)$ with $q_1 + q_2 + q_3 = q$, consider the function

$$\overline{F}(\mathbf{q}) = \max_{\mathbf{r},\mathbf{c}} F(\mathbf{q}, \mathbf{r}, \mathbf{c}), \text{ where } F(\mathbf{q}, \mathbf{r}, \mathbf{c}) := \textstyle\sum_{i=1}^3 q_i R_i \sum_{j=1}^3 q_j C_j + (B-1)\sum_{i=1}^3 q_i R_i C_i, \tag{221}$$

and the maximum is over the compact region (by restricting to $R_i = C_i = 0$ whenever $q_i = 0$)

$$\begin{gathered} \textstyle\sum_{i=1}^3 q_i R_i^{(d+1)/d} \leq 1, \ \sum_{j=1}^3 q_j C_j^{(d+1)/d} \leq 1, \\ R_1, R_2, R_3, C_1, C_2, C_3 \geq 0. \end{gathered} \tag{222}$$

Note that $\overline{F}(\mathbf{q}) > 0$, since we can set all of the R_i's and C_j's equal to x, where $qx^{(d+1)/d} = 1$. Clearly, $\overline{\Phi}(\mathbf{q}) \geq \ln \overline{F}(\mathbf{q})$. Since $\overline{\Phi^S}(\mathbf{q}, \mathbf{r}, \mathbf{c})$ is scale-free with respect to $\mathbf{r}$ and $\mathbf{c}$ (see (193)), we may scale $\mathbf{r}, \mathbf{c}$ to satisfy (222) and hence $\overline{\Phi}(\mathbf{q}) = \sup_{\mathbf{r},\mathbf{c}} \overline{\Phi^S}(\mathbf{q}, \mathbf{r}, \mathbf{c}) \leq \ln \overline{F}(\mathbf{q})$, proving that $\overline{\Phi}(\mathbf{q}) = \ln \overline{F}(\mathbf{q})$ and consequently the supremum is attained.

To prove that $\sup_{\mathbf{q}} \overline{\Phi}(\mathbf{q})$ is attained, it clearly suffices to prove that $L := \sup_{\mathbf{q}} \overline{F}(\mathbf{q})$ is attained. This can be accomplished by using variants of Berge's Maximum Theorem [8] and showing that the function $\overline{F}(\mathbf{q})$ is upper semi-continuous. We give a more direct argument, which is similar to the proof of Berge's Maximum Theorem and can also easily be adapted to show that $\overline{F}(\mathbf{q})$ is upper semi-continuous.

Note first that $L < \infty$ by a simple application of Hölder's inequality. Let $\mathbf{q}_n$ be such that $\overline{F}(\mathbf{q}_n) \uparrow L$. Since the $\mathbf{q}_n$ lie in a compact region, by restricting to a subsequence we may assume that $\mathbf{q}_n \to \mathbf{q}$. Let $\mathbf{r}_n, \mathbf{c}_n$ be maximizers for $\overline{F}(\mathbf{q}_n)$ in (221). Suppose first that $\mathbf{q}$ has positive entries. Then, for sufficiently large n, the maximizers $\mathbf{r}_n, \mathbf{c}_n$ lie in a compact set and hence a standard diagonalisation argument yields a convergent subsequence $(\mathbf{q}_{n_k}, \mathbf{r}_{n_k}, \mathbf{c}_{n_k}) \to (\mathbf{q}, \mathbf{r}, \mathbf{c})$. By continuity, $\mathbf{r}, \mathbf{c}$ must lie in the region (222) defined by $\mathbf{q}$ and moreover $\overline{F}(\mathbf{q}_{n_k}) = F(\mathbf{q}_{n_k}, \mathbf{r}_{n_k}, \mathbf{c}_{n_k}) \to F(\mathbf{q}, \mathbf{r}, \mathbf{c})$. Thus $L = F(\mathbf{q}, \mathbf{r}, \mathbf{c})$ and the supremum is attained.

Assume now that $\mathbf{q}$ has an entry equal to zero, say q_1, so that $q_{1n} \to 0$ (with the natural notation for entries of the subsequences). In this setting, R_{1n}, C_{1n} might escape to infinity, so assume that $R_{1n}, C_{1n} \uparrow \infty$, by restricting to a subsequence if necessary. However, (222) implies $q_{1n} R_{1n}^{(d+1)/d}, q_{1n} C_{1n}^{(d+1)/d} \leq 1$ and hence $q_{1n} R_{1n}, q_{1n} C_{1n} \to 0$. Note that $q_{1n}R_{1n}C_{1n} \to 0$ as well; otherwise there exists a subsequence with $q_{1n_k}R_{1n_k}C_{1n_k} \geq \varepsilon > 0$. This contradicts that $\mathbf{r}_{n_k}, \mathbf{c}_{n_k}$ maximize $F(\mathbf{q}_{n_k}, \cdot, \cdot)$, since setting $R_{1,n_k} = C_{1,n_k} = 0$ would maintain feasibility in (222) and achieve a bigger value of F for all sufficiently large k (recall that $B < 1$). Thus $q_{1n}R_{1n}, q_{1n}C_{1n}, q_{1n}R_{1n}C_{1n} \to 0$, yielding once again $L = F(\mathbf{q}, \mathbf{r}, \mathbf{c})$.

□

Proofs of Lemmas 87 and 90. We first prove Lemma 90. Let $I_R = \{i \in I \,|\, R_i > 0\}$. For $i \in I_R$, it must hold that $\partial\overline{\Phi^S}/\partial R_i = 0$. Since $q_i > 0$ for $i \in I$, it follows that

$$R_i^{1/d} \propto \textstyle\sum_j q_j C_j - (1-B)C_i \quad \text{for all } i \in I_R, \tag{223}$$

and hence

$$R_i^{(d+1)/d} \propto R_i\big(\textstyle\sum_j q_j C_j - (1-B)C_i\big) \text{ for all } i \in I.$$

Thus, for $i \in I$ it holds that

$$\frac{R_i^{(d+1)/d}}{\sum_j q_j R_j^{(d+1)/d}} = \frac{R_i\big(\sum_j q_j C_j - (1-B)C_i\big)}{\sum_j q_j R_j \sum_j q_j C_j + (B-1)\sum_j q_j R_j C_j}, \tag{224}$$

and an analogous argument for the C_i's gives

$$\frac{C_i^{(d+1)/d}}{\sum_j q_j C_j^{(d+1)/d}} = \frac{C_i\big(\sum_j q_j R_j - (1-B)R_i\big)}{\sum_j q_j R_j \sum_j q_j C_j + (B-1)\sum_j q_j R_j C_j}. \tag{225}$$

Moreover, by a direct calculation we have

$$\frac{\partial\overline{\Phi^S}}{\partial q_i} = \frac{(d+1)\big(R_i\sum_j q_j C_j + C_i\sum_j q_j R_j + (B-1)R_iC_i\big)}{\sum_j q_j R_j \sum_j q_j C_j + (B-1)\sum_j q_j R_j C_j} - \frac{dR_i^{(d+1)/d}}{\sum_j q_j R_j^{(d+1)/d}} - \frac{dC_i^{(d+1)/d}}{\sum_j q_j C_i^{(d+1)/d}}. \tag{226}$$

Plugging (224), (225) in (226) proves the first part of Lemma 90.

For the second part of Lemma 90, assume w.l.o.g. that $q_1, q_2 > 0$ and $\frac{\partial\overline{\Phi^S}}{\partial q_1} - \frac{\partial\overline{\Phi^S}}{\partial q_2} > 0$. For $\varepsilon > 0$, consider $\mathbf{q}' = (q_1 + \varepsilon, q_2 - \varepsilon, q_3)$. Since q_1, q_2 are positive, for small enough ε, $\mathbf{q}'$ has positive entries which sum to q. Moreover, for small enough ε the value of $\overline{\Phi^S}$ increases, while still maintaining feasibility in the region (192). Hence, $\mathbf{q}$ does not maximize $\overline{\Phi}$, as desired.

Lemma 87 follows easily: just use (223) and the fact that q_1, q_2, q_3 are integers to get the alignment with (18). □

Proof of Lemma 93. Suppose that $q_i > 0$ and $\sum_j q_j C_j > (1-B)C_i$. We look at the derivative $\partial\overline{\Phi^S}/\partial R_i$ evaluated at $R_i = 0$:

$$\frac{\partial\overline{\Phi^S}}{\partial R_i} = \frac{q_i(q_1C_1 + q_2C_2 + q_3C_3 - (1-B)C_i)}{\sum_j q_jR_j \sum_j q_jC_j + (B-1)\sum_j q_jR_jC_j} > 0$$

Thus, increasing the value of R_i by a sufficiently small amount, we increase the value of $\overline{\Phi^S}$. Hence, the maximum cannot be obtained at the boundary $R_i = 0$. The second part of the lemma follows immediately from the first part. □

6.3.5.1 *Uniqueness of semi-translation invariant measures (Antiferromagnetic Potts)*

In this section, we prove Lemma 80. As noted earlier, the proof extends the respective argument in [11] for colorings in the antiferromagnetic Potts model setting. That said, the technical details, due to the soft contraints, are relatively more intricate.

Proof of Lemma 80. We may assume that $R_1 \geq \ldots \geq R_q$. Then the equations easily imply $C_1 \leq \ldots \leq C_q$. Define

$$\alpha = \frac{R_1}{R_q}, \quad \beta = \frac{R_1 + \ldots + R_{q-1}}{(q-1)R_q}, \quad S = R_1 + \ldots + R_{q-1}.$$

We clearly have $\alpha \geq \beta \geq 1$, and we may assume for the sake of contradiction that $\beta > 1$. Note that

$$\alpha^{1/d} = \left(\frac{R_1}{R_q}\right)^{1/d} = 1 + \frac{(1-B)(C_q - C_1)}{C_1 + \ldots + C_{q-1} + BC_q}$$
$$C_q = (R_1 + \ldots + R_{q-1} + BR_q)^d = \big[(q-1)\beta + B\big]^d R_q^d$$
$$C_1 = (BR_1 + R_2 + \ldots + R_q)^d = \big[(q-1)\beta + 1 - (1-B)\alpha\big]^d R_q^d$$

Moreover, by Holder's inequality or otherwise, we have

$$\begin{aligned} C_1 + \ldots + C_{q-1} + BC_q &= \sum_{i=1}^{q-1} \big[S + R_q - (1-B)R_i\big]^d + B(S + BR_q)^d \\ &\geq (q-1)\Big[\frac{q-2+B}{q-1}S + (q-1)R_q\Big]^d + B(S + BR_q)^d \\ &= (q-1)\big[(q-2+B)\beta + 1\big]^d R_q^d + B\big[(q-1)\beta + B\big]^d R_q^d. \end{aligned}$$

Thus, we obtain that every solution must satisfy

$$\alpha^{1/d} \leq 1 + \frac{(1-B)\left\{\left[(q-1)\beta + B\right]^d - \left[1-(1-B)\alpha + (q-1)\beta\right]^d\right\}}{(q-1)\left[(q-2+B)\beta+1\right]^d + B\left[(q-1)\beta + B\right]^d} \iff$$

$$0 \leq 1 - \alpha^{1/d} + \frac{(1-B)\left[1 - \left(1 - \frac{(1-B)(\alpha-1)}{(q-1)\beta+B}\right)^d\right]}{(q-1)\left[1 - \frac{(1-B)(\beta-1)}{(q-1)\beta+B}\right]^d + B} =: f(\alpha, \beta, B).$$

To obtain a contradiction, our goal is to prove that for q and B as in the statement of the lemma, when $(q-1)\beta > \alpha \geq \beta > 1$, it holds that $f(\alpha, \beta, B) < 0$.

It is easy to see that f is decreasing in B. This immediately yields the lemma for $q \geq \Delta$: it holds that $f(\alpha, \beta, B) \leq f(\alpha, \beta, 0) < 0$, since the last inequality was proved by [11]. For $q \leq d$ and $B \geq \frac{d+1-q}{d+1} := B_c$, this yields

$$f(\alpha, \beta, B) \leq f\left(\alpha, \beta, B_c\right) =: g(\alpha, \beta).$$

We first prove that $g(\alpha, \beta) \leq g(\beta, \beta)$. For $q = 2$ there is nothing to prove. Hence we may assume that $d \geq q \geq 3$. Clearly it suffices to prove that g is decreasing in α. This requires a fair bit of work, so we state it as a Lemma to prove later.

Lemma 96. *For $d \geq q \geq 3$ and $B_c = \frac{d+1-q}{d+1}$, the function $g(\alpha, \beta)$ is decreasing in α for $\alpha \geq \beta > 1$.*

We finish the proof by showing that for $\beta \geq 1$, it holds that $g(\beta, \beta) \leq 0$ with equality iff $\beta = 1$. After massaging the inequality, this reduces to

$$1 \leq \left[1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta + B_c}\right]^d \left[(q-1)\left(\beta^{1/d} - 1\right) + 1 - B_c\right] + B_c\beta^{1/d} =: h(\beta)$$

Note that the inequality holds at equality for $\beta = 1$, so it suffices to prove $h'(\beta) > 0$ for $\beta > 1$, which is the assertion of the next lemma.

Lemma 97. *For $d \geq q \geq 2$ and $B_c = \frac{d+1-q}{d+1}$, the function $h(\beta)$ is increasing for $\beta \geq 1$.*

Modulo the proofs of Lemmas 96 and 97, which are given below, the proof is complete.

□

Proof of Lemma 96. We compute

$$\frac{\partial g}{\partial \alpha} = -\frac{1}{d}\alpha^{-(d-1)/d} + \frac{(1-B_c)^2}{(q-1)\beta + B_c} \cdot \frac{d\left[1 - \frac{(1-B_c)(\alpha-1)}{(q-1)\beta+B_c}\right]^{d-1}}{(q-1)\left[1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta+B_c}\right]^d + B_c}$$

Let $F(x) = x\left[1 - \frac{(1-B_c)(x-1)}{(q-1)\beta+B_c}\right]^d$ for $x \in [\beta, (q-1)\beta]$. Straightforward manipulations show that $\frac{\partial g}{\partial \alpha} < 0$ is equivalent to

$$d^2(1-B_c)^2 F(\alpha)^{(d-1)/d} \le \big[(q-1)\beta + B_c\big]\Big[(q-1)\Big(1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta+B_c}\Big)^d + B_c\Big]. \quad (227)$$

We prove that $F(x)$ is decreasing in $[\beta, (q-1)\beta]$. It is simple to check that

$$F'(x) = \left[1 - \frac{(1-B_c)(x-1)}{(q-1)\beta+B_c}\right]^{d-1} \frac{(q-1)\beta + 1 - (d+1)(1-B_c)x}{(q-1)\beta+B_c}.$$

For $x \in [\beta, (q-1)\beta]$, we have $(d+1)(1-B_c)x = qx = (q-1)x + x > (q-1)\beta + 1$, where in the last inequality we used that $\beta > 1$. It follows that $F(x)$ is indeed decreasing and thus $F(\alpha) \le F(\beta)$.

To prove (227), it thus suffices to argue that for $\beta > 1$ it holds

$$d^2(1-B_c)^2 F(\beta)^{(d-1)/d} \le \big[(q-1)\beta + B_c\big]\Big[(q-1)\Big(1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta+B_c}\Big)^d + B_c\Big]. \quad (228)$$

Note that $q - 1 + B_c = d(1-B_c)$ so that the inequality is tight for $\beta = 1$. By the weighted AM-GM inequality on A^d and 1 with weights $(q-1)$ and B_c respectively, we obtain

$$(q-1)A^d + B_c \ge (q-1+B_c)A^{d(q-1)/(q-1+B_c)} = d(1-B_c)A^{(q-1)(d+1)/q}.$$

We use this for $A = \Big(1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta+B_c}\Big)^d$ so that, after simplifications, it suffices to show that

$$d(1-B_c)\beta^{(d-1)/d} \le \big[(q-1)\beta + B_c\big]\left[1 - \frac{(1-B_c)(\beta-1)}{(q-1)\beta+B_c}\right]^{-(d+1-2q)/q}.$$

This can further be massaged into

$$G(\beta) := \beta^{(d-1)/d}\big[(q-1)\beta + B_c\big]^{-(d+1-q)/q}\big[(q-2+B_c)\beta + 1\big]^{(d+1-2q)/q} \le \frac{1}{d(1-B_c)}.$$

Once again, note that the inequality holds at equality for $\beta = 1$, so it suffices to prove that $G'(\beta) < 0$ for $\beta > 1$. This has nothing special, apart from tedious, but otherwise straightforward, calculations. We include the details briefly. Differentiating $\ln G(\beta)$, we obtain

$$\frac{G'(\beta)}{G(\beta)} = \frac{(d-1)}{d\beta} - \frac{(d+1-q)(q-1)}{q\big[(q-1)\beta + B_c\big]} + \frac{(d+1-2q)(q-2+B_c)}{q\big[(q-2+B_c)\beta+1\big]}.$$

By clearing denominators, it suffices to check that the following second order polynomial $p(\beta)$ is negative whenever $\beta > 1$:

$$p(\beta) := (d-1)q[(q-1)\beta + B_c]\big[(q-2+B_c)\beta+1\big] - (d+1-q)(q-1)d\beta\big[(q-1)\beta + B_c\big] \\ + (d+1-2q)(q-2+B_c)\beta\big[(q-1)\beta+B_c\big].$$

Using again that $q-1+B_c = d(1-B_c)$ it is easy to verify that $p(1) = 0$. The factorization of $p(\beta)$ (using the value of B_c) is given by

$$p(\beta) = -\frac{q(\beta-1)\big[\beta\big(d(q-1)^2-(q-1)\big)+d(d-q)+q-1\big]}{d+1},$$

which is obviously negative for $\beta > 1$, whenever $d \geq q \geq 2$. □

Proof of Lemma 97. We compute

$$h'(\beta) = \frac{1}{d}\beta^{-(d-1)/d}\left[(q-1)\Big(1-\frac{(1-B_c)(\beta-1)}{(q-1)\beta+B_c}\Big)^d + B_c\right] \\ - d\left[1-\frac{(1-B_c)(\beta-1)}{(q-1)\beta+B_c}\right]^{d-1}\frac{(1-B_c)(q-1+B_c)[(q-1)\beta^{1/d}-(q-2+B_c)]}{[(q-1)\beta+B_c]^2}.$$

Thus, to prove $h'(\beta) > 0$ it suffices to check (using $q-1+B_c = d(1-B_c)$ and the function F defined in Lemma 96)

$$d^3(1-B_c)^2F(\beta)^{(d-1)/d} \leq \\ \leq \frac{\big[(q-1)\beta+B_c\big]^2}{(q-1)\beta^{1/d}-(q-2+B_c)}\left[(q-1)\Big(1-\frac{(1-B_c)(\beta-1)}{(q-1)\beta+B_c}\Big)^d + B_c\right],$$

This is similar to (228) and in fact follows from (228), once we prove that

$$\frac{(q-1)\beta^{1/d}-(q-2+B_c)}{(q-1)\beta+B_c} \leq \frac{1}{d}.$$

To see the last inequality, observe that $\beta + d - 1 \geq d\beta^{1/d}$ as a consequence of the weighted AM-GM inequality (or otherwise). Hence,

$$\frac{(q-1)\beta^{1/d}-(q-2+B_c)}{(q-1)\beta+B_c} \leq \frac{(q-1)\beta+(d-1)(q-1)-d(q-2+B_c)}{d\big[(q-1)\beta+B_c\big]} = \frac{1}{d},$$

completing the proof. □

CHAPTER VII

AP-REDUCTIONS USING NON-RECONSTRUCTION

For our AP-reductions, we will use a modification of the gadget we analysed in Chapter 4. The properties of the gadget were first proved in the context of the hard-core model in [63].

7.1 A modified gadget

Recall from Section 4.1 the definition of the graph distribution $\mathcal{G}_n^r$, for integers $n > r \geq 0$:

1. $\mathcal{G}_n^r$ is supported on bipartite graphs. The two parts of the bipartite graph are labelled by $+, -$ and each is partitioned as $V^s := U^s \cup W^s$ where $|U^s| = n$, $|W^s| = r$ for $s \in \{+, -\}$. U denotes the set $U^+ \cup U^-$ and similarly W denotes the set $W^+ \cup W^-$.

2. To sample $\overline{G} \sim \mathcal{G}_n^r$, sample uniformly and independently Δ matchings: (i) $(\Delta - 1)$ random perfect matchings between $U^+ \cup W^+$ and $U^- \cup W^-$, (ii) a random perfect matching between U^+ and U^-. The edge set of $\overline{G}$ is the union of the Δ matchings. Thus, vertices in U have degree Δ, while vertices in W have degree $\Delta - 1$.

Note here the slight change of notation for a random graph from the distribution $\mathcal{G}_n^r$, since we will reserve G for the final gadget. The construction of G has two parameters $0 < \theta, \psi < 1/8$. Let $r = (\Delta - 1)^{\lfloor \theta \log_{\Delta-1} n \rfloor + 2\lfloor \frac{\psi}{2} \log_{\Delta-1} n \rfloor}$, so that $r = o(n^{1/4})$. First, sample $\overline{G}$ from the distribution $\mathcal{G}_n^r$ conditioning on $\overline{G}$ being simple. Next, for $s \in \{+, -\}$, attach t disjoint $(\Delta - 1)$-ary trees of even depth ℓ (with $t = (\Delta - 1)^{\lfloor \theta \log_{\Delta-1} n \rfloor}$ and $\ell = 2\lfloor \frac{\psi}{2} \log_{\Delta-1} n \rfloor$) to W^s, so that every vertex in W^s is a leaf of exactly one tree (this is possible since $m' = |W| = t(\Delta - 1)^\ell$). Denote by R^s the roots of the trees, so that $|R^s| = t$. The trees do not share common vertices with the graph G, apart from the vertices in W. The final graph G is the desired gadget.

Note that the definition of the phase of a configuration we gave in Section 4.1 extends to this modified gadget as well. Specifically, the phase of a configuration depends only the spins of vertices U (the large portion of the graph).

The purpose of using the trees in the modified construction is related to Remark 13. As showed in [63], the trees, together with (93), allow to obtain the pointwise conditional independence property 2 for the roots of the trees and to make the error of approximation polynomially small in n. The proof uses the non-reconstruction (see [55]) of the Gibbs measures corresponding to the dominant phases for the hard-core and ferromagnetic Potts models.

7.2 #BIS-hardness for the ferromagnetic Potts model

In this section we prove the #BIS-hardness result for the ferromagnetic Potts model. Our first objective is to describe the properties of the modified gadget G.

The regime which will be interesting to us is $B > \mathfrak{B}_o$, where Theorem 69 gives the existence of q dominant phases. Recall that the dominant phases for the ferromagnetic Potts model are of the form $(\boldsymbol{\alpha}, \boldsymbol{\alpha})$, where $\boldsymbol{\alpha}$ ranges over the permutations of $(p', \frac{1-p'}{q-1}, \ldots, \frac{1-p'}{q-1})$, for some $p' > 1/q$. Specifically, p' is the probability that the root of the infinite Δ-regular tree is assigned the spin i in the Gibbs measure corresponding to the ordered phase i. More interesting for our considerations in this section will be the same quantity in the $(\Delta-1)$-ary tree, which we denote by p. The probabillity p is the largest solution bigger than $1/q$ of the following equation

$$\frac{(q-1)p}{1-p} = \left(\frac{(B-1)p+1}{(B-1)\frac{1-p}{q-1}+1}\right)^{\Delta-1},$$

as can easily be derived from the tree recursions (8).

It would be useful to have an explicit form of the product distribution (92). For a phase i and $S \subset V$, we will denote by $Q^i_S(\cdot)$ the product distribution (92) on the set of configurations $\sigma : S \to [q]$. Specifically, for $[i] \in q$ and $\sigma : S \to [q]$, $Q^i_S(\sigma)$ is given by

$$Q^i_S(\sigma) = p^{|\sigma^{-1}(i)\cap S|}\left(\frac{1-p}{q-1}\right)^{|R\setminus\sigma^{-1}(i)|},$$

Our interest in this section will be mostly in the case $S = R$ (the roots of the trees of the gadget), though to analyze the properties of G we will also look at the case $S = W$, in order to transfer our results for the random graph $\overline{G}$ from Chapter 4.

Using Theorem 34, we prove the following lemma in Section 7.4. The main ingredient here is that the L^1 asymptotic convergence of $\mu_{\overline{G}}(\sigma_W \,|\, Y(\sigma) = i)$ to $Q^i_S(\sigma_W)$ is now an L^∞

asymptotic convergence of $\mu_G(\sigma_R \,|\, Y(\sigma) = i)$ to $Q^i_R(\sigma_R)$. Also, due to symmetries of the Potts model, the probability that the phases occur are inverse polynomially close to $1/q$.

Lemma 98. *For every $\Delta \geq 3$ and $B > \mathfrak{B}_o$, there exist constants $\theta(\Delta, B), \psi(\Delta, B) > 0$ such that the graph G satisfies the following with probability $1 - o(1)$ over the choice of the graph:*

1. *The phases occur with roughly equal probability, so that for every phase $i \in [q]$, we have*
$$\Big|\mu_G\big(Y(\sigma) = i\big) - \frac{1}{q}\Big| \leq n^{-2\theta}.$$

2. *Conditioned on the phase i, the spins of vertices in R are approximately independent, that is,*
$$\max_{\sigma_R} \Big|\frac{\mu_G\big(\sigma_R \,|\, Y = i\big)}{Q^i_R(\sigma_R)} - 1\Big| \leq n^{-2\theta}.$$

With Lemma 98 at hand, we can now give the reduction. Let $B > \mathfrak{B}_o$. Let H be a graph on n' vertices, where $n' \leq n^{\theta/4}$ and θ is as in Lemma 98. Assuming an FPRAS for the ferromagnetic Potts model on bipartite graphs of maximum degree Δ for some $B > \mathfrak{B}_o$, we will show that we can approximate $Z_H(B^*)$, the partition function of H in the ferromagnetic Potts model with temperature B^*, where B^* will be determined shortly.

To do this, we construct first a graph H^G. First, take $|H|$ disconnected copies of the gadget G in Lemma 98 and identify each copy with a vertex $v \in H$. Denote by $\hat{H}^G$ the resulting graph, G_v the copy of the gadget associated to the vertex v in H and by R^+_v, R^-_v, R_v the images of R^+, R^-, R in the gadget G_v, respectively. Finally, we denote by R the set of vertices $\cup_v R_v$. We next add the edges of H in $\hat{H}^G$. To do this, fix an arbitrary orientation of the edges of H. For each oriented edge (u, v) of H, we add an edge between one vertex in R^+_u and one vertex in R^-_v, using mutually distinct vertices for distinct edges of H. The resulting graph will be denoted by H^G. Note that H^G is bipartite and has maximum degree Δ.

We have the following connection:

Lemma 99. *Let $\Delta \geq 3$ and $B > \mathfrak{B}_o$. There exists B^* such that the following holds*

$$\big(1 - O(n^{-\theta})\big)\frac{q^{n'} Z_{H^G}(B)}{C_H\big(Z_G(B)\big)^{n'}} \leq Z_H(B^*) \leq \big(1 + O(n^{-\theta})\big)\frac{q^{n'} Z_{H^G}(B)}{C_H\big(Z_G(B)\big)^{n'}},$$

where $C_H = D^{|E(H)|}$ and $D = 1 + (B-1)\Big(\frac{2p(1-p)}{(q-1)^2} + (q-2)\frac{(1-p)^2}{(q-1)^2}\Big)$.

Proof of Lemma 99. For each $v \in H$, let $Y_v(\sigma)$ denote the phase of a configuration σ on G_v and let $\mathcal{Y} = (Y_v)_{v \in H} \in [q]^H$ be the phase vector of vertices in H. For $\mathcal{Y}' \in [q]^H$, let $Z_{H^G}(\mathcal{Y})$ be the partition function of H^G restricted to configurations with phase vector $\mathcal{Y}'$, that is

$$Z_{H^G}(\mathcal{Y}') = \sum_{\sigma: H^G \to [q]} B^{m(\sigma)} \mathbf{1}\{\mathcal{Y}(\sigma) = \mathcal{Y}'\},$$

where for a configuration σ, $m(\sigma)$ is the number of monochromatic edges under σ. We may view $\mathcal{Y}'$ as an assignment $H \to [q]$ and analogously we denote by $m(\mathcal{Y}')$ as the number of monochromatic edges in H under $\mathcal{Y}'$. We then have

$$\frac{Z_{H^G}(\mathcal{Y}')}{Z_{\hat{H}^G}(\mathcal{Y}')} = \sum_{\sigma_R} \mu_{\hat{H}^G}(\sigma_R \,|\, \mathcal{Y}(\sigma) = \mathcal{Y}') \prod_{(u,v) \in E(H^G) \setminus E(\hat{H}^G)} B^{\mathbf{1}\{\sigma(u) = \sigma(v)\}}$$

Note that $\mu_{\hat{H}^G}(\sigma_R \,|\, \mathcal{Y}(\sigma) = \mathcal{Y}') = \big(1 + O(n^{-\theta})\big)\prod_{v \in H} Q_{R_v}^{Y_v}(\sigma_{R_v})$ since $\hat{H}^G$ is a union of disconnected copies of G and in each copy of G we have property 2. It follows that

$$\begin{aligned}\frac{Z_{H^G}(\mathcal{Y}')}{Z_{\hat{H}^G}(\mathcal{Y}')} &= \big(1 + O(n^{-\theta})\big)\sum_{\sigma_R}\prod_{v \in H} Q_{R_v}^{Y_v}(\sigma_{R_v}) \prod_{(u,v) \in E(H^G) \setminus E(\hat{H}^G)} B^{\mathbf{1}\{\sigma(u) = \sigma(v)\}} \\ &= \big(1 + O(n^{-\theta})\big) A^{m(\mathcal{Y}')} D^{|E(H)| - m(\mathcal{Y}')},\end{aligned}$$

where A (resp. D) is the expected weight of an edge for two gadgets which have same (resp. different) phases. Simple calculations show that

$$A = 1 + (B-1)\Big(p^2 + \frac{(1-p)^2}{q-1}\Big), \; D = 1 + (B-1)\Big(\frac{2p(1-p)}{(q-1)^2} + (q-2)\frac{(1-p)^2}{(q-1)^2}\Big).$$

Letting $B^* = A/D$ and $C_H = D^{|E(H)|}$, we obtain

$$\frac{Z_{H^G}(\mathcal{Y}')}{Z_{\hat{H}^G}(\mathcal{Y}')} = \big(1 + O(n^{-\theta})\big) C_H (B^*)^{m(\mathcal{Y}')} \tag{229}$$

Item 1 in Lemma 98 gives that for every $\mathcal{Y}'$ it holds that

$$\big(1 - O(n^{-\theta})\big) q^{-n'} \leq \Big(\frac{1}{q} - n^{-2\theta}\Big)^{n'} \leq \frac{Z_{\hat{H}^G}(\mathcal{Y})}{Z_{\hat{H}^G}} \leq \Big(\frac{1}{q} + n^{-2\theta}\Big)^{n'} \leq \big(1 + O(n^{-\theta})\big) q^{-n'}. \tag{230}$$

We also have

$$Z_{H^G}(B) = \sum_{\mathcal{Y}} Z_{H^G}(\mathcal{Y}) = \sum_{\mathcal{Y}} \frac{Z_{H^G}(\mathcal{Y})}{Z_{\hat{H}^G}(\mathcal{Y})} Z_{\hat{H}^G}(\mathcal{Y}) = Z_{\hat{H}^G} \sum_{\mathcal{Y}} \frac{Z_{H^G}(\mathcal{Y})}{Z_{\hat{H}^G}(\mathcal{Y})} \frac{Z_{\hat{H}^G}(\mathcal{Y})}{Z_{\hat{H}^G}}. \tag{231}$$

Using the estimates (229), (230) in (231), we obtain

$$\big(1 - O(n^{-\theta})\big) q^{-n'} C_H Z_H(B^*) \leq \frac{Z_{H^G}(B)}{Z_{\hat{H}^G}(B)} \leq \big(1 + O(n^{-\theta})\big) q^{-n'} C_H.$$

The result follows after observing that $Z_{\hat{H}^G}(B) = \big(Z_G(B)\big)^{n'}$ and rearranging the inequality.

□

Proof of Theorem 18. Suppose that there exists an FPRAS for approximating the partition function with temperature B on graphs with maximum degree Δ. First sample the graph G. G satisfies the properties in Lemma 98 with probability $1 - o(1)$. Approximate the partition function of G within a multiplicative factor $1 \pm \varepsilon/10n'$. Approximate the partition function of H^G within a multiplicative factor $1 \pm \varepsilon/2$. The bounds for $Z_H(B^*)$ in Lemma 99 are then within a factor $1 \pm \varepsilon$, giving an FPRAS for approximating the partition function with temperature B^*. This, together with the result of [30], implies an FPRAS for counting independent sets in bipartite graphs. □

7.3 #BIS-hardness for the hard-core model

The following lemma captures the properties of the final gadget $G = (V, E)$ in the case of the hard-core model (it is similar to Lemma 98, the main difference is that in the first bullet the probabilities of the phases are less balanced—the reason for this is that there is less symmetry between the phases).

Lemma 100. *For every $\Delta \geq 3$ and $\lambda > \lambda_c(\Delta)$, for every $\varepsilon > 0$, there exist constants $\theta(\Delta, \lambda), \psi(\Delta, \lambda) > 0$ such that for all sufficiently large n the graph G satisfies with probability $1 - \varepsilon$ the following:*

1. *The two phases occur with roughly equal probability, so that for every phase $i \in \{\pm 1\}$, we have*

$$\Big| \mu_G\big(Y(\sigma) = i\big) - \frac{1}{2} \Big| \leq \varepsilon.$$

2. Conditioned on the phase i, the spins of vertices in R are approximately independent, that is,

$$\max_{\sigma_R}\left|\frac{\mu_G\big(\sigma_R\,|\,Y=i\big)}{Q^i_R(\sigma_R)}-1\right|\le n^{-2\theta},$$

where Q^i is the following product measure on the configurations $\sigma_R : R^+\cup R^-\to\{\pm1\}$:

$$Q^i(\sigma_R)=(p^i)^{|\sigma_R^{-1}(+1)\cap R^+|}(1-p^i)^{|\sigma_R^{-1}(-1)\cap R^+|}(p^{-i})^{|\sigma_R^{-1}(+1)\cap R^+|}(1-p^{-i})^{|\sigma_R^{-1}(-1)\cap R^+|},$$

where p^j is the probability that the root of the infinite $(\Delta-1)$-ary has spin $+1$ in the ordered phase $j\in\{\pm1\}$.

Before proceeding with the reduction we will fix some parameters. Let

$$U=1-p^+p^+,\quad W=1-p^+p^-,\quad V=1-p^-p^-. \tag{232}$$

Note that $p^-<p^+$ implies $U<W<V$. We also have

$$W^2-UV=(p^--p^+)^2>0.$$

Let $0<a<b$ be such that $a+b=1$ and

$$U^aV^b=W^{a+b}. \tag{233}$$

Note that for $a=0,b=1$ we have $W^{a+b}<U^aV^b$ whereas for $a=b=1/2$ we have $W^{a+b}>U^aV^b$. Thus such an a,b with $a<b$ exists.

Now we are ready to give the reduction. Let H be a graph with n' vertices. We will need to approximate a and b from the previous paragraph by integers (where the precision depends on H). By Dirichlet's theorem on simultaneous Diophantine approximation there exists positive integer $n'^5\le Q\le n'^{40}$ and integers $a',b'\in\{0,\dots,Q\}$ such that

$$|aQ-a'|\le n'^{-15}\quad\text{and}\quad|bQ-b'|\le n'^{-15}.$$

The gadget from Lemma 100 that we will use will have $n^\theta\ge n'^{80}$. Note that n is polynomial in n'.

Again, we construct the graph H^G. First, take $2|H|$ disconnected copies of the gadget G from Lemma 100 and identify two copies with a vertex $v\in H$. Denote by $\hat{H}^G$ the resulting graph. Let $G_{v,+1}$ and $G_{v,-1}$ denote the two copies of the gadget associated to the vertex v in H. For each v we put a perfect matching between half of the vertices R^+ in $G_{v,+}$ and half of the vertices R^+ in $G_{v,-}$, connecting the corresponding vertices. Also, for each v we put a perfect matching between half of the vertices R^- in $G_{v,+}$ and half of the vertices R^- in $G_{v,-}$, connecting the corresponding vertices. For each v call the resulting graph G_v. Note that G_v is symmetric—this will deal with the asymmetry in Lemma 100 (intuitively, the only phase assignments that will matter are ones where $G_{v,+1}$ and $G_{v,-1}$ have opposite phases— this results in two possibilities and, by symmetry, they have equal probability). The remaining (unmatched) R^- vertices from $G_{v,-1}$ will be referred to as L^- and the remaining (unmatched) R^- vertices from $G_{v,+1}$ will be referred to as L^+. Let $\tilde{H}^G$ be the union of the G_v's.

Now for each edge $\{u,v\}$ of H we 1) add matching with a' edges between the L^+ vertices of G_u and L^+ vertices of G_v and 2) add a matching with b' edges between the L^- vertices of G_u and the L^- vertices of G_v.

We have the following connection:

Lemma 101. *Let $\Delta \geq 3$ and $\lambda > \lambda_c(\Delta)$. The following holds:*

$$\big(1 - O(n^{-\theta})\big)\frac{2^{2n'} Z_{H^G}(\lambda)}{C_H\big(Z_G(\lambda)\big)^{2n'}} \leq \mathrm{BIS}(H) \leq \big(1 + O(n^{-\theta})\big)\frac{2^{2n'} Z_{H^G}(\lambda)}{C_H\big(Z_G(\lambda)\big)^{2n'}},$$

where $C_H = W^{n'k+(a+b)Q|E(H)|}$ and W is as in (232).

Proof of Lemma 101. For each $v \in H$ and $i \in \{\pm 1\}$, let $Y_{v,i}(\sigma)$ denote the phase of a configuration σ on $G_{v,i}$. Let

$$Y_v(\sigma) = \begin{cases} -1 & \text{if } Y_{v,+} = +1 \text{ and } Y_{v,-} = -1, \\ +1 & \text{if } Y_{v,+} = -1 \text{ and } Y_{v,-} = +1, \\ 0 & \text{otherwise.} \end{cases}$$

Let $\mathcal{Y} = (Y_{v,i})_{v\in H, i\in\{\pm1\}} \in \{\pm 1\}^{H\times\{\pm1\}}$ be the phase vector of vertices in H (each vertex gets two phases). For $\mathcal{Y}' \in \{\pm 1\}^{H\times\{\pm1\}}$, let $Z_{H^G}(\mathcal{Y}')$ be the partition function of H^G restricted to configurations with phase vector $\mathcal{Y}'$, that is

$$Z_{H^G}(\mathcal{Y}') = \sum_{\sigma: H^G\to\{-1,+1\}} \lambda^{|\sigma^{-1}(+1)|}\mathbf{1}\{\mathcal{Y}(\sigma) = \mathcal{Y}'\} \prod_{\{u,v\}\in E(H^G)} \mathbf{1}\{\sigma(u) \neq +1 \vee \sigma(v) \neq +1\}.$$

Let R be the union of all R^+ and R^- vertices (over all $G_{v,+}$ and $G_{v,-}$) that are matched in the construction. We have

$$\frac{Z_{H^G}(\mathcal{Y}')}{Z_{\hat{H}^G}(\mathcal{Y}')} = \sum_{\sigma_R} \mu_{\hat{H}^G}(\sigma_R \,|\, \mathcal{Y}(\sigma) = \mathcal{Y}') \prod_{\{u,v\} \in E(H^G) \setminus E(\hat{H}^G)} \mathbf{1}\{\sigma(u) \neq +1 \vee \sigma(v) \neq +1\},$$

where the sum ranges over all $\sigma_R : R \to \{\pm 1\}$.

Note that $\mu_{\hat{H}^G}(\sigma_R \,|\, \mathcal{Y}(\sigma) = \mathcal{Y}') = \big(1 + O(n^{-\theta})\big) \prod_{v \in H} Q_{R_v}^{Y_v}(\sigma_{R_v})$ since $\hat{H}^G$ is a union of disconnected copies of G and in each copy of G we have property 2. It follows that

$$\begin{aligned}\frac{Z_{H^G}(\mathcal{Y}')}{Z_{\hat{H}^G}(\mathcal{Y}')} &= \big(1 + O(n^{-\theta})\big) \sum_{\sigma_R} \prod_{v \in H} Q_{R_v}^{Y_v}(\sigma_{R_v}) \prod_{\{u,v\} \in E(H^G) \setminus E(\hat{H}^G)} \mathbf{1}\{\sigma(u) \neq +1 \vee \sigma(v) \neq +1\} \\ &= \big(1 + O(n^{-\theta})\big) W^{n'k} \prod_{v \in V(H)} \left(\frac{UV}{W^2}\right)^{(k/2)\mathbf{1}\{Y_{v,+} = Y_{v,-}\}} \times \\ &\qquad \prod_{\{u,v\} \in E(H)} F(-Y_{u,+}, -Y_{v,+})^{a'} F(Y_{u,-}, Y_{v,-})^{b'},\end{aligned}$$

where $F(+,+) = U$, $F(+,-) = W$, and $F(-,-) = V$.

For $\mathcal{Y}$ such that $Y_{v,+} = Y_{v,-}$ for some v we have

$$\frac{Z_{H^G}(\mathcal{Y}')}{Z_{\hat{H}^G}(\mathcal{Y}')} \leq \big(1 + O(n^{-\theta})\big) W^{n'k} \left(\frac{UV}{W^2}\right)^{k/2} V^{(a'+b')|E(H)|}. \tag{234}$$

Assuming $Y_{u,+} \neq Y_{u,-}$ and $Y_{v,+} \neq Y_{v,-}$ (or equivalently $Y_u \neq 0$ and $Y_v \neq 0$) we have

$$\frac{F(-Y_{u,+}, -Y_{v,+})^{a'} F(Y_{u,-}, Y_{v,-})^{b'}}{W^{(a+b)Q}} \in \begin{cases} [V^{-2n'^{-5/2}}, V^{2n'^{-5/2}}] & \text{if } Y_u \neq +1 \vee Y_v \neq +1, \\ [0, (U^b V^a / W^{a+b})^{n'^5}] & \text{if } Y_u = Y_v = +1. \end{cases} \tag{235}$$

Hence if $Y_u \neq 0$ and $Y_v \neq 0$ but there exists $\{u,v\} \in H$ such that $Y_u = Y_v = +1$ then

$$\frac{Z_{H^G}(\mathcal{Y}')}{Z_{\hat{H}^G}(\mathcal{Y}')} \leq \big(1 + O(n^{-\theta})\big) W^{n'k + (a+b)Q|E(H)|} \left(\frac{U^b V^a}{W^{a+b}}\right)^{n'^5} V^{(a'+b')|E(H)|}. \tag{236}$$

On the other hand, if $Y_u \neq 0$ and $Y_v \neq 0$ and for all $\{u,v\} \in H$ we have $Y_u \neq +1 \vee Y_v \neq +1$, then by (235) the edges of H contribute a factor in $[V^{-2n'^{-1/2}}, V^{2n'^{-1/2}}]$. It follows that

$$\frac{Z_{H^G}(\mathcal{Y}')}{Z_{\hat{H}^G}(\mathcal{Y}')} \in \big(1 + O(n^{-\theta})\big) W^{n'k + (a+b)Q|E(H)|} [V^{-2n'^{-1/2}}, V^{2n'^{-1/2}}]. \tag{237}$$

Item 1 in Lemma 100 gives that for every $\mathcal{Y}'$ it holds that

$$\left(\frac{1}{2} - \varepsilon\right)^{2n'} \leq \frac{Z_{\hat{H}^G}(\mathcal{Y}')}{Z_{\hat{H}^G}} \leq \left(\frac{1}{2} + \varepsilon\right)^{2n'}. \tag{238}$$

We also have

$$Z_{H^G}(\lambda) = \sum_{\mathcal{Y}} Z_{H^G}(\mathcal{Y}) = \sum_{\mathcal{Y}} \frac{Z_{H^G}(\mathcal{Y})}{Z_{\hat{H}^G}(\mathcal{Y})} Z_{\hat{H}^G}(\mathcal{Y}) = Z_{\hat{H}^G} \sum_{\mathcal{Y}} \frac{Z_{H^G}(\mathcal{Y})}{Z_{\hat{H}^G}(\mathcal{Y})} \frac{Z_{\hat{H}^G}(\mathcal{Y})}{Z_{\hat{H}^G}}. \tag{239}$$

Using estimates (234), (235), (236), (237), and (238) we obtain that the contribution of "bad" $\mathcal{Y}$ (that is, one with $Y_v = 0$ or with a pair $Y(u) = Y(v) = +1, \{u, v\} \in E(H)$) is $\exp(-\Theta(n'^5))$ smaller than contribution of "good" $\mathcal{Y}$. For good $\mathcal{Y}$ we have, by symmetry,

$$\frac{Z_{\tilde{H}^G}(\mathcal{Y})}{Z_{\tilde{H}^G}} = \Big(\frac{1}{2} + \exp(-\Theta(n'^5))\Big)^{n'}. \tag{240}$$

Now using (237) we obtain the result. $\square$

Proof of Theorem 17. Suppose that there exists an FPRAS for approximating the partition function with activity λ on bipartite graphs of maximum degree Δ. First sample the graph G. G satisfies the properties in Lemma 100 with probability $1 - \varepsilon'$, where ε' can be taken arbitrarily small, say $\varepsilon' = 1/3$. Approximate the partition function of G within a multiplicative factor $1 \pm \varepsilon/10n'$. Approximate the partition function of H^G within a multiplicative factor $1 \pm \varepsilon/2$. The bounds for $\mathrm{BIS}(H)$ in Lemma 99 are then within a factor $1 \pm \varepsilon$, giving an FPRAS for counting independent sets in bipartite graphs. $\square$

7.4 Proving the properties of the modified gadgets

In this section, we prove the properties of the gadgets we use, as stated in Lemmas 98 and 100. The proofs follow the same approach as in [63, Theorem 2.1]. We will need to argue however more thoroughly for Item 1 in Lemma 98 and Item 1 in Lemma 100, since in [63] a cruder bound for the probability that a phase appears was enough. In the case of the Potts model, the more delicate bound will follow from the model's symmetries, while in the case of the hard-core model the bound will follow from the small subgraph conditioning results of Chapter 4.

Let Σ_G^i be the set of configurations on G which have phase i, i.e., $\Sigma_G^i := \{\sigma : V \to [q] \mid Y(\sigma) = i\}$. Moreover, let Σ_G^o be the set of configurations σ with $\big|\arg\max_{i\in[q]} |\sigma^{-1}(i) \cap U|\big| \geq 2$, that is, Σ_G^o consists of these configurations whose phase was determined by breaking a tie. As we will illustrate, our goal is to show that Σ_G^o has exponentially smaller contribution to the partition function of G than Σ_G^i for every $i \in [q]$.

To capture this, for a subset $\Sigma \subseteq \Omega_G$ of the configuration space, denote by $Z_G(\Sigma)$ the partition function restricted to configurations in Σ, that is,

$$Z_G(\Sigma) = \sum_{\sigma \in \Sigma} w_G(\sigma).$$

Let π be a permutation of the colors $[q]$ which maps color i to color j. For a configuration σ, we denote by $\pi(\sigma)$ the configuration $\pi \circ \sigma$. Clearly, for every configuration $\sigma \in \Sigma^i_G \backslash \Sigma^o_G$ we have $\pi(\sigma) \in \Sigma^j_G \backslash \Sigma^o_G$. It follows that for every two colors i, j we have $Z_G(\Sigma^i_G \backslash \Sigma_0) = Z_G(\Sigma^j_G \backslash \Sigma^0_G)$. Since

$$Z_G = Z_G(\Sigma^o_G) + \sum_i Z_G(\Sigma^i_G \backslash \Sigma^o_G),$$

to get the inequality in Item 1 of Lemma 98 it suffices to show that $Z_G(\Sigma^o_G)$ is smaller than Z_G by an arbitrary polynomial factor. In particular, note that the definition of the phase of a configuration makes sense for configurations on $\overline{G}$ as well. For convenience, we will henceforth use Z^o_G, Z^i_G as shorthands for $Z_G(\Sigma^o_G), Z_G(\Sigma^i_G)$ and $Z^o_{\overline{G}}, Z^i_{\overline{G}}$ for their counterparts in $\overline{G}$. We will show that $Z^o_{\overline{G}}$ is exponentially smaller than $Z^i_{\overline{G}}$. To transfer this to G, note that for every configuration on $\overline{G}$, the multiplicative contribution of the configurations on the trees to the partition function is at most $\exp(n^{1/2})$, which gives $Z^o_G \le \exp(n^{1/2}) Z^i_{\overline{G}}$. It will turn out that this is at least an $\exp(n^{1/2})$ factor smaller than Z^i_G. Summing over $i \in [q]$ will yield the desired bound.

To formalize the argument, we will have to capture how the partition functions $Z_{\overline{G}}$ and Z_G interplay. Due to the Markov property, this happens only through vertices in W. We will partition the sets $\Sigma^o_{\overline{G}}, \Sigma^i_{\overline{G}}$ according to the configuration η on W. In particular, $\Sigma^o_{\overline{G}}(\eta)$ will be those configurations σ in $\Sigma^o_{\overline{G}}$ such that $\sigma_W = \eta$ and $Z^o_{\overline{G}}(\eta)$ will be the contribution to the partition function of $\overline{G}$ from configurations in $\Sigma^o_{\overline{G}}(\eta)$. Define similarly $\Sigma^i_{\overline{G}}(\eta)$ and $Z_{\overline{G}}(\eta)$.

We need a final piece of notation. Denote by J the graph induced by the edges in $E(G) \backslash E(\overline{G})$. Note that J is the union of the trees (with vertices in W included). Let $Z_J(\eta)$ be the contribution to the partition function of J from configurations σ (on J) such that

$\sigma_W = \eta$. We are now able to put these definitions into work. In particular, we have that

$$Z^i_G = \sum_{\eta: W \to [q]} Z^i_{\overline{G}}(\eta) Z_J(\eta).$$

We will need the following lemma, whose proof follows from the small subgraph conditioning method of Chapter 4 and the phase diagram (Theorem 69) for the ferromagnetic Potts model (note that for ferromagnetic models, the dominant phases have the same quantitative structure in bipartite graphs). The only part of the lemma which we have not accounted in Chapter 4 is Item (ii), which follows from a straightforward application of Markov's inequality.

Lemma 102. *Let $\mathcal{G} := \mathcal{G}_n$ denote the distribution of the random bipartite graph $\overline{G}$. For $B > \mathfrak{B}_o$, it holds that*

(i) *There exist constants C_1, C_2, T depending only on q, B, Δ, such that for every $i \in [q]$ and $\eta : W \to [q]$,*

$$\begin{aligned} \mathbf{E}_{\mathcal{G}}\left[Z^i_{\overline{G}}\right] &= (1+o(1)) C_1 n^{q-1} C_2^{2m'} \exp(nT), \\ \mathbf{E}_{\mathcal{G}}[Z^i_{\overline{G}}(\eta)] &= (1+o(1)) Q^i_W(\eta) \mathbf{E}_{\mathcal{G}}\left[Z^i_{\overline{G}}\right]. \end{aligned} \tag{241}$$

(ii) *For all sufficiently small $\varepsilon > 0$ and sufficiently large n, for every $\eta : W \to [q]$ and $i \in [q]$,*

$$\mathbf{E}_{\mathcal{G}}\left[Z^o_{\overline{G}}(\eta)\right] \leq \exp(-\varepsilon n) \mathbf{E}_{\mathcal{G}}\left[Z^i_{\overline{G}}\right]. \tag{242}$$

(iii) $\max_{i \in [q],\, \eta: W \to [q]} \Pr_{\mathcal{G}}\left(Z^i_{\overline{G}}(\eta) < \frac{1}{n} \mathbf{E}_{\mathcal{G}}[Z^i_{\overline{G}}(\eta)] \right) \to 0$ *as* $n \to \infty$.

Using Lemma 102 we can give the proof of Lemma 98.

Proof of Lemma 98. Item 2 of the lemma follows exactly the approach in [63]. The required reconstruction results to push the approach in [63] are given in [47, Proof of Theorem 1.4] (ferromagnetic Potts model on the tree with constant boundary condition). Together with Lemma 102, the proof of [63, Theorem 2.1] extends almost verbatim to our case as well.

To get Item 1, we use Lemma 102. In particular, Markov's inequality yields

$$\Pr_{\mathcal{G}^r_n}\left(\sum_{\eta: W \to [q]} Z^o_{\overline{G}}(\eta) Z_J(\eta) > n \sum_{\eta: W \to [q]} Z_J(\eta) \mathbf{E}_{\mathcal{G}}[Z^o_{\overline{G}}(\eta)] \right) \to 0 \text{ as } n \to \infty. \tag{243}$$

Item (iii) of Lemma 102 yields for every $i \in [q]$

$$\Pr_{\mathcal{G}_n^r}\left(\sum_{\eta:W\to[q]} Z^i_{\overline{G}}(\eta)Z_J(\eta) < \frac{1}{2n}\sum_{\eta:W\to[q]} Z_J(\eta)\mathbf{E}_{\mathcal{G}}[Z^i_{\overline{G}}(\eta)]\right) \to 0 \text{ as } n\to\infty. \tag{244}$$

We also have the crude bound $\max_\eta Z_J(\eta) \le \exp(o(n^{1/2}))\min_\eta Q^i_W(\eta)Z_J(\eta)$. Combining (241), (242), (243), (244) and the symmetry argument described in the beginning of the section yields the first item of the lemma, concluding the proof. □

Proof of Lemma 100. The second item of the lemma is proved in [63], extended in [26] and [27]. The first item in our current version of the lemma has a more delicate bound than in [63], where the more precise bound was not needed. We show how to obtain the desired bound.

The small subgraph conditioning method in Chapter 4 gives that for all sufficiently large n, there exist random variables W_{mn}, a function of the number of cycles of length $2, 4, \dots, 2m$ in $\overline{G}$, such that $Z^i_{\overline{G}}(\eta) = (W_{mn} \pm \varepsilon')\mathbf{E}_{\overline{G}}[Z^i_{\overline{G}}(\eta)]$ with probability $1-\varepsilon$. As we saw, when the graph is bipartite, the random variable W_{mn} is bounded by a constant $c > 0$ for all m. Let ε' be sufficiently smaller than c. It follows that for any $\varepsilon' > 0$ sufficiently smaller than c and any $t < 1$, it holds that

$$\lim_{m\to\infty}\limsup_{n\to\infty}\Pr_{\mathcal{G}_n^r}\left(\sum_{\eta:W\to[q]} Z^i_{\overline{G}}(\eta)Z_J(\eta) < t(W_{mn}\pm\varepsilon')\sum_{\eta:W\to[q]} Z_J(\eta)\mathbf{E}_{\mathcal{G}}[Z^i_{\overline{G}}(\eta)]\right) = 0. \tag{245}$$

and

$$\lim_{m\to\infty}\limsup_{n\to\infty}\Pr_{\mathcal{G}_n^r}\left(\sum_{\eta:W\to[q]} Z^i_{\overline{G}}(\eta)Z_J(\eta) > \frac{1}{t}(W_{mn}\pm\varepsilon')\sum_{\eta:W\to[q]} Z_J(\eta)\mathbf{E}_{\mathcal{G}}[Z^i_{\overline{G}}(\eta)]\right) = 0. \tag{246}$$

Note that

$$\sum_{\eta:W\to[q]} Z_J(\eta)\mathbf{E}_{\mathcal{G}}[Z^+_{\overline{G}}(\eta)] = \sum_{\eta:W\to[q]} Z_J(\eta)\mathbf{E}_{\mathcal{G}}[Z^-_{\overline{G}}(\eta)].$$

It follows that letting $\varepsilon' \to 0$ and then $t \uparrow 1$, the two bounds in (245) and (246) are within a factor of $1 \pm \varepsilon/4$ from each other and the result follows. □

CHAPTER VIII

EXTENSIONS TO RANDOM REGULAR GRAPHS

In this chapter, we discuss how our results extend to random Δ-regular graphs. We are going to establish that for ferromagnetic models our techniques can be used to completely analyze the Gibbs distribution on a random Δ-regular graph, in a completely analogous way as we did for bipartite graphs. For models which are not ferromagnetic, we will need the extra assumption of semi-translation invariant uniqueness to analyze the Gibbs distribution.

8.1 Preliminaries

8.1.1 The distribution on Δ-regular Graphs

We will use the standard pairing model, introduced in [9] (see also [5]), to study random Δ-regular bipartite graphs.

For $\Delta \geq 3$ and Δn even, conside the set $[\Delta n]$. Elements of $[\Delta n]$ will be called *points* and a *pairing* of the points is a perfect matching of $[\Delta n]$. The pairing model $\mathcal{G}_{n,\Delta}$ is a probability distribution on Δ-regular multigraphs. The distribution is generated by the following random process. First, sample a uniformly random pairing of the points. Given the pairing, construct a graph G with n vertices by identifying for $i = 1, \dots, n$ the points $\Delta(i-1)+1, \dots, \Delta i$ as a single vertex. The edges of G are induced by the pairing in the natural way.

Let $G \sim \mathcal{G}_{n,\Delta}$. G is *simple* if it does not contain parallel edges or self loops. Every simple graph G corresponds to exactly $(\Delta!)^n$ pairings (note that this is not true if G is not simple). Let $\mathcal{G}^s_{n,\Delta}$ denote the conditional distribution on the event that G is simple; it follows that $\mathcal{G}^s_{n,\Delta}$ is the uniform distribution over Δ-regular graphs with n vertices. For fixed Δ, the probability that G is simple tends to $\exp\big(-(\Delta^2-1)/4\big)$ as $n \to \infty$ ([5, 9]). Consequently, a property holds asymptotically almost surely over the distribution $\mathcal{G}_{n,\Delta}$ iff it holds asymptotically almost surely over the distribution $\mathcal{G}^s_{n,\Delta}$; the distributions $\mathcal{G}_{n,\Delta}$ and $\mathcal{G}^s_{n,\Delta}$ are thus *contiguous* (see [36, Section 9.6] for an account of contiguity).

We will invariably work with the pairing model and translate the results to $\mathcal{G}^s_{n,\Delta}$ per our needs. With a slight abuse of terminology, due to the contiguity of $\mathcal{G}_{n,\Delta}$ and $\mathcal{G}^s_{n,\Delta}$, we shall refer to $\mathcal{G}_{n,\Delta}$ as the distribution of uniformly random Δ-regular graphs; since our results hold a.a.s. over $\mathcal{G}_{n,\Delta}$, this should not lead to any confusions. We will also omit Δ from notation whenever the context is clear.

8.1.2 Ferromagnetic spin systems

Before giving the definition of a ferromagnetic spin system in terms of the interaction matrix $\mathbf{B}$, we note that there is no loss of generality to assume in this chapter that $\mathbf{B}$ is ergodic (irreducible and aperiodic), as a consequence of the fact that a random Δ-regular graph is almost surely connected and non-bipartite. Indeed, if $\mathbf{B}$ is reducible by a suitable permutation of the labels of colors, $\mathbf{B}$ can be put in a block diagonal form where each of the blocks is either irreducible or zero. The original spin system can be studied by considering the induced sub-models of each block which correspond to irreducible symmetric matrices (where our results apply). These sub-models are "non-interacting" for connected graphs G, that is, the partition function for the original model is simply the sum of the partition functions of each sub-model. Similarly, if $\mathbf{B}$ is periodic, its period must be two since $\mathbf{B}$ is symmetric; such a model is only interesting on bipartite graphs.

To define ferromagnetic systems, recall that in Section 5.5.1 we defined antiferromagnetic spin systems as those spin systems whose interaction matrix $\mathbf{B}$ has a single positive eigenvalue and the rest are negative. Analogously, we use the following definition of ferromagnetic spin systems.

Definition 8. *A model is called* **ferromagnetic** *if* $\mathbf{B}$ *is positive definite. Equivalently we have that all of its eigenvalues are positive and also that*

$$\mathbf{B} = \hat{\mathbf{B}}^{\intercal}\hat{\mathbf{B}}, \tag{15}$$

for some $q \times q$ matrix $\hat{\mathbf{B}}$.

The most alluring aspect of this definition is that for a ferromagnetic model, neighboring vertices prefer to have the same spin. More generally, and analogously to Lemma 60, we have the following simple application of the Cauchy-Schwarz inequality.

Lemma 103. *Let* $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^q_{\geq 0}$ *with* $\|\mathbf{z}_1\|_1 = \|\mathbf{z}_2\|_1 = 1$. *For ferromagnetic* $\mathbf{B}$, *we have*

$$(\mathbf{z}_1^{\intercal}\mathbf{B}\mathbf{z}_1)(\mathbf{z}_2^{\intercal}\mathbf{B}\mathbf{z}_2) \geq (\mathbf{z}_1^{\intercal}\mathbf{B}\mathbf{z}_2)^2. \tag{247}$$

Equality holds iff $\mathbf{z}_1 = \mathbf{z}_2$. *Recall that for antiferromagnetic* $\mathbf{B}$, *the inequality is reversed.*

Observe that if we plug in the inequality (247) the vectors with a single 1 in the positions i and j respectively, we obtain that any two spins i, j induce a ferromagnetic two-spin system.

Proof of Lemma 103. Just use the decomposition (15) and use the Cauchy-Scharz inequality on the vectors $\hat{\mathbf{B}}\mathbf{z}_1$ and $\hat{\mathbf{B}}\mathbf{z}_2$. □

We also remind the reader that the definition of ferromagnetism/antiferromagnetism in terms of the signature of the interaction matrix is invariant in the presence of external fields. In particular, for Δ-regular graphs, any external field can be pushed into the interaction matrix $\mathbf{B}$ with a congruence transformation of the matrix $\mathbf{B}$. The resulting interaction matrix, by Sylvester's law of inertia, has the same number of positive, zero and negative eigenvalues and in particular remains ferromagnetic/antiferromagnetic.

8.1.3 First and second moments

For a q-spin system with interaction matrix $\mathbf{B}$, our goal is to understand the Gibbs distribution on a random Δ-regular graph $G = (V, E)$ by looking at the distribution of spin values in V. Let $n = |V|$. For a configuration $\sigma : V \to [q]$, we denote the set of vertices assigned spin i by $\sigma^{-1}(i)$. For a q-dimensional probability vector $\boldsymbol{\alpha}$, let

$$\Sigma^{\boldsymbol{\alpha}} = \left\{\sigma : V \to \{1, \dots, q\} \,\middle|\, |\sigma^{-1}(i)| = \alpha_i n \text{ for } i = 1, \dots, q\right\},$$

that is, configurations in $\Sigma^{\boldsymbol{\alpha}}$ assign $\alpha_i n$ vertices the spin i. We will be interested in the total weight $Z_G^{\boldsymbol{\alpha}}$ of configurations in $\Sigma^{\boldsymbol{\alpha}}$, namely

$$Z_G^{\boldsymbol{\alpha}} = \textstyle\sum_{\sigma \in \Sigma^{\boldsymbol{\alpha}}} w_G(\sigma).$$

We study $Z_G^{\boldsymbol{\alpha}}$ by looking at the moments $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$ and $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha}})^2]$, where the expectation is over the distribution of the random Δ-regular bipartite graph, from hereon denoted by $\mathcal{G}$.

Denote the leading term of the first and second moments as:

$$\begin{aligned}\Psi_1(\boldsymbol{\alpha}) = \Psi_1^{\mathbf{B}}(\boldsymbol{\alpha}) &:= \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\big[Z_G^{\boldsymbol{\alpha}}\big].\\ \Psi_2(\boldsymbol{\alpha}) = \Psi_2^{\mathbf{B}}(\boldsymbol{\alpha}) &:= \lim_{n\to\infty} \frac{1}{n} \log \mathbf{E}_{\mathcal{G}}\Big[\big(Z_G^{\boldsymbol{\alpha}}\big)^2\Big].\end{aligned} \tag{248}$$

8.2 Results for general spin systems

8.2.1 Reformulating the first moment

We state an analog of Theorem 5 for the class of random Δ-regular graphs. For random Δ-regular graphs, we will need the *depth-one tree recursions* stated in Chapter 1:

$$\widehat{R}_i \propto \Big(\sum_{j=1}^{q} B_{ij} R_j\Big)^{\Delta-1}. \tag{7}$$

The fixpoints of the tree recursions are those $\mathbf{r} = (R_1, \dots, R_q)$ such that: $\widehat{R}_i \propto R_i$ for all $i \in [q]$. We refer to a fixpoint $\mathbf{r}$ of the tree recursions as *Jacobian attractive* if the Jacobian at $\mathbf{r}$ has spectral radius less than 1.

It will also be useful to reformulate the function Ψ_1 into the function Φ_1, as we did for random bipartite graphs. The function Φ_1 for random regular graphs takes the form

$$\Phi_1(\mathbf{r}) = \frac{\Delta}{2} \ln\Big(\sum_{i=1}^{q}\sum_{j=1}^{q} B_{ij} R_i R_j\Big) - (\Delta-1)\ln\Big(\sum_{i=1}^{q} R_i^{\Delta/(\Delta-1)}\Big), \tag{249}$$

where $\mathbf{r} = (R_1, \dots, R_q)^{\intercal} \geq 0$. Let $p := \Delta/(\Delta-1)$. Note that (249) has the following appealing form

$$\exp(2\Phi_1(\mathbf{r})/\Delta) = \frac{\mathbf{r}^{\intercal}\mathbf{B}\mathbf{r}}{\|\mathbf{r}\|_p^2}, \tag{250}$$

where $\|\mathbf{r}\|_p = (\sum_{i=1}^{n} R_i^p)^{1/p}$.

Lemma 104. *There is a one-to-one correspondence between the fixpoints of the depth-one tree recursions and the critical points of Φ_1 (both considered for $R_i \geq 0$ in the projective space, that is, up to scaling by a constant). The following transformation $\mathbf{r} \mapsto \boldsymbol{\alpha}$ given by:*

$$\alpha_i = R_i^{\Delta/(\Delta-1)} / \textstyle\sum_i R_i^{\Delta/(\Delta-1)} \tag{251}$$

yields a one-to-one-to-one correspondence between the critical points of $\Phi_1(\mathbf{r})$ and the critical points of $\Psi_1(\boldsymbol{\alpha})$ (in the region defined by $\alpha_i \geq 0$ and $\sum_i \alpha_i = 1$). Moreover, for the corresponding critical points $\mathbf{r}$ and $\boldsymbol{\alpha}$ one has

$$\Phi_1(\mathbf{r}) = \Psi_1(\boldsymbol{\alpha}). \tag{252}$$

Finally, the local maxima of Φ_1 and Ψ_1 happen at the critical points (that is, there are no local maxima on the boundary).

The proof of Lemma 104 follows the same arguments as the proof of Theorem 5 in Section 3.2 with C_j identified with R_j and β_j identified with α_j and is therefore omitted.

8.2.2 Second-moment analysis in the semi-translation invariant uniqueness regime

For general models on random regular graphs, we will be able to analyze the second moment in the semi-translation invariant uniqueness regime, i.e., when the system of equations

$$R_i \propto \Big(\sum_{j=1}^{q} B_{ij} C_j\Big)^{\Delta-1}, \quad C_j \propto \Big(\sum_{j=1}^{q} B_{ij} R_i\Big)^{\Delta-1} \tag{8}$$

has a unique positive solution (up to scaling the values of the R_i's and C_j's). In particular, this implies (trivially) that the depth-one tree recursions (7) have a unique fixed point. By Lemma 104, we have that the maximizer of the function $\Psi_1(\boldsymbol{\alpha})$ is unique, i.e., there is unique dominant phase $\boldsymbol{\alpha}^*$.

We prove the following lemma.

Lemma 105. *Suppose that $\mathbf{B}$ specifies a model in the semi-translation invariant uniqueness regime. Then for the unique dominant phase $\boldsymbol{\alpha}^*$, it holds that $\Psi_2(\boldsymbol{\alpha}^*) = 2\Psi_1(\boldsymbol{\alpha}^*)$.*

Proof. We first show that $\exp(2\Psi_1(\boldsymbol{\alpha}^*)/\Delta) = \|\mathbf{B}\|_{p\to\Delta}$. From (250) and (252), we obtain:

$$\max_{\boldsymbol{\alpha}} \exp(2\Psi_1(\boldsymbol{\alpha})/\Delta) = \max_{\mathbf{r}} \frac{\mathbf{r}^\intercal \mathbf{B} \mathbf{r}}{\|\mathbf{r}\|_p^2} \leq \max_{\mathbf{r},\mathbf{c}} \frac{\mathbf{r}^\intercal \mathbf{B} \mathbf{c}}{\|\mathbf{r}\|_p \|\mathbf{c}\|_p}. \tag{253}$$

Note that the last inequality is trivial; we just enlarged the region we consider. By Theorem 5, the maximum of the rhs is achieved at a semi-translation invariant fixpoint. Since $\mathbf{B}$ specifies a model in the semi-translation invariant uniqueness regime, this must coincide with the translation invariant fixpoint and hence the maximum in the rhs of (253) must occur at $\mathbf{r} = \mathbf{c}$. We thus obtain that (253) is satisfied at equality. On the other hand, the rhs in (253) is equal to $|\mathbf{B}\|_{p\to\Delta}$, while the lhs is $\exp(2\Psi_1(\boldsymbol{\alpha}^*)/\Delta)$, thus proving the desired claim.

Using that the second moment corresponds to a spin model with interaction matrix $\mathbf{B} \otimes \mathbf{B}$, by the same token, one has the bounds

$$\Psi_2(\boldsymbol{\alpha}^*) \le \max_{\boldsymbol{\alpha}\in\triangle_q} \exp(\Psi_2(\boldsymbol{\alpha})/\Delta) \le \max_{\boldsymbol{\alpha}\in\triangle_{q^2}} \exp(\Psi_1^{\mathbf{B}\otimes\mathbf{B}}(\boldsymbol{\alpha})/\Delta) \le \|\mathbf{B}\otimes\mathbf{B}\|_{p\to\Delta}, \tag{254}$$

Note that from $\mathbf{E}[X^2] \ge (\mathbf{E}[X])^2$, we have (trivially) $\Psi_2(\boldsymbol{\alpha}^*) \ge 2\Psi_1(\boldsymbol{\alpha}^*)$. Since $\|\mathbf{B}\otimes\mathbf{B}\|_{p\to\Delta} = \|\mathbf{B}\|_{p\to\Delta}^2$ ([7, Proposition 10.3]), we obtain that (254) holds at equality. This implies $\Psi_2(\boldsymbol{\alpha}^*) = 2\Psi_1(\boldsymbol{\alpha}^*)$, as wanted. □

8.3 Results for ferromagnetic spin systems

For ferromagnetic models on random regular graphs, we will use the Cholesky decomposition of the interaction matrix $\mathbf{B}$ to obtain a matrix norm formulation of the first moment. This will allow us to obtain the following analogue of Theorem 4, which we prove in Section 8.3.1. We restate the relevant theorems.

Theorem 8. *For a ferromagnetic model,*

$$\max_{\boldsymbol{\alpha}} \Psi_2(\boldsymbol{\alpha}) = 2\max_{\boldsymbol{\alpha}} \Psi_1(\boldsymbol{\alpha}).$$

Specifically, for dominant phases $\boldsymbol{\alpha}$*,* $\Psi_2(\boldsymbol{\alpha}) = 2\Psi_1(\boldsymbol{\alpha})$*.*

We will also obtain the following analogue of Theorem 2 in Section 8.3.2.

Theorem 7. *For a ferromagnetic model, Jacobian attractive fixpoints of the (depth-one) tree recursions are in one-to-one correspondence with the Hessian local maxima of* Ψ_1*.*

The above connection fails for antiferromagnetic models, e.g., for the antiferromagnetic Potts model the uniform distribution is a global maximum but it is not a stable fixpoint of the tree recursions for small enough temperature. (In fact, for antiferromagnetic models every solution of the depth-one tree recursions is a local maximum, see Remark 17.)

8.3.1 Second-moment analysis

Since we assume that $\mathbf{B}$ is ferromagnetic we have $\mathbf{B} = \hat{\mathbf{B}}^{\intercal}\hat{\mathbf{B}}$ and hence we can write

$$\exp(\Phi_1(\mathbf{r})/\Delta) = \frac{\|\hat{\mathbf{B}}\mathbf{r}\|_2}{\|\mathbf{r}\|_p}, \tag{255}$$

where Φ_1 was defined in (249). The next lemma describes the connection between Φ_1 and Ψ_1. We note that Φ_1 is not a reparameterization of Ψ_1, however they do agree at the critical points. This is sufficient for our purpose: to understand the maxima of Ψ_1 it is enough to understand the maxima of Φ_1. The maximization

$$\max_{\mathbf{r}\geq 0} \frac{\|\hat{\mathbf{B}}\mathbf{r}\|_2}{\|\mathbf{r}\|_p} = \max_{\mathbf{r}} \frac{\|\hat{\mathbf{B}}\mathbf{r}\|_2}{\|\mathbf{r}\|_p} = \|\hat{\mathbf{B}}\|_{p\to 2}, \tag{256}$$

is the induced $p \to 2$ matrix norm of $\hat{\mathbf{B}}$. The first equality in (256) follows from the fact that the maximum on the right-hand-side of (250) is achieved for non-negative $\mathbf{r}$ (this follows from the fact that $\mathbf{B}$ has non-negative entries).

Lemma 5 allows us to reexpress the optimization problem associated with the first moment in terms of matrix norms.

Lemma 106. *Let* $\mathbf{B}$ *be the interaction matrix of a ferromagnetic spin system. We have*

$$\max_{\boldsymbol{\alpha}} \Psi_1(\boldsymbol{\alpha}) = \Delta \ln \|\hat{\mathbf{B}}\|_{\frac{\Delta}{\Delta-1}\to 2}.$$

Proof. Using Lemma 5 and equations (255) and (256), we obtain

$$\max_{\boldsymbol{\alpha}} \exp(\Psi_1(\boldsymbol{\alpha})/\Delta) = \max_{\mathbf{r}} \exp(\Phi_1(\mathbf{r})/\Delta) = \|\hat{\mathbf{B}}\|_{p\to 2}.$$

□

Recall, the definition of Ψ_2 (see (248)) corresponding to the leading term of the second moment. A key fact is that Ψ_2 is given by a constrained first moment calculation on a "paired-spin" model where the interaction matrix in this model is the tensor product of the original interaction matrix with itself (see Remark 19 in Section 8.4). The second moment considers a pair of configurations, say σ and σ', which are constrained to have a given phase $\boldsymbol{\alpha}$. We capture this constraint using a vector $\boldsymbol{\gamma}$ corresponding to the overlap between σ and σ', in particular, γ_{ij} is the number of vertices with spin i in σ and spin j in σ'.

Recall, $\Psi_1^{\mathbf{B}}$ indicates the dependence of the function Ψ_1 on the interaction matrix $\mathbf{B}$; to simplify the notation we will drop the exponent if it is $\mathbf{B}$. More precisely,

$$\Psi_2(\boldsymbol{\alpha}) = \max_{\boldsymbol{\gamma}} \Psi_1^{\mathbf{B}\otimes\mathbf{B}}(\boldsymbol{\gamma}), \tag{257}$$

where the optimization in (257) is constrained to $\boldsymbol{\gamma}$ such that

$$\textstyle\sum_i \gamma_{ik} = \alpha_k \quad \text{and} \quad \sum_k \gamma_{ik} = \alpha_i. \tag{258}$$

Ignoring the two constraints in (258) can only increase the value of (257) and hence

$$\max_{\boldsymbol{\alpha}} \exp(\Psi_2(\boldsymbol{\alpha})/\Delta) \leq \max_{\boldsymbol{\gamma}} \exp(\Psi_1^{\mathbf{B}\times\mathbf{B}}(\boldsymbol{\gamma})/\Delta) = \|\hat{\mathbf{B}} \otimes \hat{\mathbf{B}}\|^2_{\frac{\Delta}{\Delta-1}\to 2}. \tag{259}$$

For induced norms $\|\cdot\|_{p\to q'}$ with $p \leq q'$ it is known (Proposition 10.1 in [7]) that

$$\|\hat{\mathbf{B}} \otimes \hat{\mathbf{B}}\|_{p\to q'} = \|\hat{\mathbf{B}}\|_{p\to q'} \|\hat{\mathbf{B}}\|_{p\to q'}. \tag{260}$$

Now we are ready to prove Theorem 8.

Proof of Theorem 8. Combining Lemma 106 and equations (259),(260) we obtain:

$$\exp(\Psi_2(\boldsymbol{\alpha})/\Delta) = \max_{\boldsymbol{\gamma}} \exp(\Psi_1^{\mathbf{B}\times\mathbf{B}}(\boldsymbol{\gamma})/\Delta) \leq \|\hat{\mathbf{B}}\|^2_{\frac{\Delta}{\Delta-1}\to 2} = 2\max_{\boldsymbol{\alpha}} \exp(\Psi_1(\boldsymbol{\alpha})/\Delta).$$

This proves that if $\boldsymbol{\alpha}$ maximizes Ψ_1, we have $\Psi_2(\boldsymbol{\alpha}) \leq 2\Psi_1(\boldsymbol{\alpha})$. The reverse inequality is trivial, yielding Theorem 8. $\square$

Remark 16. *We will illustrate the necessity of the ferromagnetism assumption in Theorem 8 by giving an example of an antiferromagnetic model for which the second moment fails. Consider proper* 3*-colorings of random* 10*-regular graphs. As the size of the graph goes to infinity the probability of it being* 3 *colorable goes to zero. The intuitive effect of this is that to achieve a large value in the "paired-spin" model it is better to correlate the coordinates to agree. In terms of* Ψ_1 *and* Ψ_2 *we have that the maximum in the first moment is achieved for* $\alpha_1 = \alpha_2 = \alpha_3 = 1/3$ *with* $\Psi_1 = 5\ln 2 - 4\ln 3 < 0$*. To obtain a lower bound on the maximum in the second moment we take* $\gamma_{11} = \gamma_{22} = \gamma_{33} = 1/3$*, which yields* $\Psi_2 = \Psi_1 > 2\Psi_1$*. The argument actually applies whenever* $\Psi_1 < 0$ *(for models whose interaction matrices have* 0*'s and* 1*'s). By continuity (taking small* B *in antiferromagnetic Potts model) one can obtain an example of a model without hard constraints for which the second moment fails.*

8.3.2 Jacobian stability and Hessian maxima

The proof of Theorem 7 is essentially identical to that of Theorem 2 after identifying C_j's with R_j's and β_j's with α_j's. Here, we give an overview of the proof including the important formulas.

Our starting point is the one-to-one correspondence between fixpoints of the tree recursions and the critical points of Ψ_1, as given in Lemma 104. We show, roughly, that the stability of a fixpoint is equivalent to the local maximality of the corresponding critical point. This will be done by relating the Jacobian of the tree recursions at a fixpoint with the Hessian of Ψ_1 at the corresponding critical point. More precisely, we show that the Jacobian has spectral radius less than 1 (a sufficient condition for stability) if and only in the Hessian is negative definite (a sufficient condition for local maximality). Both constraints on the matrices are independent of the choice of local coordinates (that is, they are invariant under similarity transformations), however to make the connection between the Jacobian and the Hessian apparent we will have to choose the local coordinates very carefully. A further technical complication is that the tree recursions are in the projective space and that the optimization of Ψ_1 is constrained.

We give a high level overview of the Jacobian; the proofs for the Δ-regular case follow the same reasoning as for the bipartite Δ-regular case, see Section 3.4, after simply changing C_j's to R_j's and β_j's to α_j's. Assume that $R_1, \dots, R_q$ is a fixpoint of the tree recursions. Now we consider an infinitesimal perturbation of the fixpoint $R_1 + \varepsilon R_1', \dots, R_q + \varepsilon R_q'$ and see how it is mapped by the tree recursions. Let $\alpha_i := \sum_j B_{ij} R_i R_j$. The right parametrization (choice of local coordinates) is to take $R_i' = r_i R_i / \sqrt{\alpha_i}$, where $r_1, \dots, r_q$ determines the perturbation. Note that $R_i/\sqrt{\alpha_i}$ depends on the fixpoint. The tree recursions map (in the projective space) the perturbation as follows:

$$\left(R_1 + \varepsilon r_1 \frac{R_1}{\sqrt{\alpha_1}}, \dots, R_q + \varepsilon r_q \frac{R_q}{\sqrt{\alpha_q}}\right) \mapsto \left(R_1 + \varepsilon \hat{r}_1 \frac{R_1}{\sqrt{\alpha_1}}, \dots, R_q + \varepsilon \hat{r}_1 \frac{R_q}{\sqrt{\alpha_q}}\right) + O(\varepsilon^2), \quad (261)$$

where $\hat{r}_i$'s are given by the following linear transformation

$$\hat{r}_i = (\Delta - 1) \sum_{j=1}^{q} \frac{B_{ij} R_i R_j}{\sqrt{\alpha_i \alpha_j}} r_j, \quad (262)$$

and where the r_i's are required to satisfy

$$\sum_{i=1}^{q} \sqrt{\alpha_i} r_i = 0. \tag{263}$$

The condition (263) is invariant under the map (262) and corresponds to choosing the representative of $R_1, \ldots, R_q$ with $\sum_i \sum_j B_{ij} R_i R_j = 1$.

Next we give a high level description of the Hessian; again, this is almost identical to Section 3.4 after identifying C_j's with R_j's and β_j's with α_j's. Recall that Ψ_1 is a function of $\alpha_1, \ldots, \alpha_q$. There is an alternative parameterization of Ψ_1: instead of $\alpha_1, \ldots, \alpha_q$ (restricted to $\sum \alpha_i = 1$) we use $R_1, \ldots, R_q$ (restricted to $\sum_i \sum_j B_{ij} R_i R_j = 1$) and use the following

$$\alpha_i = \sum_j B_{ij} R_i R_j \quad \text{for all } i \in [q]. \tag{264}$$

Every $\boldsymbol{\alpha}$ can be achieved using parameterization by $\mathbf{R}$. Let $\alpha_1, \ldots, \alpha_q$ be a critical point of Ψ_1 and let $R_1, \ldots, R_q$ satisfy (264). We are going to evaluate Ψ_1 in a small neighborhood around $\alpha_1, \ldots, \alpha_q$. It is equivalent (and easier to understand) to perturb the $R_1, \ldots, R_q$ to $R_1 + \varepsilon R'_1, \ldots, R_q + \varepsilon R'_q$ and evaluate at the point given by (264). Again, the correct parameterization is to take $R'_i = r_i R_i / \sqrt{\alpha_i}$. This yields the following expression for the value of Ψ_1 at the perturbed point

$$\Psi_1(\alpha_1, \ldots, \alpha_q) + \varepsilon^2 \sum_{i=1}^{q} \left(r_i + \sum_{j=1}^{q} \frac{B_{ij} R_i R_j}{\sqrt{\alpha_i \alpha_j}} r_j \right) \left(\sum_{j=1}^{q} (\Delta - 1) \frac{B_{ij} R_i R_j}{\sqrt{\alpha_i \alpha_j}} r_j - r_i \right) + O(\varepsilon^3). \tag{265}$$

Note that there is no linear term, since we are at a critical point. Recall that the α_i have to satisfy $\sum_i \alpha_i = 1$ which corresponds to the restriction (263).

Now we are ready to prove Theorem 2. Let $\mathbf{L}$ be a linear map such that the Jacobian (262) is $(\Delta - 1)\mathbf{L}$. The Hessian of Ψ_1 is then $(\mathbf{I} + \mathbf{L})((\Delta - 1)\mathbf{L} - \mathbf{I})$. Finally let S be the linear subspace defined by (263).

Proof of Theorem 7. The constraint for the fixpoint to be Jacobian attractive is that $(\Delta - 1)L$ on S has spectral radius less than 1, see equation (261). The constraint for the critical point to be Hessian maximum is that the eigenvalues of $(I + L)((\Delta - 1)L - I)$ on S are negative, see equation (265).

Note that L is symmetric and if $\mathbf{B}$ is positive semidefinite then $\mathbf{L}$ is positive semidefinite (since $\mathbf{L}$ is congruent to $\mathbf{B}$; $\mathbf{L}$ is obtained by multiplying $\mathbf{B}$ by a diagonal matrix on the left and on the right). Hence $\mathbf{L}$ has non-negative real spectrum. Note that S is invariant under $\mathbf{L}$ and hence the spectrum of $\mathbf{L}$ on S is a subset of the spectrum of $\mathbf{L}$ (it is still non-negative real; the restriction wiped out the eigenvalue 1).

The constraint for the fixpoint to be Jacobian attractive, in terms of eigenvalues, is: for each eigenvalue x of $\mathbf{L}$ on S

$$(\Delta - 1)|x| < 1. \tag{266}$$

The constraint for the critical point to be Hessian maximum, in terms of eigenvalues, is: for each eigenvalue x of $\mathbf{L}$ on S

$$(1 + x)((\Delta - 1)x - 1) < 0. \tag{267}$$

Note that conditions (266) and (267) are equivalent (since $x \geq 0$). □

Remark 17. *For antiferromagnetic models every critical point of Ψ_1 is a local maximum. Indeed, we need only to prove that equation* (267) *is satisfied for every critical point. The matrix $\mathbf{L}$ has non-negative entries hence* 1 *is the largest eigenvalue and all the other eigenvalues have magnitude less than* 1 *(since $\mathbf{B}$ is ergodic). Moreover the matrix $\mathbf{L}$ has the same signature as $\mathbf{B}$ (since they are congruent) and hence the eigenvalues other than* 1 *are negative. These* 2 *facts imply* (267).

Remark 18. *Note that one direction of the implication in Theorem 2, namely, that a Jacobian attractive fixpoint is Hessian local maximum, holds for every model (without the ferromagnetism assumption), since* (266) *always implies* (267)*. However, for the reverse implication, the ferromagnetic assumption is essential. For example, in an antiferromagnetic model, by Remark 17, every critical point is a local maximum. For the antiferromagnetic Potts model, the only critical point is the uniform vector and hence it is always a local maximum for every value of B. On the other hand, it is straightforward to check that for the antiferromagnetic Potts model the uniform fixpoint is Jacobian unstable when $B < \frac{\Delta - q}{\Delta}$.*

8.4 Expressions for Ψ_1 *and* Ψ_2

In this section, we derive expressions for the first and second moments of $Z_G^{\boldsymbol{\alpha}}$, which will allow us to derive explicit expressions for the functions $\Psi_1(\boldsymbol{\alpha})$ and $\Psi_2(\boldsymbol{\alpha})$. Our exposition here is such that it provides a straightforward alignment with the analogous expressions in Chapter 2. The minor differences are due to the model of Δ-regular random graphs, which in this chapter is the pairing model $\mathcal{G} := \mathcal{G}_{n,\Delta}$ defined in Section 8.1.1.

Recall, $\triangle_t$ denotes the simplex

$$\triangle_t = \{(x_1, x_2, \dots, x_t) \in \mathbb{R}^t \mid \textstyle\sum_{i=1}^t x_i = 1 \text{ and } x_i \geq 0 \text{ for } i = 1, \dots, t\}. \tag{268}$$

Let $G \sim \mathcal{G}$. For a configuration $\sigma : V \to \{1, \dots, q\}$, we shall denote the set of vertices assigned color i by $\sigma^{-1}(i)$. For $\boldsymbol{\alpha} \in \triangle_q$ and $n\boldsymbol{\alpha} \in \mathbb{Z}^q$, let

$$\Sigma^{\boldsymbol{\alpha}} = \{\sigma : V \to \{1, \dots, q\} \,\big|\, |\sigma^{-1}(i)| = \alpha_i n, \text{ for } i = 1, \dots, q\},$$

that is, $\Sigma^{\boldsymbol{\alpha}}$ is the set of configurations σ which assign $\alpha_i n$ vertices of V the color i, for each $i \in [q]$. We are interested in the total weight $Z_G^{\boldsymbol{\alpha}}$ of configurations in $\Sigma^{\boldsymbol{\alpha}}$, namely

$$Z_G^{\boldsymbol{\alpha}} = \textstyle\sum_{\sigma \in \Sigma^{\boldsymbol{\alpha}}} w_G(\sigma).$$

Note that $Z_G^{\boldsymbol{\alpha}}$ is a r.v., and as indicated earlier, we will look at its moments $\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$ and $\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha}})^2]$.

We begin with the first moment. For $\sigma \in \Sigma^{\boldsymbol{\alpha}}$ and $i, j \in [q]$, let $e_{ij}n$ denote the number of edges matching vertices in $\sigma^{-1}(i)$ and $\sigma^{-1}(j)$. Clearly, $e_{ij} = e_{ji}$. It will be notationally convenient to reparameterize the variables e_{ij} as follows: for $i \neq j$ we set $e_{ij} = \Delta x_{ij}$ and for $i = j$ we set $e_{ii} = \Delta x_{ii}/2$. Note that the x_{ij} are r.v., since they depend on the choice of the random graph G.For future use, we denote by $\mathbf{x}_G(\sigma)$ the random vector $(x_{11}, \dots, x_{qq})$.

The number of perfect matchings between $2n$ vertices will be denoted by $(2n)!!$. It is well known and easy to see that $(2n)!! = (2n)!/(n!2^n)$. Under the convention that $0^0 \equiv 1$, and using the notation of Section 2.2.1 we then have

$$\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}] = \binom{n}{\boldsymbol{\alpha}n} \sum_{\mathbf{x}} \prod_i \binom{\Delta\alpha_i n}{\Delta \mathbf{x}_{i\cdot} n} \frac{\left[\prod_{i \neq j} (\Delta x_{ij} n)!\right]^{1/2} \prod_i (\Delta x_{ii} n)!!}{(\Delta n)!!} \prod_{i,j} B_{ij}^{\Delta x_{ij} n/2}, \tag{269}$$

where the sum ranges over all the possibles values of the random vector $\mathbf{x}_G(\sigma)$, that is, $\mathbf{x} = (x_{11}, \ldots, x_{qq})$ satisfying:

$$\begin{aligned} \textstyle\sum_j x_{ij} &= \alpha_i \quad \big(\forall i \in [q]\big), \\ x_{ij} &= x_{ji} \geq 0 \quad \big(\forall i, j \in [q]\big). \end{aligned} \tag{270}$$

The first line in (269) accounts for the cardinality of $\Sigma^{\boldsymbol{\alpha}}$, while the second line is $\mathbf{E}_{\mathcal{G}}[w_G(\sigma)]$ for a fixed $\sigma \in \Sigma^{\boldsymbol{\alpha}}$, since by symmetry we may focus on any fixed σ. The first product is the number of ways to choose a partition of the points which is consistent with the values prescribed by $\mathbf{x}$, the fraction is the probability that the random matching connects the points as prescribed, and the last product is the weight of the configuration σ conditioned on $\mathbf{x}$.

We next deal with the second moment of $Z_G^{\boldsymbol{\alpha}}$. The desired expression may be derived analogously to (23). For $(\sigma_1, \sigma_2) \in \Sigma^{\boldsymbol{\alpha}} \times \Sigma^{\boldsymbol{\alpha}}$, we need to compute $\mathbf{E}_{\mathcal{G}}[w_G(\sigma_1) w_G(\sigma_2)]$. For $i, k \in [q]$, let $\gamma_{ik} n = |\sigma_1^{-1}(i) \cap \sigma_2^{-1}(k)|$. The vector $\boldsymbol{\gamma}$ captures the overlap of the configurations σ_1, σ_2. Denote by $e_{ikjl} n$ the number of edges matching vertices in $\sigma_1^{-1}(i) \cap \sigma_2^{-1}(k)$ and $\sigma_1^{-1}(i) \cap \sigma_2^{-1}(l)$. We reparameterize as follows: for $(i,k) \neq (j,l)$ we set $e_{ikjl} = \Delta y_{ikjl}$ and for $(i,k) = (j,l)$ we set $e_{ikjl} = \Delta y_{ikjl}/2$. Using the notation of Section 2.2.1, we have:

$$\begin{aligned} &\mathbf{E}_{\mathcal{G}}[(Z_G^{\boldsymbol{\alpha}})^2] = \sum_{\boldsymbol{\gamma}} \binom{n}{\boldsymbol{\gamma} n} \\ &\sum_{\mathbf{y}} \prod_{i,k} \binom{\Delta \gamma_{ik} n}{\Delta \mathbf{y}_{ik\cdot} n} \frac{\left[\prod_{(i,j) \neq (k,l)} (\Delta y_{ikjl} n)!\right]^{1/2} \prod_{i,k} (\Delta y_{ikik} n)!!}{(\Delta n)!!} \prod_{i,j} \left(B_{ij} B_{kl}\right)^{\Delta y_{ikjl} n/2}, \end{aligned} \tag{271}$$

where the sums range over $\boldsymbol{\gamma} = (\gamma_{11}, \ldots, \gamma_{qq})$, $\mathbf{y} = (y_{1111}, \ldots, y_{qqqq})$ satisfying

$$\begin{aligned} \textstyle\sum_k \gamma_{ik} &= \alpha_i \quad \big(\forall i \in [q]\big), \\ \textstyle\sum_i \gamma_{ik} &= \alpha_k \quad \big(\forall k \in [q]\big), \\ \textstyle\sum_{j,l} y_{ikjl} &= \gamma_{ik} \quad \big(\forall (i,k) \in [q]^2\big) \end{aligned} \tag{272}$$

$$\gamma_{ik} \geq 0 \quad \big(\forall (i,k) \in [q]^2\big), \quad y_{ikjl} = y_{jlik} \geq 0 \quad \big(\forall (i,k,j,l) \in [q]^4\big).$$

The sums in (269) and (271) are typically exponential in n. The most critical component of our arguments is to find the quantitative structure of configurations which determine the exponential order of the moments. Formally, we study the limits of $\frac{1}{n}\log \mathbf{E}_{\mathcal{G}}\left[Z_G^{\boldsymbol{\alpha}}\right]$ and $\frac{1}{n}\log \mathbf{E}_{\mathcal{G}}\left[(Z_G^{\boldsymbol{\alpha}})^2\right]$ as $n \to \infty$. These limits are obtainable using Stirling's approximation. In particular, we shall use that for a constant $c > 0$, we have

$$\frac{1}{n}\log\left[(cn)!\right] \sim c\log n + c\log c - c \quad \text{and} \quad \frac{1}{n}\log\left[(2cn)!!\right] \sim c\log n + c\log c + c(\log 2 - 1). \quad (273)$$

Under the usual conventions that $\ln 0 \equiv -\infty$ and $0 \ln 0 \equiv 0$, the above formulas are correct even in the degenerate case $c = 0$.

We now derive asymptotics for the first moment $\mathbf{E}_{\mathcal{G}}\left[Z_G^{\boldsymbol{\alpha}}\right]$ in order to obtain the function $\Psi_1(\boldsymbol{\alpha})$. Applying (273) yields:

$$\Psi_1(\alpha) := \lim_{n\to\infty}\frac{1}{n}\log \mathbf{E}_{\mathcal{G}}\left[Z_G^{\alpha}\right] = \max_{\mathbf{X}} \Phi_1(\boldsymbol{\alpha}, \mathbf{X}), \quad (274)$$

$$\begin{aligned} \text{where} \quad \Phi_1(\alpha, \mathbf{x}) &:= (\Delta - 1) f_1(\boldsymbol{\alpha}) + \Delta g_1(\mathbf{x}) \\ f_1(\boldsymbol{\alpha}) &:= \textstyle\sum_i \alpha_i \ln \alpha_i \\ g_1(\mathbf{x}) &:= \tfrac{1}{2}\textstyle\sum_{i,j} x_{ij} \ln B_{ij} - \tfrac{1}{2}\textstyle\sum_{i,j} x_{ij}\ln x_{ij}, \end{aligned}$$

defined on the region (270).

Completely analogously, for the second moment we obtain:

$$\Psi_2(\alpha) := \lim_{n\to\infty}\frac{1}{n}\log \mathbf{E}_{\mathcal{G}}\left[(Z_G^{\boldsymbol{\alpha}})^2\right] = \max_{\boldsymbol{\gamma}}\max_{\mathbf{y}} \Phi_2(\boldsymbol{\gamma}, \mathbf{y}), \quad (275)$$

$$\begin{aligned} \text{where} \quad \Phi_2(\boldsymbol{\gamma}, \mathbf{y}) &:= (\Delta - 1) f_2(\boldsymbol{\gamma}) + \Delta g_2(\mathbf{y}) \\ f_2(\boldsymbol{\gamma}) &:= \textstyle\sum_{i,k} \gamma_{ik} \ln \gamma_{ik} \\ g_2(\mathbf{y}) &:= \tfrac{1}{2}\textstyle\sum_{i,k,j,l} y_{ikjl} \ln(B_{ij}B_{kl}) - \tfrac{1}{2}\textstyle\sum_{i,k,j,l} y_{ikjl}\ln y_{ikjl}, \end{aligned}$$

defined on the region (272).

Remark 19. *As for random bipartite graphs, we can interpret the second moment as the first moment of a paired-spin model with interaction matrix* $\mathbf{B} \otimes \mathbf{B}$*. Indeed, in* (275)*, we can interpret* $B_{ij}B_{kl}$ *as the activity between the paired spins* (i, k) *and* (j, l)*, thus giving the desired alignment.*

8.5 Small subgraph conditioning method

In this section, we give the outline for the proof of the following lemma.

Lemma 107. *For every ferromagnetic model* $\mathbf{B}$*, if* $\boldsymbol{\alpha}$ *is dominant and the corresponding fixpoint is Jacobian attractive (c.f. Section 8.3) with probability* $1 - o(1)$ *over the choice of the graph* $G \sim \mathcal{G}_{n,\Delta}$*, it holds that* $Z_G^{\boldsymbol{\alpha}} \geq \frac{1}{n}\mathbf{E}\big[Z_G^{\boldsymbol{\alpha}}\big]$.

The proof of Lemma 107 is a minor modification of the arguments in Section 4.2.1 which were carried out for random Δ-regular bipartite graphs. Here, we just need to account for the Δ-regular case which turns out to be completely analogous. As such, we give only the statements of the necessary lemmas and briefly sketch the proofs.

The main tool we are going to use is the small graph conditioning method, which we described in Section 2.3. For convenience of the reader, we reiterate the statement of the theorem (in a slightly weaker form this time; namely, we only state the conclusion (C1)).

Theorem 108. *For* $i = 1, 2, \ldots,$ *let* $\lambda_i > 0$ *and* $\delta_i > -1$ *be constants and assume that for each* n *there are random variables* X_{in}*,* $i = 1, 2, \ldots,$ *and* Y_n*, all defined on the same probability space* $\mathcal{G} = \mathcal{G}_n$ *such that* X_{in} *is non-negative integer valued,* $Y_n \geq 0$ *and* $\mathbf{E}\big[Y_n\big] > 0$ *(for* n *sufficiently large). Furthermore, the following hold:*

(A1) $X_{in} \xrightarrow{d} Z_i$ *as* $n \to \infty$*, jointly for all* i*, where* $Z_i \sim \mathrm{Po}(\lambda_i)$ *are independent Poisson random variables;*

(A2) *for every finite sequence* $j_1, \ldots, j_m$ *of non-negative integers,*

$$\frac{\mathbf{E}_{\mathcal{G}}\big[Y_n[X_{1n}]_{j_1} \cdots [X_{mn}]_{j_m}\big]}{\mathbf{E}_{\mathcal{G}}\big[Y_n\big]} \to \prod_{i=1}^{m} \big(\lambda_i(1+\delta_i)\big)^{j_i} \quad \text{as } n \to \infty; \tag{276}$$

(A3) $\sum_i \lambda_i \delta_i^2 < \infty$*;*

(A4) $\mathbf{E}_{\mathcal{G}}\big[Y_n^2\big]/\big(\mathbf{E}_{\mathcal{G}}[Y_n]\big)^2 \leq \exp\big(\sum_i \lambda_i \delta_i^2\big) + o(1)$ *as* $n \to \infty$*;*

Let $r(n)$ *be a function such that* $r(n) \to 0$ *as* $n \to \infty$*. It holds that* $Y_n > r(n)\mathbf{E}_{\mathcal{G}}\big[Y_n\big]$ *asymptotically almost surely.*

To obtain Lemma 107, it should be clear that it suffices to verify the conditions of Theorem 108 for the random variables $Z_G^{\boldsymbol{\alpha}}$. Recall, we restrict our attention to $\boldsymbol{\alpha}$ which are Hessian dominant. For $G \sim \mathcal{G}(n, \Delta)$, let $X_i = X_{in}$ be the number of cycles of length i in G, $i = 1, 2, \ldots$.

The most technical part of this verification is assumption (A4) which requires computing the precise asymptotics of the moments. This in turn reduces to certain determinants which are not completely trivial. Nevertheless, the arguments have been carried out in full generality in Sections 4.3 and 4.4. The only minor modification required in the present case is to account for the random Δ-regular graph setting instead of the bipartite random Δ-regular graph setting studied in Sections 4.3 and 4.4. The arguments there extend in a straightforward manner. We thus obtain the following lemmas.

Lemma 109. *Assumption (A1) holds with* $\lambda_i = \frac{(\Delta-1)^i}{2i}$.

Lemma 110. *Assumption (A2) holds with* $\delta_i = \sum_{j=1}^{q-1} \mu_j^i$, *where* $\mu_1, \mu_2, \ldots, \mu_{q-1}$ *are the eigenvalues different than 1 of the matrix*

$$\mathbf{M} = \left\{ \frac{B_{ij} R_i R_j}{\sqrt{\alpha_i \alpha_j}} \right\}_{i,j=1}^{q}.$$

For Hessian dominant $\boldsymbol{\alpha}$, *it holds that the* μ_i *are positive and strictly smaller than* $1/(\Delta-1)$.

Lemma 111. *Assumption (A3) holds with*

$$\sum_i \lambda_i \delta_i^2 = \prod_{i=1}^{q-1} \prod_{j=1}^{q-1} \big(1 - (\Delta - 1)\mu_i \mu_j\big)^{-1/2},$$

where the μ_i*'s are as in Lemma 110.*

Lemma 112. *For a ferromagnetic model, for all Hessian dominant* $\boldsymbol{\alpha}$ *it holds that*

$$\frac{\mathbf{E}_{\mathcal{G}}\big[(Z_G^{\boldsymbol{\alpha}})^2\big]}{\big(\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]\big)^2} \to \prod_{i=1}^{q-1} \prod_{j=1}^{q-1} \big(1 - (\Delta - 1)\mu_i \mu_j\big)^{-1/2},$$

where μ_i *are as in Lemma 110.*

Proof of Lemma 107. Apply Lemmas 109–112 to obtain that the assumptions of Theorem 108 hold. The lemma follows by the conclusion of Theorem 108, for $r(n) = 1/n$. □

8.6 Proofs of Theorems 9 and 10

Using Lemma 107, we now give the proofs of Theorems 9 and 10.

Proofs of Theorems 9 and 10. Let $\boldsymbol{\alpha}$ be a Hessian dominant phase, whose existence is guaranteed by the assumptions. By Lemma 107, with probability $1-o(1)$ over the choice of the graph, we have $Z_G^{\boldsymbol{\alpha}} \geq \frac{1}{n}\mathbf{E}\big[Z_G^{\boldsymbol{\alpha}}\big]$, which implies $\frac{1}{n}\log Z_G \geq \Psi_1(\boldsymbol{\alpha}) + o(1)$.

Moreover, since the model is ferromagnetic, for Δ-regular graphs G with n vertices, $\frac{1}{n}\log Z_G \geq C$ for some constant $C > -\infty$ (explicitly, one can take $C := \frac{\Delta}{2}\log\max_{i\in[q]} B_{ii}$, see the remarks after the statement of Theorem 9 in the Introduction). We thus obtain

$$\liminf_{n\to\infty}\frac{1}{n}\mathbf{E}_{\mathcal{G}}[\log Z_G] \geq \liminf_{n\to\infty}\big[(1-o(1))\Psi_1(\boldsymbol{\alpha}) + o(1)C\big] = \Psi_1(\boldsymbol{\alpha}).$$

By Jensen's inequality, we also have

$$\limsup_{n\to\infty}\frac{1}{n}\mathbf{E}_{\mathcal{G}}[\log Z_G] \leq \lim_{n\to\infty}\frac{1}{n}\log \mathbf{E}_{\mathcal{G}}[Z_G].$$

All that remains to show is that $\frac{1}{n}\log \mathbf{E}_{\mathcal{G}}[Z_G] = \Psi_1(\boldsymbol{\alpha}) + o(1)$. This is straightforward; if we decompose Z_G as $Z_G = \sum_{\boldsymbol{\alpha}'} Z_G^{\boldsymbol{\alpha}'}$, we obtain $\exp(o(n))\mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}] \geq \mathbf{E}_{\mathcal{G}}[Z_G] \geq \mathbf{E}_{\mathcal{G}}[Z_G^{\boldsymbol{\alpha}}]$. Note the $\exp(o(n))$ is there to allow for dominant phases which are not Hessian.

This concludes the proof of Theorem 9. The proof of Theorem 10 is completely analogous. $\square$

CHAPTER IX

CONCLUSIONS

This thesis studied the complexity of counting in spin systems on graphs of constant maximum degree Δ in terms of phase transitions on the infinite tree $\mathbb{T}_\Delta$. We gave strong NP-hardness results for a wide class of multi-spin systems, including the particularly interesting cases of the k-colorings and the q-state antiferromagnetic Potts models. To obtain these inapproximability results, we analyzed spin systems on random bipartite Δ-regular graphs and identified a natural class of models (ferromagnetic) where our analysis extends to random Δ-regular graphs.

We conclude by listing several problems that are closely related to the present thesis, together with a short discussion of possible approaches to tackle them.

Problem 1. *Give a proof that the uniqueness threshold on the infinite Δ-regular tree for the q-state antiferromagnetic Potts model is at $B = \frac{\Delta - q}{\Delta}$.*

Note that in this thesis we proved that the semi-translation invariant uniqueness threshold is at $B = \frac{\Delta - q}{\Delta}$ (see Lemma 80), but proving that this threshold coincides with the actual uniqueness threshold is harder. In principle, an approach analogous to Jonasson's proof for k-colorings should be possible [41], where it is shown that the tree recursions converge to the uniform fixpoint under arbitrary boundary conditions. It should be noted however that, for the Potts model, the extra parameter B makes these arguments harder to carry out. On a related note, it would be particularly interesting if one could formulate a fairly general criterion to capture the convergence of the tree recursions for a natural class of spin systems, e.g., systems with no hard constraints.

Problem 2. *Analyze the dominant phases on random bipartite Δ-regular graphs for the k-colorings and the q-state antiferromagnetic Potts models when k, q are odd. Note that this would remove the conditions in Theorems 12 and 13 for k, q to be even.*

Recall that analyzing the dominant phases on random bipartite Δ-regular graphs reduces to computing $\|\mathbf{B}\|_{p\to\Delta}$, where $p = \Delta/(\Delta-1)$ and $\mathbf{B}$ is the interaction matrix corresponding to the k-colorings and the q-state antiferromagnetic Potts models. The difficulty of solving this problem can already be seen from our work in Section 6.3.2: the solution to the natural relaxation of the underlying optimization problem is given for all values of q, B, Δ by a triple $(q/2, q/2, 0)$ (see Lemmas 88 and 89). Thus, for odd q, one has to directly tackle the integrality issue. A reasonable approach is to look at the different types of fixpoints of the tree recursions (see Definition 12) and compare the values of the function Φ. While 2-supported fixpoints of the tree recursions are fairly simple to compare, this is no longer the case for 3-supported fixpoints. Local maximum considerations (see Theorem 2) can help rule out certain types of 3-supported fixpoints, but unfortunately not all of them. Experimental evidence suggests that the dominant phases correspond to either a $(\lfloor \frac{q}{2} \rfloor, \lceil \frac{q}{2} \rceil, 0)$-fixpoint or a $(\lfloor \frac{q}{2} \rfloor, \lfloor \frac{q}{2} \rfloor, 1)$-fixpoint of the tree recursions (the latter seems to be relevant only for $q = 3$ and small values of Δ).

Problem 3. *Do dominant phases on random bipartite Δ-regular graphs correspond to* extremal *Gibbs measures on the infinite tree $\mathbb{T}_\Delta$?*

We have already displayed in Chapter 1 that dominant phases correspond to Gibbs measures on the infinite tree $\mathbb{T}_\Delta$. For the hard-core model (as well as antiferromagnetic 2-spin systems) it is known that the Gibbs measures corresponding to dominant phases are in fact extremal; see, e.g., [55] for an alternative characterization of extremality in terms of the reconstruction problem. This is also the case for the ferromagnetic Potts model. In an unpublished result with Daniel Štefankovič and Eric Vigoda, we have also verified the extremality of the Gibbs measures corresponding to the dominant phases for the k-colorings and q-state antiferromagnetic Potts models (for k, q even). It would be interesting to generalize these model-specific results, if possible, to general spin systems (general interaction matrix $\mathbf{B}$). Such a generalization would yield that the properties of Sly's original gadget [63] hold for general spin systems, allowing to obtain connections between the complexity of approximating the partition function on different spin systems (similarly to Theorems 17 and 18). Moreover, it would allow to remove the condition in the general inapproximability Theorem 14 that the dominant phases are permutation-symmetric.

Problem 4. *Obtain generalizations of our results for random (bipartite) regular graphs to graph sequences converging locally to a tree.*

Our analysis of the partition function for random regular graphs exploits the graph distribution (expressions for the moments of the partition function, short cycle distribution). Can one analyze the partition function for graph sequences converging locally to a tree, in the sense of [53, 64, 18]? The approaches therein are based on interpolation schemes rather than the analysis of the moments that we used in this thesis.

Problem 5. *Does the existence of multiple dominant phases on random bipartite regular graphs (which correspond to semi-translation invariant measures on $\mathbb{T}_\Delta$ which are not translation invariant) imply hardness of approximating the partition function?*

Note that the general Theorem 14 applies to antiferromagnetic models, but Theorem 15 hints that the eigenvalue restriction on the interaction matrix $\mathbf{B}$ might not be necessary. To remove the restriction on the eigenvalues of $\mathbf{B}$, one would need to exploit our analysis of random bipartite graphs to devise a different reduction. We remark here that for ferromagnetic models (whose interaction matrices have only positive eigenvalues), it is fairly straightforward to show that all dominant phases correspond to translation invariant measures on $\mathbb{T}_\Delta$.

REFERENCES

[1] Achlioptas, D. and Moore, C., "The chromatic number of random regular graphs," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques* (Jansen, K., Khanna, S., Rolim, J. D. P., and Ron, D., eds.), vol. 3122 of *Lecture Notes in Computer Science*, pp. 219–228, Springer Berlin Heidelberg, 2004.

[2] Aldous, D. and Steele, J., "The objective method: Probabilistic combinatorial optimization and local weak convergence," in *Probability on Discrete Structures* (Kesten, H., ed.), vol. 110 of *Encyclopaedia of Mathematical Sciences*, pp. 1–72, Springer Berlin Heidelberg, 2004.

[3] Alimonti, P. and Kann, V., "Some APX-completeness results for cubic graphs," *Theoretical Computer Science*, vol. 237, no. 12, pp. 123 – 134, 2000.

[4] Bayati, M., Gamarnik, D., and Tetali, P., "Combinatorial approach to the interpolation method and scaling limits in sparse random graphs," *The Annals of Probability*, vol. 41, pp. 4080–4115, 11 2013.

[5] Bender, E. A. and Canfield, E. R., "The asymptotic number of labeled graphs with given degree sequences," *Journal of Combinatorial Theory, Series A*, vol. 24, no. 3, pp. 296–307, 1978.

[6] Benjamini, I. and Schramm, O., "Recurrence of distributional limits of finite planar graphs," *Electron. J. Probab.*, vol. 6, pp. no. 23, 1–13, 2001.

[7] Bennett, G., "Schur multipliers," *Duke Mathematical Journal*, vol. 44, no. 3, pp. 603–639, 1977.

[8] Berge, C., *Espaces Topologiques: Fonctions Multivoques.* (French) Collection Universitaire de Mathématiques, Vol. III, Dunod, Paris, 1959. (In English) *Topological spaces: Including a treatment of multi-valued functions, vector spaces and convexity.* Translated from the French original by E. M. Patterson. Reprint of the 1963 translation. Dover Publications, Inc., Mineola, NY, 1997.

[9] Bollobás, B., "A probabilistic proof of an asymptotic formula for the number of labelled regular graphs," *European Journal of Combinatorics*, vol. 1, no. 4, pp. 311–316, 1980.

[10] Bollobás, B., *Random graphs.* Springer, 1998.

[11] Brightwell, G. R. and Winkler, P., "Random colorings of a Cayley tree," in *Contemporary Combinatorics* (Bollobás, B., ed.), Bolyai Society Mathematical Studies, pp. 247–276, Springer, 2002.

[12] Bulatov, A. and Grohe, M., "The complexity of partition functions," in *Automata, Languages and Programming* (Daz, J., Karhumki, J., Lepist, A., and Sannella, D., eds.), vol. 3142 of *Lecture Notes in Computer Science*, pp. 294–306, Springer Berlin Heidelberg, 2004.

[13] Catlin, P. A., "A bound on the chromatic number of a graph," *Discrete Mathematics*, vol. 22, no. 1, pp. 81–83, 1978.

[14] Coja-Oghlan, A., Efthymiou, C., and Hetterich, S., "On the chromatic number of random regular graphs," *CoRR*, vol. abs/1308.4287, 2013.

[15] de Bruijn, N. G., *Asymptotic methods in analysis*. Dover Publications, Inc., New York, third ed., 1981.

[16] de la Harpe, P., *Topics in Geometric Group Theory*. Chicago Lectures in Mathematics, University of Chicago Press, 2000.

[17] Dembo, A. and Montanari, A., "Gibbs measures and phase transitions on sparse random graphs," *Brazilian Journal of Probability and Statistics*, vol. 24, no. 2, pp. 137–211, 2010.

[18] Dembo, A., Montanari, A., Sly, A., and Sun, N., "The replica symmetric solution for Potts models on d-regular graphs," *Communications in Mathematical Physics*, vol. 327, no. 2, pp. 551–575, 2014.

[19] Dembo, A., Montanari, A., and Sun, N., "Factor models on locally tree-like graphs," *The Annals of Probability*, vol. 41, pp. 4162–4213, 11 2013.

[20] Ding, J., Sly, A., and Sun, N., "Maximum independent sets in random regular graphs," *CoRR*, vol. abs/1310.4787, 2013.

[21] Dobrushin, R. L., "Gibbsian random fields for lattice systems with pairwise interactions," *Functional Analysis and Its Applications*, vol. 2, no. 4, pp. 292–301, 1968.

[22] Dyer, M., Frieze, A., and Jerrum, M., "On counting independent sets in sparse graphs," *SIAM Journal on Computing*, vol. 31, no. 5, pp. 1527–1541, 2002.

[23] Dyer, M., Goldberg, L. A., Greenhill, C., and Jerrum, M., "On the relative complexity of approximate counting problems," in *Approximation Algorithms for Combinatorial Optimization* (Jansen, K. and Khuller, S., eds.), vol. 1913 of *Lecture Notes in Computer Science*, pp. 108–119, Springer Berlin Heidelberg, 2000.

[24] Dyer, M. and Greenhill, C., "The complexity of counting graph homomorphisms," *Random Structrures & Algorithms*, vol. 17, pp. 260–289, Oct. 2000.

[25] Emden-Weinert, T., Hougardy, S., and Kreuter, B., "Uniquely colourable graphs and the hardness of colouring graphs of large girth," *Combinatorics, Probability and Computing*, vol. 7, pp. 375–386, 12 1998.

[26] Galanis, A., Ge, Q., Štefankovič, D., Vigoda, E., and Yang, L., "Improved inapproximability results for counting independent sets in the hard-core model," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques* (Goldberg, L. A., Jansen, K., Ravi, R., and Rolim, J. D., eds.), vol. 6845 of *Lecture Notes in Computer Science*, pp. 567–578, Springer Berlin Heidelberg, 2011.

[27] Galanis, A., Štefankovič, D., and Vigoda, E., "Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models," *CoRR*, vol. abs/1203.2226, 2012.

[28] GALANIS, A., ŠTEFANKOVIČ, D., and VIGODA, E., "Inapproximability for antiferromagnetic spin systems in the tree non-uniqueness region," *CoRR*, vol. abs/1305.2902, 2013.

[29] GEORGII, H.-O., *Gibbs measures and phase transitions*, vol. 9. Walter de Gruyter, 2011.

[30] GOLDBERG, L. A. and JERRUM, M., "Approximating the partition function of the ferromagnetic Potts model," in *Automata, Languages and Programming*, pp. 396–407, Springer, 2010.

[31] GOLDBERG, L. A., JERRUM, M., and PATERSON, M., "The computational complexity of two-state spin systems," *Random Structures & Algorithms*, vol. 23, no. 2, pp. 133–154, 2003.

[32] HÄGGSTRÖM, O., "The random-cluster model on a homogeneous tree," *Probability Theory and Related Fields*, vol. 104, no. 2, pp. 231–253, 1996.

[33] HORN, R. A. and JOHNSON, C. R., *Matrix analysis*. Cambridge University Press, 2012.

[34] ISING, E., "Beitrag zur theorie des ferromagnetismus," *Zeitschrift für Physik*, vol. 31, no. 1, pp. 253–258, 1925.

[35] JANSON, S., "Random regular graphs: asymptotic distributions and contiguity," *Combinatorics, Probability & Computing*, vol. 4, pp. 369–405, 1995.

[36] JANSON, S., ŁUCZAK, T., and RUCINSKI, A., *Random graphs*, vol. 45. John Wiley & Sons, 2011.

[37] JENSEN, T. R. and TOFT, B., *Graph coloring problems*, vol. 39. John Wiley & Sons, 2011.

[38] JERRUM, M. and SINCLAIR, A., "Polynomial-time approximation algorithms for the Ising model," *SIAM Journal on Computing*, vol. 22, no. 5, pp. 1087–1116, 1993.

[39] JERRUM, M. R., VALIANT, L. G., and VAZIRANI, V. V., "Random generation of combinatorial structures from a uniform distribution," *Theoretical Computer Science*, vol. 43, pp. 169–188, 1986.

[40] JOHANSSON, A., "Asymptotic choice number for triangle free graphs," Tech. Rep. 91-5, DIMACS, 1996.

[41] JONASSON, J., "Uniqueness of uniform random colorings of regular trees," *Statistics & Probability Letters*, vol. 57, no. 3, pp. 243–248, 2002.

[42] KELLY, F. P., "Loss networks," *The Annals of Applied Probability*, pp. 319–378, 1991.

[43] LANFORD, O.E., I. and RUELLE, D., "Observables at infinity and states with short range correlations in statistical mechanics," *Communications in Mathematical Physics*, vol. 13, no. 3, pp. 194–215, 1969.

[44] Li, L., Lu, P., and Yin, Y., "Correlation decay up to uniqueness in spin systems," in *Proceedings of the Twenty-fourth Annual ACM-SIAM Symposium on Discrete Algorithms* (Khanna, S., ed.), SODA '13, pp. 67–84, SIAM, 2013.

[45] Liu, J., Lu, P., and Zhang, C., "The complexity of ferromagnetic two-spin systems with external fields," *CoRR*, vol. abs/1402.4346, 2014.

[46] Luenberger, D. G. and Ye, Y., *Linear and nonlinear programming*, vol. 116. Springer, 2008.

[47] Martinelli, F., Sinclair, A., and Weitz, D., "Fast mixing for independent sets, colorings and other models on trees," in *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '04, (Philadelphia, PA, USA), pp. 456–465, Society for Industrial and Applied Mathematics, 2004.

[48] Maslov, V. P. and Fedoryuk, M. V., "Logarithmic asymptotic of the Laplace integrals," *Mathematical Notes*, vol. 30, no. 5, pp. 880–883, 1981.

[49] Mézard, M. and Montanari, A., *Information, physics, and computation.* Oxford University Press, 2009.

[50] Molloy, M. and Reed, B., "Colouring graphs when the number of colours is nearly the maximum degree," in *Proceedings of the Thirty-third Annual ACM symposium on Theory of Computing*, pp. 462–470, ACM, 2001.

[51] Molloy, M. and Reed, B., *Graph colouring and the probabilistic method*, vol. 23. Springer, 2002.

[52] Molloy, M. S. O., Robalewska, H., Robinson, R. W., and Wormald, N. C., "1-factorizations of random regular graphs," *Random Structures & Algorithms*, vol. 10, no. 3, pp. 305–321, 1997.

[53] Montanari, A., Mossel, E., and Sly, A., "The weak limit of Ising models on locally tree-like graphs," *Probability Theory and Related Fields*, vol. 152, no. 1-2, pp. 31–51, 2012.

[54] Moser, R. A. and Tardos, G., "A constructive proof of the general Lovász local lemma," *J. ACM*, vol. 57, pp. 11:1–11:15, Feb. 2010.

[55] Mossel, E., "Survey: Information flow on trees," *DIMACS series in Discrete Mathematics and Theoretical Computer Science*, vol. 63, pp. 155–170, 2004.

[56] Mossel, E., Weitz, D., and Wormald, N., "On the hardness of sampling independent sets beyond the tree threshold," *Probability Theory and Related Fields*, vol. 143, no. 3-4, pp. 401–439, 2009.

[57] Onsager, L., "Crystal statistics. I. A two-dimensional model with an order-disorder transition," *Phys. Rev.*, vol. 65, pp. 117–149, Feb 1944.

[58] Papadimitriou, C. H., *Computational Complexity.* Addison-Wesley, 1994.

[59] Peruggi, F., di Liberto, F., and Monroy, G., "Phase diagrams of the q-state potts model on bethe lattices," *Physica A: Statistical Mechanics and its Applications*, vol. 141, no. 1, pp. 151–186, 1987.

[60] Robinson, R. W. and Wormald, N. C., "Almost all regular graphs are hamiltonian," *Random Structures & Algorithms*, vol. 5, no. 2, pp. 363–374, 1994.

[61] Sinclair, A., Srivastava, P., and Thurley, M., "Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs," in *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '12, pp. 941–953, SIAM, 2012.

[62] Sly, A., "Uniqueness thresholds on trees versus graphs," *The Annals of Applied Probability*, vol. 18, no. 5, pp. pp. 1897–1909, 2008.

[63] Sly, A., "Computational transition at the uniqueness threshold," in *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, FOCS '10, pp. 287–296, IEEE Computer Society, 2010. Preprint of the full version is available from the arXiv at: http://arxiv.org/abs/1005.5584.

[64] Sly, A. and Sun, N., "The computational hardness of counting in two-spin models on d-regular graphs," in *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, FOCS '12, pp. 361–369, IEEE Computer Society, 2012.

[65] Tits, J., "Sur le groupe des automorphismes dun arbre," in *Essays on Topology and Related Topics*, pp. 188–211, Springer Berlin Heidelberg, 1970.

[66] Valiant, L., "The complexity of computing the permanent," *Theoretical Computer Science*, vol. 8, no. 2, pp. 189 – 201, 1979.

[67] Wainwright, M. J. and Jordan, M. I., "Graphical models, exponential families, and variational inference," *Foundations and Trends® in Machine Learning*, vol. 1, no. 1-2, pp. 1–305, 2008.

[68] Weitz, D., "Counting independent sets up to the tree threshold," in *Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing*, STOC '06, pp. 140–149, ACM, 2006.

[69] Wong, R., *Asymptotic approximations of integrals*, vol. 34. SIAM, 2001.

[70] Wormald, N. C., "Models of random regular graphs," *London Mathematical Society Lecture Note Series*, pp. 239–298, 1999.