SOME PROBLEMS INVOLVING H-COLORINGS OF GRAPHS

A Dissertation

Submitted to the Graduate School

of the University of Notre Dame

in Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

by

John Alan Engbers

David Galvin, Director

Graduate Program in Mathematics

Notre Dame, Indiana

April 2013

SOME PROBLEMS INVOLVING H-COLORINGS OF GRAPHS

Abstract

by

John Alan Engbers

For graphs G and H, an H-*coloring* of G, or *homomorphism* from G to H, is an edge-preserving map from the vertices of G to the vertices of H. H-colorings generalize such graph theory notions as proper colorings and independent sets. In this dissertation, we consider four questions involving H-colorings of graphs.

Recently, Galvin [27] showed that the maximum number of independent sets in a an n vertex minimum degree δ graph occurs (for sufficiently large n) when $G = K_{\delta,n-\delta}$. First, we show this result holds for level sets: for all triples (n, δ, t) with $\delta \leq 3$ and $t \geq 3$, no n-vertex graph with minimum degree δ admits more independent sets of size t than $K_{\delta,n-\delta}$, and we obtain the same conclusion for $\delta > 3$ and $t \geq 2\delta + 1$.

Second, we begin the project of generalizing Galvin's result to arbitrary H. Writing $\hom(G, H)$ for the number of H-colorings of G, we show that for $\delta = 1$ and $\delta = 2$ and fixed H,

$$\hom(G, H) \leq \max\{\hom(K_{\delta+1}, H)^{\frac{n}{\delta+1}}, \hom(K_{\delta,\delta}, H)^{\frac{n}{2\delta}}, \hom(K_{\delta,n-\delta}, H)\}$$

for any n vertex minimum degree δ graph G (for sufficiently large n). For $\delta \geq 3$ (and sufficiently large n), we provide a class of H for which $\hom(G, H) \leq \hom(K_{\delta,n-\delta}, H)$ for any G in this family.

Third, for a given H, $k \in V(H)$, and regular bipartite G, we consider the proportion of vertices of G that get mapped to k in a uniformly chosen H-coloring of G. We find numbers $0 \leq a^-(k) \leq a^+(k) \leq 1$ with the property that for all such G, with high probability the proportion is between $a^-(k)$ and $a^+(k)$, and we give examples where these extremes are achieved.

Fourth, we study the set of H-colorings of the even discrete torus $\mathbb{Z}_m^d$. For any H and fixed m, we show that the space of H-colorings of $\mathbb{Z}_m^d$ may be partitioned into a subset of negligible size (as d grows) and a collection of subsets indexed by certain pairs $(A, B) \in V(H)^2$, with each H-coloring in the subset indexed by (A, B) having almost all vertices in one partition class mapped to A and almost all vertices in the other partition class mapped to B.

Dedication

To Ruth, Samuel, and Luke

CONTENTS

FIGURES

ACKNOWLEDGMENTS

No project of this magnitude comes to fruition without the help and support of many people. In particular, I'd like to thank

my advisor, David Galvin, for his encouragement, his patience, his guidance, his view of mathematics, and his instrumental role in inducing my passion for discrete mathematics;

my colleagues throughout the years, especially Jon Dent, Matt Dawson, Jon Gray, Justin Grieves, Chris Sass, Eric Savage, Harold Smith, and Martha Precup;

my mentors for previous projects, Gerard Venema and Bob Daverman, who showed me beautiful mathematics and always gave very helpful advice;

my thesis defense committee of Andrzej Dudek, Roxana Smarandache, and Zoltán Toroczkai;

my friends, including the wonderful people at South Bend CRC, with special thanks to my small group at SBCRC;

my parents, for constant encouragement, wisdom, support, and love;

and most importantly my family: Ruth, Samuel, and Luke — I owe you three more than you'll ever know.

CHAPTER 1

INTRODUCTION

1.1 Introduction

In this dissertation we consider a variety of problems centered on extremal and structural properties of H-colorings of graphs. We begin by describing these problems while attempting to defer many of the technicalities to later chapters.

For $G = (V(G), E(G))$ a simple, loopless graph, and $H = (V(H), E(H))$ a graph without multiple edges but perhaps with loops, an *H-coloring* of G, or *homomorphism* from G to H, is a function $f : V(G) \to V(H)$ that preserves adjacency, that is, which satisfies $f(u) \sim_H f(v)$ whenever $u \sim_G v$ (where $\sim_*$ denotes adjacency in $*$). See Figure 1.1. The term H-coloring is quite natural as we think of coloring the vertices of G, with the palette of available colors being the vertices of H and the edges of H telling us what pairs of colors are allowed to appear on adjacent vertices of G. We write $\mathrm{Hom}(G, H)$ for the set of all H-colorings of G, and let $\mathrm{hom}(G, H) = |\mathrm{Hom}(G, H)|$. (Unless explicitly stated otherwise, all graphs in this dissertation will be finite.)

H-colorings have a natural statistical physics interpretation as configurations in *hard-constraint spin systems*. Here, the vertices of G are thought of as sites that are occupied by particles, with edges of G representing pairs of bonded sites. The vertices of H are the different types of particles (or spins), and the occupation rule is that bonded sites must be occupied by pairs of particles that are adjacent in H. A legal configuration in such a spin model is exactly an H-coloring of G.

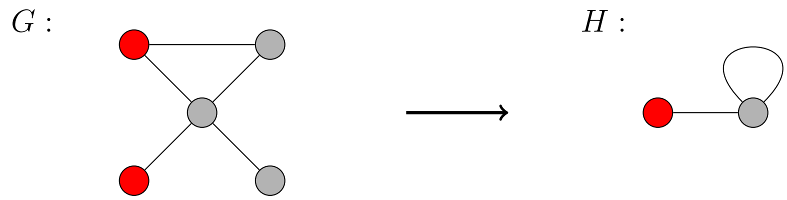

Figure 1.1. An example of an H-coloring of G.

When $H = K_q$ (where K_q is the complete loopless graph on q vertices; see the right-hand side of Figure 1.2), an element of $\mathrm{Hom}(G, K_q)$ is a coloring of the vertices of G from a palette of q colors so that adjacent vertices receive different colors, which is called a *proper q-coloring* of G. These have been widely studied throughout the history of graph theory; for example, the four-color theorem states that all planar graphs admit a proper 4-coloring. In statistical physics, a proper q-coloring is a configuration in the *zero-temperature q-state anti-ferromagnetic Potts model.* Other applications of vertex colorings are found in scheduling and assignment problems and solutions to Sudoku.

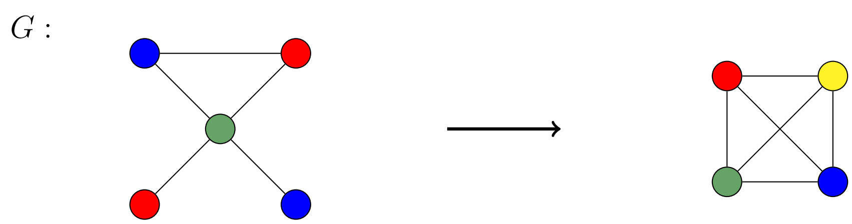

Figure 1.2. An H-coloring of G using $H = K_4$.

Another important example occurs when $H = H_{\rm WR}$ (where $H_{\rm WR}$ is the completely looped path on three vertices; see the right-hand side of Figure 1.3). The set $\mathrm{Hom}(G, H_{\rm WR})$ coincides with the state space of the *Widom-Rowlinson model* of statistical physics (or *WR model*), introduced in [58] as a model of liquid-vapor phase transitions. It models the placement of two repelling (but not self-repelling) particles on sites, where not every site needs a particle.

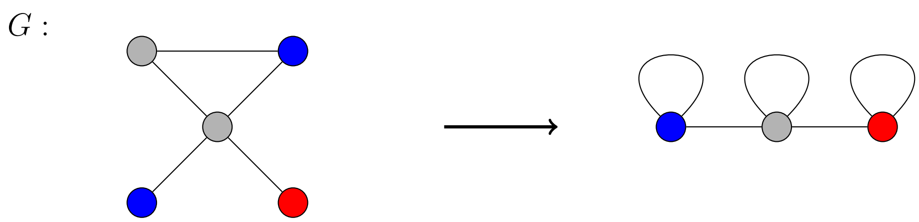

Figure 1.3. An H-coloring of G using the Widom-Rowlinson graph $H = H_{\rm WR}$.

A third example of an H-coloring occurs when $H = H_{\rm ind}$ (where $H_{\rm ind}$ consists of two vertices joined by an edge, with a loop at one of the vertices; see the right-hand side of Figure 1.4). An element of $\mathrm{Hom}(G, H_{\rm ind})$ yields a set of vertices of G which spans no edges (via the preimage of the unlooped vertex), which is called an *independent set* (or *stable set*) of G. In statistical physics, an independent set is a configuration in the *hard-core lattice gas model*, a model of the occupation of space by large particles.

See for example [10] for a discussion of some of these models from a combinatorial point of view, and [62] for a statistical physics oriented discussion.

G :

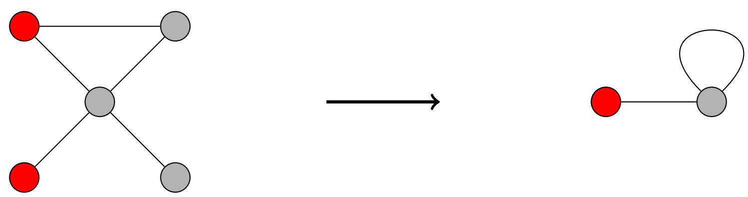

Figure 1.4. An H-coloring of G using $H = H_{\text{ind}}$.

Let $i(G) = \hom(G, H_{\text{ind}})$ be the total number of independent sets in G. Granville, motivated by a question in combinatorial group theory, asked which graph in the family of n-vertex d-regular graphs has the largest value of $i(G)$ (see [2] for more details of the combinatorial group theory background). An approximate answer — $i(G) \le 2^{n/2(1+o(1))}$ for all such G, where $o(1) \to 0$ as $d \to \infty$ — was given by Alon in [2], and he speculated a more exact result, that the maximizing graph, at least in the case $2d|n$, is the disjoint union of $n/2d$ copies of $K_{d,d}$. See Figure 1.5. This speculation was confirmed for *bipartite* G by Kahn [39] (and recently for general regular G by Zhao [67]).

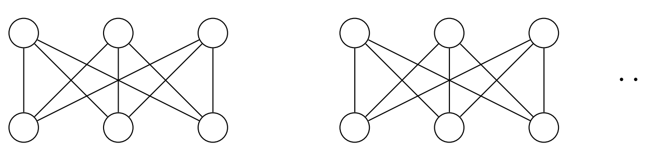

Figure 1.5. The disjoint union of copies of $K_{d,d}$.

Furthermore, Kahn showed in [39, 40] that the occupation probabilities induced by a randomly chosen independent set are suitably close to the distribution which picks a partition class and chooses an independent set at random from within that partition class. This is a somewhat remarkable and non-obvious phenomenon which gives a fairly complete description of most independent sets in any n-vertex d-regular graph.

This is the point of departure for much of the dissertation: we consider related extremal questions for both independent sets of a fixed size in graphs and also for H-colorings of graphs, and we consider related structural questions for H-colorings of graphs.

1.2 Extremal questions

An extremal question in graph theory has the following general form: given a parameter f of graphs (a function f that assigns to each G a real number $f(G)$) and a family $\mathcal{G}$ of graphs, what are the extremal values of $f(G)$ as G ranges over $\mathcal{G}$, and which graphs achieve the extremes? Granville's extremal question for independent sets is the case where $f(G) = i(G)$ and $\mathcal{G}$ is the family of n-vertex d-regular graphs (and where he only looks to maximize $i(G)$). Extremal questions can be modified by changing either the parameter or the family.

For instance, the extremal question for maximizing the number of independent sets can be modified by changing the family of graphs under consideration. A few of the other families that have been considered (in addition to n-vertex d-regular graphs described above) include n-vertex trees [56], n-vertex m-edge graphs [15], n-vertex graphs with fixed independence number [69], n-vertex graphs with a given number of cut-edges [36], and n-vertex claw-free graphs [54].

For a fixed $t \in \{0, 1, \ldots, |V(G)|\}$, we let $i_t(G)$ denote the number of independent sets of size t in G. The values $i_t(G)$ are the coefficients of the independence polynomial of a graph (the polynomial $P_G(x) = \sum_t i_t(G)x^t$). A related extremal question to ask is: for each t, what is the largest value attained by $i_t(G)$ as G ranges over a family of graphs? This question has also been studied for various families of graphs, including n-vertex d-regular graphs (some asymptotic results in [12]), n-vertex m-edge graphs [14], and n-vertex trees [65].

Notice that any n-vertex tree has minimum degree 1, and so another natural family to study is the family of n-vertex graphs with fixed minimum degree δ. Denote this family by $\mathcal{G}(n, \delta)$. Our extremal questions become: which graphs G in $\mathcal{G}(n, \delta)$ maximize $i(G)$ and $i_t(G)$ for each t? Sapozhenko first studied the question for $i(G)$ in bipartite graphs in $\mathcal{G}(n, \delta)$ for large δ in [60].

Since a set of vertices is independent if no edge is present among those vertices, it is natural to conjecture that the extremal graph would have the least number of edges possible (so it would be δ-regular, or close to δ-regular). This is surprisingly not the case. In [27], Galvin showed that for $n \geq 8\delta^2$ (and conjectured for $n \geq 2\delta$) that the unique graph in $\mathcal{G}(n, \delta)$ which admits the largest number of independent sets is $K_{\delta, n-\delta}$, the complete bipartite graph with δ vertices in one partition class and $n - \delta$ vertices in the other partition class. See Figure 1.6. In that paper, he conjectured that the same graph also maximizes the number of independent sets of each nontrivial size $t \geq 3$, and showed this for $\delta = 1$ and $n \geq 2$. The case $t = 2$ is an anomoly as it simply counts the number of non-edges in the graph G, and so in this case the maximizer is indeed the graph with the least number of edges. (Also note that the cases $t = 0$ and $t = 1$ have values 1 and n, respectively, for any $G \in \mathcal{G}(n, \delta)$.)

Chapter 3 presents joint work with D. Galvin, in which we prove this conjecture for a large range of n and t.

Theorem 1.2.1. *1. Fix $\delta \in \{2, 3\}$, $n \geq 2\delta$, and $t \geq 3$. Then for any $G \in \mathcal{G}(n, \delta)$,*

$$i_t(G) \leq i_t(K_{\delta, n-\delta}) = \binom{n-\delta}{t} + \binom{\delta}{t}.$$

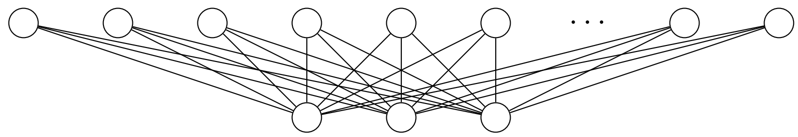

Figure 1.6. The complete bipartite graph $K_{3,n-3}$.

2. Fix $\delta \geq 4$, $n \geq 3\delta + 1$, and $t \geq 2\delta + 1$. Then for any $G \in \mathcal{G}(n, \delta)$,

$$i_t(G) \leq i_t(K_{\delta,n-\delta}) = \binom{n-\delta}{t}.$$

The precise statement of what we prove is the content of Theorem 3.1.4, where we also characterize the cases of equality. A key ingredient in the proof of Theorem 3.1.4 is our study of critical graphs (graphs where the deletion of any edge or vertex lowers the minimum degree), which are an important subfamily of the family of graphs with a fixed minimum degree. These results appear in [22], which has been submitted for publication.

Realizing that we can view independent sets as a particular instance of an H-coloring, the following question is quite natural: given a family of graphs $\mathcal{G}$, which G in $\mathcal{G}$ maximize $\hom(G, H)$ for each H? For the family of n-vertex m-edge graphs, this question was posed for proper q-colorings ($H = K_q$) in the 1980's by Linial [48] and Wilf [64]. The answer has not fully been resolved; for recent progress see e.g. [49] and the references therein. Results for some other choices of H appear in [15, 16].

For n-vertex d-regular bipartite graphs, recall that Kahn [39] showed that $i(G) \leq i(K_{d,d})^{n/2d}$ for any G in this family. Using the identification of an independent set with an H_{ind}-coloring of G, we can write this as $\hom(G, H_{\text{ind}}) \leq \hom(K_{d,d}, H_{\text{ind}})^{n/2d}$. Galvin and Tetali [32] generalized his entropy techniques and showed that for any H and G in this family, $\hom(G, H) \leq \hom(K_{d,d}, H)^{\frac{n}{2d}}$. (Notice that when $2d|n$ the bound is achieved by $\frac{n}{2d}K_{d,d}$, the disjoint union of $n/2d$ copies of $K_{d,d}$.) The fact that this holds for every H is quite striking.

For the family of n-vertex d-regular (not necessarily bipartite) graphs, Kahn's bound $\hom(G, H_{\text{ind}}) \leq \hom(K_{d,d}, H_{\text{ind}})^{n/2d}$ continues to hold for all G in this larger family (as demonstrated by Zhao [67]). It is tempting to believe that Galvin and Tetali's bound for arbitrary H extends to this larger family as well. There is indeed a large class of H for which $\hom(G, H) \leq \hom(K_{d,d}, H)^{\frac{n}{2d}}$ holds among all n-vertex d-regular G (see [28, 68]; evidence for proper q-colorings is given in [29]). However, Galvin [28] has demonstrated triples (n, d, H) for which $\hom(G, H) \leq \hom(K_{d+1}, H)^{\frac{n}{d+1}}$ for all G in this family.

For $\mathcal{G}(n, \delta)$, the family of n-vertex graphs with minimum degree δ, the extremal graph for independent sets in this family (with n large) is $G = K_{\delta,n-\delta}$ [27]. In Chapter 4, we extend this result to H-colorings for a large class of H, and for small fixed δ we show that all other H are maximized by either $G = \frac{n}{2\delta}K_{\delta,\delta}$ or $G = \frac{n}{\delta+1}K_{\delta+1}$ (when $2\delta(\delta+1)|n$).

For the statement of this theorem, we assume that loops count once toward the degree of a vertex, and we let Δ_H denote the maximum degree of H.

Theorem 1.2.2. *1. Fix H and $\delta \in \{1, 2\}$. There exists a c (depending on H and δ) such that for $n > c$ and $G \in \mathcal{G}(n, \delta)$, we have*

$$\hom(G, H) \leq \max\{\hom(K_{\delta,\delta}, H)^{\frac{n}{2\delta}}, \hom(K_{\delta+1}, H)^{\frac{n}{\delta+1}}, \hom(K_{\delta,n-\delta}, H)\}.$$

2. Fix H and δ. If H satisfies $\sum_{v\in V(H)} d(v) < (\Delta_H)^2$, then there exists a c (depending on H) such that for $n > c^\delta$ and $G \in \mathcal{G}(n, \delta)$, we have

$$\hom(G, H) \leq \hom(K_{\delta,n-\delta}, H).$$

The precise statement of what we prove is the content of Theorems 4.1.5, 4.1.6, and 4.1.7, where we also characterize the graphs that achieve equality. The proofs of the $\delta = 1$ and $\delta = 2$ results involve analyzing edge-critical graphs (graphs with the property that the deletion of any edge lowers the minimum degree). The proof of the general δ result partitions $\mathcal{G}(n, \delta)$ based on the size of a maximum matching in the graph, shows that a large matching cannot be present in an extremal graph, and analyzes those graphs with a small maximum matching size. These results appear in [19], and are being prepared for publication.

1.3 Structural questions

Chapters 3 and 4 deal with extremal questions for H-colorings of graphs; Chapters 5 and 6 address the rather different question of the *typical* appearance of an H-coloring of a graph G. At this point, we'll motivate the results of Chapters 5 and 6 using the language of proper q-colorings; the theorems and proofs in those chapters will be suitably generalized to H-colorings.

Recall that a proper q-coloring of a graph G is a coloring of the vertices of G from a palette of q colors so that adjacent vertices receive different colors. Suppose that we fix a regular bipartite graph G on n vertices, and fix a particular equipartition $U \cup W$. (We will assume this equipartition is fixed for any regular bipartite G under consideration; if G is connected, then U and W are essentially unique.) A result from Galvin and Tetali [32], extending Kahn's beautiful entropy techniques for independent sets [39], gives nearly matching upper and lower bounds on the total number of proper q-colorings of G. However, their result gives little insight into the typical structure of these colorings, and it is natural to ask what a uniformly chosen proper q-coloring of G typically looks like.

In [39], Kahn answers this for independent sets (in fact, for *weighted* independent sets). He showed that independent sets exhibit *phase coexistence* in the sense that most independent sets in G tend to come either mostly from U or mostly from W. In particular, he showed that the size of an independent set is close to $n/4$, which is the expected size for an independent set chosen according to the distribution which half the time picks an independent set exclusively from U and half the time picks exclusively from W.

The analogous distribution for a proper q-coloring of G is the following: color G by using some fixed $\lfloor \frac{q}{2} \rfloor$ colors independently on the vertices of the partition class U (or W) and using the remaining $\lceil \frac{q}{2} \rceil$ colors independently on the vertices of the partition class W (U). We call a coloring obtained from this distrubution a *pure proper q-coloring*. It seems reasonable to think that most proper q-colorings of G should be, in some sense, close to some pure proper q-coloring.

Chapter 5 represents joint work with D. Galvin where we show that this is true in the sense that the typical proper q-coloring of G has each color appearing on approximately the same proportion of vertices as would be expected from a pure proper q-coloring of G.

Theorem 1.3.1. *Given an n-vertex, d-regular bipartite graph G, almost all proper q-colorings have each color appearing on a proportion between $\frac{1}{q+1}$ and $\frac{1}{q-1}$ of the vertices (for q odd) or on a proportion close to $\frac{1}{q}$ of the vertices (for q even).*

So for even q, almost all proper q-colorings of a regular bipartite graph are "almost equitable". Notice that by the symmetry of the colors the expected number of vertices receiving a particular color is n/q.

The precise statement of what we prove, which is suitably generalized to weighted H-colorings, is the content of Theorem 5.1.2 (the corollary addressing proper q-colorings is Corollary 5.1.5). In this more general result, some H are similar to the q odd case of Theorem 1.3.1, where the proportion of vertices receiving a particular color lies in some interval. In these cases, the random regular bipartite graph shows that these intervals are sharp (see Theorem 5.1.4). However, most H are similar to the q even case of Theorem 1.3.1, where the proportion of vertices receiving a particular color is concentrated around a single value. We also remark that the condition of regularity can be relaxed somewhat (see Theorem 5.4.1), and in Corollary 5.4.3 we apply this to show that if edge percolation is run on G with probability p, then for many H there is a threshold occurring at $p = 1/d$. These results appear in the *Journal of Combinatorial Theory, Series B* [20].

For certain graphs and their set of proper q-colorings we have results about the actual distributions of the colors among the two partition classes and not just their occurrence probabilities. This is true for the random regular bipartite graph, as is shown in the proof of Theorem 5.1.4 by critically using its excellent expansion properties.

In Chapter 6, which is joint work with D. Galvin, we refine the entropy techniques of Kahn [40] to obtain results for the d-dimensional discrete hypercube Q_d (the graph on vertex set $\{0,1\}^d$ with edges joining strings that differ in one coordinate; see Figure 1.7), which has much weaker expansion than the random regular bipartite graph. Q_d is an widely studied graph which often acts as a test-bed for results on the harder to study infinite regular lattice $\mathbb{Z}^d$. Related results have been obtained on Q_d for proper 3-colorings [23, 55] and independent sets [43].

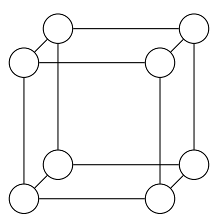

Figure 1.7. The 3-dimensional discrete hypercube Q_3.

Theorem 1.3.2. *All but a vanishing proportion of proper q-colorings of Q_d have $\lfloor \frac{q}{2} \rfloor$ colors each appearing on a proportion $\left(1 \pm 2^{-\Omega(d)}\right) \frac{1}{\lfloor q/2 \rfloor}$ of the vertices in one partition* *class, and the remaining* $\lceil \frac{q}{2} \rceil$ *colors each appearing on a proportion* $\left(1 \pm 2^{-\Omega(d)}\right) \frac{1}{\lceil q/2 \rceil}$ *of the vertices in the other partition class.*

In other words, almost all proper q-colorings of Q_d differ from some pure proper q-coloring of Q_d on only some small number of vertices. We are actually able to prove Theorem 1.3.2 in a far more general setting, that of weighted H-colorings of the even discrete torus $\{0, 1, \ldots, m-1\}^d$ (for $m \geq 2$ even). The precise statement of what we prove is the content of Theorems 6.1.1 and 6.1.2.

One application of Theorem 1.3.2 is the following surprising long-range influence result.

Theorem 1.3.3. *Let* $v \neq w$ *be vertices in* Q_d*. Choose a uniform proper* q*-coloring of* Q_d *conditioned on the information that* w *receives color* 1*. Then the asymptotic probability that* v *receives color* 1 *is either* $\frac{2}{q}$ *(if* v *and* w *are in the same partition class) or* 0 *(if* v *and* w *are not in the same partition class).*

Note that without conditioning on the color of w, the probability that v receives color 1 is exactly $1/q$ (by symmetry).

Theorem 1.3.3 generalizes to weighted H-colorings of the discrete even torus as well; the precise statement of what we prove is the content of Theorem 6.2.1 (the corollary specifically addressing proper q-colorings is Corollary 6.2.4). These results appear in the *Journal of Combinatorial Theory, Series B* [21].

Chapter 2 presents much of the notation that we will use in this dissertation. Chapters 3, 4, 5, and 6 (containing the substantial results) can be read essentially independently of each other; they all use the notation developed in Section 2.1. Chapter 3 requires the additional notation and results from Section 2.2, Chapter 4 requires the additional notation and results from Sections 2.2 and 2.3, and each of Chapters 5 and 6 require the additional notation from Section 2.3.

CHAPTER 2

NOTATION, CONVENTIONS, AND PRELIMINARY MATERIAL

2.1 Basic definitions, notation, and conventions

In this section, we gather together some of the basic definitions and notation that we will use. Any additional notation will be defined as it is needed. For graph theory basics, see e.g. [5], [17].

Let $G = (V(G), E(G))$ be a finite, undirected graph with no multi-edges. Any graph named G (or some derivative of G) will always be loopless; any graph named H (or some derivative of H) may have loops. Throughout we will assume $|V(G)| = n$. For v, w in $V(G)$, we write $v \sim w$ (or $v \sim_G w$ to emphasize the graph G) if there is an edge from v to w, and $v \nsim w$ if there is not. Set $N(v) = \{x : x \sim v\}$ ($N(v)$ is the *neighborhood* of v). For $A, B \subset V(G)$ we use $N(A)$ for $\cup_{v \in A} N(v)$, and write $A \sim B$ if $a \sim b$ for all $a \in A$ and $b \in B$.

An *independent set* $I \subseteq V(G)$ is a set of vertices that spans no edges, and the *size* of an independent set is $|I|$. The number of independent sets of size t in G is denoted $i_t(G)$, and the total number of independent sets in G is denoted $i(G)$ (so $i(G) = \sum_t i_t(G)$).

A graph is *bipartite* if its vertex set can be partitioned as $\mathcal{E} \cup \mathcal{O}$ so that $\mathcal{E}$ and $\mathcal{O}$ are independent sets.

For a vertex $v \in V(G)$, we denote the *degree* of v by $d(v)$ (or $d_G(v)$ to emphasize the graph G). For degrees in graphs H (where loops are allowed), we utilize the convention that loops count *once* toward the degree (i.e. $d(v) = |N(v)|$). The minimum and maximum degrees of graphs G are denoted $\delta(G)$ and $\Delta(G)$, respectively. A graph is *regular* if $\delta(G) = \Delta(G)$ and *d-regular* if $\delta(G) = \Delta(G) = d$.

We let K_n denote the *complete graph* on n vertices, $K_{a,b}$ the *complete bipartite graph* with a vertices in one partition class and b vertices in the other partition class (with the *star* on n vertices being the special case $K_{1,n-1}$), P_n the *path* on n vertices, and C_n the *cycle* on n vertices. For $m \in \mathbb{N}$ we let mG denote the graph consisting of m disjoint copies of the graph G.

We use $G[Y]$ to denote the subgraph induced by a subset Y of the vertices, and $E(Y)$ to denote the edge set of this subgraph. When appropriate we abuse notation by failing to distinguish between a graph and the set of vertices of a graph; for example we will write $d_Y(v)$ instead of $d_{G[Y]}(v)$.

For $t \in \mathbb{N}$ and $x \in \mathbb{R}$, we let $x^{\underline{t}}$ indicate the falling power $x(x-1)\cdots(x-(t-1))$, and we let $\binom{x}{t} = \frac{x^{\underline{t}}}{t!}$.

We use the standard Bachmann-Landau notation, with $f = o(g)$ and $f = \omega(g)$ indicating, respectively, that $f/g \to 0$ and $f/g \to \infty$ as some variable (often n) approaches infinity; $f = O(g)$ and $f = \Omega(g)$ indicating, respectively, that $|f| < C|g|$ and $|f| > C|g|$ for some constant C; and $f = \Theta(g)$ indicating that both $f = O(g)$ and $f = \Omega(g)$ hold.

2.2 Definitions, notation, and preliminary material for minimum degree δ graphs

In Chapters 3 and 4, we will be dealing with n-vertex graphs with minimum degree δ. Here we provide some definitions and structural results that will be used.

For integers $\delta \geq 1$ and $n \geq 1$, we let $\mathcal{G}(n, \delta)$ denote the set of all graphs on n vertices with minimum degree δ. For a graph $G \in \mathcal{G}(n, \delta)$, we let $V_{=\delta}$ $(= V_{=\delta}(G))$ denote the set of vertices with degree δ, and $V_{>\delta}$ $(= V_{>\delta}(G))$ denote the set of vertices of G with degree larger than δ (so $V_{=\delta}$ and $V_{>\delta}$ partition $V(G)$).

A graph G with minimum degree δ is *edge-critical (for δ)* if for any edge e in G, the minimum degree of $G - e$ is $\delta - 1$. It is *vertex-critical (for δ)* if for any vertex v in G, the minimum degree of $G - v$ is $\delta - 1$. If it is both edge- and vertex-critical, we say that G is *critical (for δ)*.

Edge-critical graphs in $\mathcal{G}(n, 1)$ are disjoint unions of stars, and the only critical graph in $\mathcal{G}(n, 1)$ is the matching graph $\frac{n}{2}K_2$.

We have an inductive decomposition of all edge-critical graphs in $\mathcal{G}(n, 2)$.

Lemma 2.2.1. *Fix $\delta = 2$. Let G be a n-vertex edge-critical graph. Either*

1. *G is a disjoint union of cycles or*
2. *$V(G)$ can be partitioned as $Y_1 \cup Y_2$ with $1 \leq |Y_1| \leq n - 3$ in such a way that Y_1 induces a path, Y_2 induces a graph with minimum degree 2, each endvertex of the path induced by Y_1 has exactly one edge to Y_2, the endpoints of these two edges in Y_2 are either the same or non-adjacent, and there are no other edges from Y_1 to Y_2 (see Figure 2.1).*

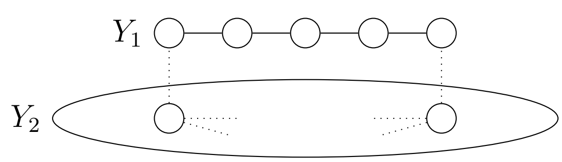

Figure 2.1. A example of a path of length 5 ($|Y_1| = 5$) given in Lemma 2.2.1.

Proof. If G is not a disjoint union of cycles, then it has some vertices of degree greater than 2. If some component contains exactly one such vertex, say v, then by parity considerations $d(v)$ is even and at least 4. Since all degrees are even, the edge set in that component may be partitioned into cycles. Take any cycle through v and remove v from it to get a path whose vertex set can be taken to be Y_1.

Suppose now that some component of G has at least two vertices with degree larger than 2. Since G is edge-critical, $V_{>\delta}$ forms an independent set and so there is a path on at least 3 vertices joining distinct vertices $v_1, v_2 \in V_{>\delta}$, all of whose internal vertices $u_1, \dots, u_k$ have degree 2 (the shortest path joining two vertices in $V_{>\delta}$ would work). We may now take $Y_1 = \{u_1, \dots, u_k\}$. Note that the Y_2 endpoints (v_1 and v_2) of the edges from u_1 and u_k to Y_2 are both in $V_{>\delta}$ and so are non-adjacent. □

Remark. If we restrict to the subclass of critical graphs in $\mathcal{G}(n, 2)$, then we obtain the same conclusion with $2 \leq |Y_1| \leq n-3$ (instead of $1 \leq |Y_1| \leq n-3$). Indeed, $|Y_1| = 1$ implies that deleting the single vertex in Y_1 leaves a graph with minimum degree 2, which contradicts the assumption that G is vertex-critical.

Corollary 2.2.2. *Fix $\delta = 2$. Let G be a n-vertex edge-critical graph. Then G may be constructed via the following iterative procedure:*

- *Start with a non-empty collection of disjoint cycles;*
- *Next, iteratively add a collection paths of length $k \geq 2$ which connect to existing vertices of the graph* only *at the endpoints of the path;*
- *Finally, add a collection of paths of length* 1 *to the graph (all at the same time).*

Proof. By Lemma 2.2.1, paths can be removed inductively until a collection of disjoint cycles remain, and so we may construct any graph G starting with the cycles. Reversing this, we may iteratively add paths of length $k \geq 1$ to produce G. The content of this corollary is that G may be constructed by adding all paths of length $k \geq 2$ before the paths of length 1, and the paths of length 1 may all be added at the same time.

Why is this possible? Adding a path of length 1 creates a vertex of degree 2 adjacent to two vertices of degree at least 3. Since G is edge-critical, no future path will connect to this vertex of degree 2. □

2.3 Definitions and notation for H-colorings

Chapters 4, 5, and 6 deal with (weighted) H-colorings of certain graphs and classes of graphs. We provide some notation here.

Recall that for a graph G and a graph (possibly with loops) H an H*-coloring* of G (or *graph homomorphism*) is an edge preserving map from $V(G)$ to $V(H)$. We emphasize that any definition pertaining to graphs G from Section 2.1 will continue to hold for graphs H, with the exception that we assume H may have loops and we assume $|V(H)| = q$ (so, for example, for $A, B \subset V(H)$ we let $A \sim B$ indicate that $a \sim b$ for all $a \in A$ and $b \in B$). We let $\mathrm{Hom}(G, H)$ denote the set of all H-colorings of G and $\hom(G, H) = |\mathrm{Hom}(G, H)|$. We frequently refer to elements of $V(H)$ as *colors*, and say that a vertex of G is *colored* k if its image in the H-coloring under consideration is k.

The graph H_{WR} will denote the fully looped path on 3 vertices (see the right-hand side of Figure 1.3) and the graph H_{ind} will denote the graph consisting of an edge and one looped endvertex (see the right-hand side of Figure 1.4).

From a statistical physics standpoint, there is a very natural family of probability distributions that can be put on $\mathrm{Hom}(G, H)$. Fix a set of positive weights $\Lambda = \{\lambda_i : i \in V(H)\}$ indexed by the vertices of H. We think of the magnitude of λ_k as measuring how likely particle k is to appear at each site. This can be formalized by giving each $f \in \mathrm{Hom}(G, H)$ weight $w_\Lambda(f) = \prod_{v \in V(G)} \lambda_{f(v)}$ and probability

$$p_\Lambda(f) = \frac{w_\Lambda(f)}{Z_\Lambda(G, H)}$$

where $Z_\Lambda(G, H) = \sum_{f \in \mathrm{Hom}(G,H)} w_\Lambda(f)$ is the appropriate normalizing constant or *partition function* of the model. When all weights are 1, $Z_\Lambda(G, H) = \hom(G, H)$ and p_Λ is uniform measure. In this special case, we will often omit the reference to Λ (e.g. we will write $p(f)$ for $p_\Lambda(f)$). Interestingly, several proofs of results about *unweighted* H-colorings (including those in Chapter 5) *require* passing to the weighted model first; other examples of this phenomenon may be found in e.g. [39, 41].

For $S \subseteq \mathrm{Hom}(G,H)$ and $T \subseteq V(H)$ we write $w_\Lambda(S)$ for $\sum_{f\in S} w_\Lambda(f)$ and λ_T for $\sum_{k\in T}\lambda_k$ (so if $\lambda_i = 1$ for all i, then $w_\Lambda(S) = |S|$ and $\lambda_T = |T|$). Set

$$\eta_\Lambda(H) = \max\left\{\lambda_A\lambda_B \colon A, B \subseteq V(H),\ A \sim B\right\}.$$

When $\lambda_i = 1$ for all i, $\eta(H)$ is (essentially) measuring the size of the largest bipartite subgraph of H, where size is measured by the number of edges; $\eta_\Lambda(H)$ is then measuring the size of the largest weighted bipartite subgraph of H. Set

$$\mathcal{M}_\Lambda(H) = \left\{(A,B) \in V(H)^2 \colon A \sim B,\ \lambda_A\lambda_B = \eta_\Lambda(H)\right\};$$

in other words, $\mathcal{M}_\Lambda(H)$ is counting the number of different realizations of the largest weighted bipartite subgraph in H.

If G is a regular bipartite graph with fixed bipartition $\mathcal{E}\cup\mathcal{O}$, then given $A, B \subset V(H)$ with $A \sim B$, a *pure-(A,B) coloring* is an $f \in \mathrm{Hom}(G,H)$ with $f(u) \in A$ for all $u \in \mathcal{E}$ and $f(v) \in B$ for all $v \in \mathcal{O}$. Notice that if $|\mathcal{E}| = |\mathcal{O}| = k$, then there are $\eta_\Lambda(H)^k$ pure-(A,B) colorings of G.

In Chapter 6, we'll focus on H-colorings of the even discrete torus, which is the graph on vertex set $V = \{0,1,\ldots,m-1\}^d$ (for m even and $d \geq 1$) with edge set E consisting of all pairs of strings that differ by exactly 1 (mod m) on exactly one coordinate. For $m \geq 4$ it is $2d$-regular and bipartite. In the case $m = 2$, the even discrete torus is d-regular and bipartite, and it is isomorphic to the familiar *Hamming cube* or *discrete hypercube* (the graph on vertex set $\{0,1\}^d$ with edge set consisting of all pairs of strings that differ on exactly one coordinate). For this special case we use the more familiar notation Q_d.

Section 6.4.2 contains additional technical notation that we use in Chapter 6.

CHAPTER 3

INDEPENDENT SETS OF A FIXED SIZE

3.1 Introduction and statement of results

Recall that an independent set in a graph G is a set of vertices spanning no edges, and that $i(G)$ denotes the number of independent sets in G. In [56] this quantity is referred to as the *Fibonacci number* of G, motivated by the fact that for the path graph P_n its value is a Fibonacci number. It has also been studied in the field of molecular chemistry, where it is referred to as the *Merrifield-Simmons* index of G [52].

A natural extremal enumerative question to ask is the following: as G ranges over some family $\mathcal{G}$, what is the maximum value of $i(G)$, and which graphs G achieve this maximum?

This question has been addressed for numerous families. Prodinger and Tichy [56] considered the family of n-vertex trees, and showed that the maximum is uniquely attained by the star $K_{1,n-1}$. Kahn [39] considered the family of n-vertex d-regular bipartite graphs and showed that when $2d|n$ the maximizing graph is $\frac{n}{2d}K_{d,d}$, the disjoint union of $n/2d$ copies of $K_{d,d}$; Zhao [67] extended Kahn's result to the family of n-vertex d-regular graphs. The family of n-vertex, m-edge graphs was considered by Cutler and Radcliffe in [15], and they observed that it is a corollary of the Kruskal-Katona theorem [42, 44] that the lex graph $L(n,m)$ (on vertex set $\{1,\dots,n\}$, with edges being the first m pairs in lexicographic order) maximizes $i(G)$ in this class. Zykov [69] considered the family of graphs with a fixed number of vertices and fixed independence number, and showed that the maximum is attained by the complement of a certain Turán graph. (Zykov was actually considering cliques in a graph with given clique number, but by complementation this is equivalent to considering independent sets in a graph with given independence number.) Other papers addressing questions of this kind include [36, 47, 54, 60].

Having resolved the question of maximizing $i(G)$ for G in a particular family, it is natural to ask which graph maximizes $i_t(G)$, the number of independent sets of size t in G, for each possible t. For many families, it turns out that the graph which maximizes $i(G)$ also maximizes $i_t(G)$ for all t. Wingard [65] showed this for trees, Zykov [69] showed this for graphs with a given independence number (see [14] for a short proof), and Cutler and Radcliffe [14] showed this for graphs on a fixed number of edges (again, as a corollary of Kruskal-Katona). In [39], Kahn conjectured that for all $2d|n$ and all t, no n-vertex, d-regular graph admits more independent sets of size t than the disjoint union of $n/2d$ copies of $K_{d,d}$; this conjecture remains open, although asymptotic evidence appears in [12].

In this chapter, we consider the family $\mathcal{G}(n, \delta)$ of n-vertex graphs with minimum degree δ, and we look at maximizing $i(G)$ and $i_t(G)$ for graphs in $\mathcal{G}(n, \delta)$. First, we consider the extremal problem for $i(G)$. Intuitively, one might imagine that since removing edges increases the count of independent sets, the graph in $\mathcal{G}(n, \delta)$ that maximizes the count of independent sets would be δ-regular (or close to δ-regular), but this turns out not to be the case. The following result is due to Galvin [27].

Theorem 3.1.1 (Galvin, 2011 [27]). *For $n \geq 2$ and $G \in \mathcal{G}(n, 1)$, we have $i(G) \leq i(K_{1,n-1})$. For $\delta \geq 2$, $n \geq 8\delta^2$ and $G \in \mathcal{G}(n, \delta)$, we have $i(G) \leq i(K_{\delta,n-\delta})$.*

What about maximizing $i_t(G)$ for each t? Unlike the family of n-vertex trees [65], n-vertex m-edge graphs [14], and n-vertex d-regular graphs (conjectured in [39], with asymptotic evidence in [12]), the family $\mathcal{G}(n, \delta)$ is an example of a family for which the maximizer of the total count is *not* the maximizer for each individual t. Indeed, consider the case $t = 2$. Maximizing the number of independent sets of size two is the same as minimizing the number of edges, and it is easy to see that for all fixed δ and sufficiently large n, there are n-vertex graphs with minimum degree at least δ which have fewer edges than $K_{\delta,n-\delta}$ (consider for example a δ-regular graph, or one which has one vertex of degree $\delta + 1$ and the rest of degree δ). However, we expect that anomalies like this occur for very few values of t. Indeed, the following conjecture is made in [27].

Conjecture 3.1.2 (Galvin, 2011 [27])**.** *For each $\delta \geq 1$ there is a $C(\delta)$ such that for all $t \geq C(\delta)$, $n \geq 2\delta$ and $G \in \mathcal{G}(n, \delta)$, we have*

$$i_t(G) \leq i_t(K_{\delta,n-\delta}) = \binom{n-\delta}{t} + \binom{\delta}{t}.$$

The case $\delta = 1$ of Conjecture 3.1.2 is proved in [27], with $C(1)$ as small as it possible can be, namely $C(1) = 3$. In [1], Alexander, Cutler and Mink looked at the subfamily $\mathcal{G}^{\text{bip}}(n, \delta)$ of *bipartite* graphs in $\mathcal{G}(n, \delta)$, and resolved the conjecture in the strongest possible way for this family.

Theorem 3.1.3 (Alexander, Cutler, Mink 2012 [1])**.** *For $\delta \geq 1$, $n \geq 2\delta$, $t \geq 3$ and $G \in \mathcal{G}^{\text{bip}}(n, \delta)$, we have $i_t(G) \leq i_t(K_{\delta,n-\delta})$.*

This provides good evidence for the truth of the strongest possible form of Conjecture 3.1.2, namely that we may take $C(\delta) = 3$.

We make significant progress towards this strongest possible form of Conjecture 3.1.2. We completely resolve the cases $\delta = 2$ and $\delta = 3$, and for larger δ we deal with all but a small fraction of cases.

Theorem 3.1.4. *1. For $\delta = 2$, $t \geq 3$ and $G \in \mathcal{G}(n, 2)$, we have $i_t(G) \leq i_t(K_{2,n-2})$. For $n \geq 5$ and $3 \leq t \leq n-2$ we have equality iff $G = K_{2,n-2}$ or G is obtained from $K_{2,n-2}$ by joining the two vertices in the partite set of size 2.*

2. *For $\delta = 3$, $t \geq 3$ and $G \in \mathcal{G}(n,3)$, we have $i_t(G) \leq i_t(K_{3,n-3})$. For $n \geq 6$ and $t = 3$ we have equality iff $G = K_{3,n-3}$; for $n \geq 7$ and $4 \leq t \leq n-3$ we have equality iff G is obtained from $K_{3,n-3}$ by adding some edges inside the partite set of size 3.*

3. *For $\delta \geq 3$, $t \geq 2\delta+1$ and $G \in \mathcal{G}(n,\delta)$, we have $i_t(G) \leq i_t(K_{\delta,n-\delta})$. For $n \geq 3\delta+1$ and $2\delta+1 \leq t \leq n-\delta$ we have equality iff G is obtained from $K_{\delta,n-\delta}$ by adding some edges inside the partite set of size δ.*

(Note that there is some overlap between parts 2 and 3 above.) Recently, Law and McDiarmid [45] have found a proof of Conjecture 3.1.2 that holds for all n sufficiently large, $\delta = o(n^{1/3})$, and $t \geq 3$; Theorem 3.1.4 part 3 holds for a larger range of n but a smaller range of t.

In Section 3.2 we make some easy preliminary observations that will be helpful in the proof of Theorem 3.1.4. We will prove the case $\delta = 2$ (part 1 of Theorem 3.1.4) in Section 3.3. We begin Section 3.4 with the proof of part 3 of Theorem 3.1.4, and then explain how the argument can be improved (within the class of critical graphs). This improvement is an important ingredient in the proof of the case $\delta = 3$ (part 2 of Theorem 3.1.4) whose proof we present in Section 3.5.

We also note that part 1 of Theorem 3.1.4 provides an alternate proof of the $\delta = 2$ case of the total count of independent sets, originally proved in [27].

Corollary 3.1.5. *For $n \geq 4$ and $G \in \mathcal{G}(n,2)$, we have $i(G) \leq i(K_{2,n-2})$. For $n = 4$ and $n \geq 6$ there is equality iff $G = K_{2,n-2}$.*

Proof. The result is trivial for $n = 4$. For $n = 5$, it is easily verified by inspection, and we find that both C_5 and $K_{2,3}$ have the same total number of independent sets. So we may assume $n \geq 6$.

We clearly have $i(K'_{2,n-2}) < i(K_{2,n-2})$, where $K'_{2,n-2}$ is the graph obtained from $K_{2,n-2}$ by joining the two vertices in the partite set of size 2. For all $G \in \mathcal{G}(n,2)$ different from both $K_{2,n-2}$ and $K'_{2,n-2}$, Theorem 3.1.4 part 1 tells us that $i_t(G) \leq i_t(K_{2,n-2}) - 1$ for $3 \leq t \leq n-2$. For $t = 0, 1, n-1$ and n we have $i_t(G) = i_t(K_{2,n-2})$

(with the values being 1, n, 0 and 0 respectively). We have $i_2(G) \leq \binom{n}{2} - n$ (this is the number of non-edges in a 2-regular graph), and so

$$i_2(G) \leq i_2(K_{2,n-2}) + \binom{n}{2} - n - \binom{n-2}{2} - 1 = i_2(K_{2,n-2}) + n - 4. \qquad (3.1)$$

Putting all this together we get $i(G) \leq i(K_{2,n-2})$.

If G is not 2-regular then we have strict inequality in (3.1) and so $i(G) < i(K_{2,n-2})$. If G is 2-regular, then (as we will show presently) we have $i_3(G) < i_3(K_{2,n-2}) - 1$ and so again $i(G) < i(K_{2,n-2})$. To see the inequality concerning independent sets of size 3 note that in any 2-regular graph the number of independent sets of size 3 that include a fixed vertex v is the number of non-edges in the graph induced by the $n-3$ vertices $V \setminus \{v, x, y\}$ (where x and y are the neighbors of v), which is at most $\binom{n-3}{2} - (n-4)$. It follows that

$$i_3(G) \leq \frac{1}{3}\left(n\left(\binom{n-3}{2} - (n-4)\right)\right) < \binom{n-2}{3} - 1.$$

□

We also obtain some results in the range where $n < 2\delta$. Note that in the range $n \geq 2\delta$ we (conjecturally) maximize the count of independent sets by extracting as large an independent set as possible (one of size $n - \delta$). In the range $n < 2\delta$ this is still the largest independent set size, but now it is possible to have many disjoint independent sets of this size. The following conjecture seems quite reasonable.

Conjecture 3.1.6. *For $\delta \geq 1$, $\delta + 1 \leq n \leq 2\delta$, and $G \in \mathcal{G}(n, \delta)$, we have $i(G) \leq i(K_{n-\delta,n-\delta,\ldots,n-\delta,x})$, where $0 \leq x < n - \delta$ satisfies $n \equiv x \pmod{n-\delta}$.*

When $n - \delta$ divides n (that is, $x = 0$), we prove Conjecture 3.1.6 and answer the related question for $i_t(G)$ for every t.

Theorem 3.1.7. *For $\delta \geq 1$, $\delta + 1 \leq n \leq 2\delta$ with $(n-\delta)|n$, $t \in \mathbb{Z}$, and $G \in \mathcal{G}(n,\delta)$, we have $i(G) \leq i(K_{n-\delta,n-\delta,\ldots,n-\delta})$ and $i_t(G) \leq i_t(K_{n-\delta,n-\delta,\ldots,n-\delta})$.*

The case $n = 2\delta$ was originally proved in [1]. Our proof of Theorem 3.1.7 is a consequence of an upper bound on $i_t(G)$ for any $G \in \mathcal{G}(n,\delta)$ and will be given in the discussion around (3.2).

3.2 Preliminary observations

For integers n, δ and t, let $P(n,\delta,t)$ denote the statement that for every $G \in \mathcal{G}(n,\delta)$, we have $i_t(G) \leq i_t(K_{\delta,n-\delta})$. An important observation is that if we prove $P(n,\delta,t)$ for some triple (n,δ,t) with $t \geq \delta+1$, we automatically have $P(n,\delta,t+1)$. The proof introduces the important idea of considering *ordered* independent sets, that is, independent sets in which an order is placed on the vertices.

Lemma 3.2.1. *For $\delta \geq 2$ and $t \geq \delta + 1$, if $G \in \mathcal{G}(n,\delta)$ satisfies $i_t(G) \leq i_t(K_{\delta,n-\delta})$ then $i_{t+1}(G) \leq i_{t+1}(K_{\delta,n-\delta})$. Moreover, if $t < n - \delta$ and $i_t(G) < i_t(K_{\delta,n-\delta})$ then $i_{t+1}(G) < i_{t+1}(K_{\delta,n-\delta})$.*

Corollary 3.2.2. *For $\delta \geq 2$ and $t \geq \delta + 1$, $P(n,\delta,t) \Rightarrow P(n,\delta,t+1)$.*

Proof. Fix $G \in \mathcal{G}(n,\delta)$. By hypothesis, the number of ordered independent sets in G of size t is at most $(n-\delta)^{\underline{t}}$. For each ordered independent set of size t in G there are at most $n - (t+\delta)$ vertices that can be added to it to form an ordered independent set of size $t+1$ (no vertex of the independent set can be chosen, nor can any neighbor of any particular vertex in the independent set; see Figure 3.1).

This leads to a bound on the number of ordered independent sets in G of size $t+1$ of $(n-\delta)^{\underline{t}}(n-(t+\delta)) = (n-\delta)^{\underline{t+1}}$. Dividing by $(t+1)!$, we find that $i_{t+1}(G) \leq \binom{n-\delta}{t+1} = i_{t+1}(K_{\delta,n-\delta})$.

If we have $i_t(G) < \binom{n-\delta}{t}$ then we have strict inequality in the count of ordered independent sets of size t, and so also as long as $n - (\delta + t) > 0$ we have strict

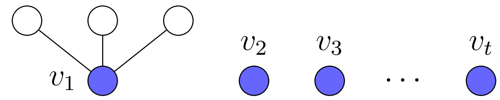

Figure 3.1. An illustration of the restrictions when adding a vertex to the ordered independent set $(v_1, v_2, v_3 \dots, v_t)$.

inequality in the count for $t+1$, and so $i_{t+1}(G) < \binom{n-\delta}{t+1}$. $\square$

Given Corollary 3.2.2, in order to prove $P(n, \delta, t)$ for all $t \geq t(\delta)$ it will be enough to prove $P(n, \delta, t(\delta))$. Many of our proofs will be by induction on n, and will be considerably aided by the following simple observation.

Lemma 3.2.3. *Fix $t \geq 3$. Suppose we know $P(n-1, \delta, t)$, and let $G \in \mathcal{G}(n, \delta)$ be such that there is $v \in V(G)$ with $G - v \in \mathcal{G}(n-1, \delta)$ (i.e. $G - v$ has minimum degree δ). Then $i_t(G) \leq i_t(K_{\delta, n-\delta})$. Equality can only occur if all of 1) $i_t(G - v) = i_t(K_{\delta, n-1-\delta})$, 2) $G - v - N(v)$ is empty (has no edges), and 3) $d(v) = \delta$ hold.*

Proof. Counting first the independent sets of size t in G that do not include v and then those that do, and bounding the former by our hypothesis on $P(n-1, \delta, t)$ and the latter by the number of subsets of size $t-1$ in $G - v - N(v)$, we have (with E_k the empty graph on k vertices)

$$
\begin{aligned}
i_t(G) &= i_t(G - v) + i_{t-1}(G - v - N(v)) \\
&\leq i_t(K_{\delta, n-1-\delta}) + i_{t-1}(E_{n-1-d(v)}) \\
&= \binom{n-1-\delta}{t} + \binom{\delta}{t} + \binom{n-1-\delta}{t-1} \\
&= \binom{n-\delta}{t} + \binom{\delta}{t} \\
&= i_t(K_{\delta, n-\delta}).
\end{aligned}
$$

The statement concerning equality is evident. $\square$

Lemma 3.2.3 allows us to focus on graphs with the property that each vertex has a neighbor of degree δ. Another simple lemma further restricts the graphs that must be considered.

Lemma 3.2.4. *If G' is obtained from G by deleting edges, then for each t we have $i_t(G) \leq i_t(G')$.*

Lemmas 3.2.3 and 3.2.4 allow us to concentrate mostly on critical graphs. In Section 2.2 (specifically Lemma 2.2.1) we obtained structural information about critical graphs in the case $\delta = 2$, while much of Section 3.5 is concerned with the same problem for $\delta = 3$.

By imagining counting ordered independent sets first, an easy upper bound on the number of independent sets of size $t \geq 2$ in a graph with minimum degree δ is

$$i_t(G) \leq \frac{n(n-(\delta+1))(n-(\delta+2))\cdots(n-(\delta+(t-1)))}{t!}. \tag{3.2}$$

This bound assumes that each vertex has degree δ, and moreover that all of the vertices in any independent set share the same δ neighbors. This upper bound is in fact tight in the situation of Theorem 3.1.7 (and is the proof of the theorem), but is not tight in general.

We will obtain better upper bounds by considering more carefully when these two conditions actually hold, as having many vertices which share the same neighborhood forces those vertices in the neighborhood to have large degree (when $n \geq 2\delta$). Most of the proofs proceed by realizing that a critical graph must have at least one of a small list of different structures in it, and we exploit the presence of a structure to significantly dampen the easy upper bound in (3.2).

3.3 Proof of Theorem 3.1.4, part 1 ($\delta = 2$)

Recall that we want to show that for $\delta = 2$, $t \geq 3$ and $G \in \mathcal{G}(n, 2)$, we have $i_t(G) \leq i_t(K_{2,n-2})$, and that for $n \geq 5$ and $3 \leq t \leq n-2$ we have equality iff $G = K_{2,n-2}$ or $K'_{2,n-2}$ (obtained from G by joining the two vertices in the partite set of size 2). We concern ourselves initially with the inequality, and discuss the cases of equality at the end. By Corollary 3.2.2, it is enough to consider $t = 3$, and we will prove this case by induction on n, the base cases $n \leq 5$ being trivial. So from here on we assume that $n > 5$ and that $P(m, 2, 3)$ has been established for all $m < n$, and let $G \in \mathcal{G}(n, 2)$ be given. By Lemmas 3.2.3 and 3.2.4 we may assume that G is critical.

We first state a well-known lemma (see e.g. [35]).

Lemma 3.3.1. *Let $k \geq 1$ and $0 \leq t \leq k+1$. In the k-path P_k we have*

$$i_t(P_k) = \binom{k+1-t}{t}.$$

Let $k \geq 3$ and $0 \leq t \leq k-1$. In the k-cycle C_k we have

$$i_t(C_k) = \binom{k-t}{t} + \binom{k-t-1}{t-1}.$$

Armed with Lemmas 3.3.1 and 2.2.1 (in particular the remark following Lemma 2.2.1) we now show that for critical G we have

$$i_3(G) < i_3(K_{2,n-2}) = \binom{n-2}{3}.$$

If G is the n-cycle, then we are done by Lemma 3.3.1. If G is a disjoint union of cycles, then choose one such, of length k, and call its vertex set Y_1, and let $Y_2 = V \setminus Y_1$. We will count the number of independent sets of size 3 in G by considering how the independent set splits across Y_1 and Y_2.

By Lemma 3.3.1, there are $\binom{k-3}{3}+\binom{k-4}{2}$ independent sets of size 3 in Y_1 (note that this is still a valid upper bound when $k=3$), and by induction there are at most $\binom{n-k-2}{3}$ independent sets of size 3 in Y_2. There are $\left(\binom{k-1}{2}-1\right)(n-k)$ independent sets with two vertices in Y_1 and one in Y_2 (the first factor here simply counting the number of non-edges in a k-cycle). Finally, there are $k\left(\binom{n-k-1}{2}-1\right)$ independent sets with one vertex in Y_1 and two in Y_2 (the second factor counting the number of non-edges in a 2-regular graph on $n-k$ vertices). The sum of these bounds is easily seen to be $\binom{n-2}{3}-k$, so strictly smaller than $\binom{n-2}{3}$.

We may now assume that G is not 2-regular. Let Y_1 be as constructed in Lemma 2.2.1. Since we are considering critical G, by the remark following Lemma 2.2.1 we may assume that $|Y_1|\geq 2$. Denote by v_1, v_2 the neighbors in Y_2 of the endpoints of the path. Note that it is possible that $v_1=v_2$, but if not then by Lemma 2.2.1 we have $v_1 \nsim v_2$. We will again upper bound $i_3(G)$ by considering the possible splitting of independent sets across Y_1 and Y_2.

By Lemma 3.3.1, there are $\binom{k-2}{3}$ independent sets of size 3 in Y_1, and by induction there are at most $\binom{n-k-2}{3}$ independent sets of size 3 in Y_2.

The number of independent sets of size 3 in G that have two vertices in Y_1 and one in Y_2 is at most

$$\binom{k-3}{2}(n-k)+\left(\binom{k-1}{2}-\binom{k-3}{2}\right)(n-k-1). \tag{3.3}$$

The first term above counts those independent sets in which neither endpoint of the k-path is among the two vertices from Y_1, and uses Lemma 3.3.1. The second term upper bounds the number of independent sets in which at least one endpoint of the k-path is among the two vertices from Y_1, and again uses Lemma 3.3.1. (Note that when $k=2$ the application of Lemma 3.3.1 is not valid, since when we remove the endvertices we are dealing with a path of length 0, outside the range of validity of the lemma; however, the bound in (3.3) is valid for $k = 2$ since it equals 1 in this case.) Finally, the number of independent sets of size 3 in G that have one vertex in Y_1 and two in Y_2 is at most

$$\left((k-2)\left(\binom{n-k}{2} - |E(Y_2)|\right)\right) + \sum_{i=1}^{2}\left(\binom{n-k-1}{2} - |E(Y_2)| + d_{Y_2}(v_i)\right).$$

The first term here counts the number of independent sets in which the one vertex from Y_1 is not an endvertex, the second factor being simply the number of non-edges in $G[Y_2]$. The second term counts those with the vertex from Y_1 being the neighbor of v_i, the second factor being the number of non-edges in $G[Y_2] - v_i$.

The sum of all of these bounds, when subtracted from $\binom{n-2}{3}$, simplifies to

$$-(k-1)n + k^2 + k - 3 + k|E(Y_2)| - d_{Y_2}(v_1) - d_{Y_2}(v_2), \tag{3.4}$$

a quantity which we wish to show is strictly positive.

Suppose first that Y_1 can be chosen so that $v_1 \neq v_2$. Recall that in this case $v_1 \nsim v_2$, so $d_{Y_2}(v_1) + d_{Y_2}(v_2) \leq |E(Y_2)|$. Combining this with $|E(Y_2)| \geq n - k$ we get that (3.4) is at most $2k - 3$, which is indeed strictly positive for $k \geq 2$.

If $v_1 = v_2 = v$, then we first note that

$$|E(Y_2)| = \frac{1}{2}\sum_{w \in Y_2} d_{Y_2}(w) \geq \frac{d_{Y_2}(v)}{2} + (n - k - 1)$$

(since $G[Y_2]$ has minimum degree 2). Inserting into (3.4) we find that (3.4) is at most

$$n - 3 + \left(\frac{k}{2} - 2\right) d_{Y_2}(v). \tag{3.5}$$

This is clearly strictly positive for $k \geq 4$, and for $k = 3$ strict positivity follows from $d_{Y_2}(v) < 2(n-3)$, which is true since in fact $d_{Y_2}(v) < n - 3$ in this case.

If $k = 2$, then (3.5) is strictly positive unless $d_{Y_2} = n - 3$ (the largest possible value it can take in this case). There is just one critical graph G with the property that for all possible choices of Y_1 satisfying the conclusions of Lemma 2.2.1 (and the subsequent remark) we have $|Y_1| = 2$, $v_1 = v_2 = v$ and $d_{Y_2}(v) = n - 3$; this is the windmill graph (see Figure 3.2) consisting of $(n-1)/2$ triangles with a single vertex in common to all the triangles, and otherwise no overlap between the vertex sets (note that the degree condition on v forces G to be connected). A direct count gives $(n-1)(n-3)(n-5)/6 < \binom{n-2}{3}$ independent sets of size 3 in this particular graph.

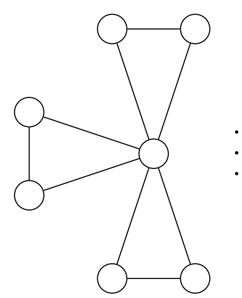

Figure 3.2. The windmill graph.

This completes the proof that $i_t(G) \leq i_t(K_{2,n-2})$ for all $t \geq 3$ and $G \in \mathcal{G}(n, 2)$. We now turn to considering the cases where equality holds in the range $n \geq 5$ and $3 \leq t \leq n - 2$. For $t = 3$ and $n = 5$, by inspection we see that we have equality iff $G = K_{2,3}$ or $K'_{2,3}$ (obtained from $K_{2,3}$ by adding an edge inside the partite set of size 2). For larger n, we prove by induction that equality can be achieved only for these two graphs. If a graph G is not edge-critical, we delete edges until we obtain a graph G' which is edge-critical, using Lemma 3.2.4 to get $i_t(G) \leq i_t(G')$. If G' is critical, then the discussion in this section shows that we cannot achieve equality. If G' is not vertex-critical, Lemma 3.2.3 and our induction hypothesis shows that we only achieve equality for G' if there is $v \in V(G')$ with $G' - v = K_{2,n-3}$ or $K'_{2,n-3}$, $G' - v - N(v)$ empty, and $d(v) = 2$. First, notice that $G' - v = K'_{2,n-3}$ implies that G' is not edge-critical, so equality can only occur when $G' - v = K_{2,n-3}$. If $G' - v = K_{2,n-3}$, the second and third conditions tell us that $N(v)$ is exactly the partite set of size 2 in $K_{2,n-3}$, that is, that $G' = K_{2,n-2}$. From here it is evident that equality can only occur for $G = K_{2,n-2}$ or $K'_{2,n-2}$.

Now for each fixed $n \geq 5$, we conclude from Lemma 3.2.1 that for $3 \leq t \leq n-2$ we *cannot* have equality unless $G = K_{2,n-2}$ or $K'_{2,n-2}$; and since the equality is trivial for these two cases, the proof is complete.

3.4 Proof of Theorem 3.1.4 part 3 ($\delta \geq 3$)

Throughout this section we set $h = |V_{>\delta}|$ and $\ell = |V_{=\delta}|$; note that $h + \ell = n$. We begin this section with the proof of Theorem 3.1.4 part 3; we then show how the method used may be improved to obtain a stronger result within the class of critical graphs (Lemma 3.4.1 below), a result which will play a role in the proof of Theorem 3.1.4, part 2 ($\delta = 3$) that will be given in Section 3.5.

Recall that we are trying to show that for $\delta \geq 3$, $t \geq 2\delta + 1$ and $G \in \mathcal{G}(n, \delta)$, we have $i_t(G) \leq i_t(K_{\delta,n-\delta})$, and that for $n \geq 3\delta + 1$ and $2\delta + 1 \leq t \leq n - \delta$ there is equality iff G is obtained from $K_{\delta,n-\delta}$ by adding some edges inside the partite set of size δ. As with Theorem 3.1.4 part 1 we begin with the inequality and discuss cases of equality at the end.

By Corollary 3.2.2 it is enough to consider $t = 2\delta + 1$. We prove $P(n, \delta, 2\delta + 1)$ by induction on n. For $n < 3\delta + 1$ the result is trivial, since in this range all $G \in \mathcal{G}(n, \delta)$ have $i_t(G) = 0$. It is also trivial for $n = 3\delta + 1$, since the only graphs G in $\mathcal{G}(n, \delta)$ with $i_t(G) > 0$ in this case are those that are obtained from $K_{\delta,n-\delta}$ by the addition of some edges inside the partite set of size δ, and all such G have $i_t(G) = 1 = i_t(K_{\delta,n-\delta})$. So from now on we assume $n \geq 3\delta + 2$ and that $P(m, \delta, 2\delta + 1)$ is true for all $m < n$, and we seek to establish $P(n, \delta, 2\delta + 1)$.

By Lemmas 3.2.3 and 3.2.4 we may restrict attention to G which are critical (for minimum degree δ). To allow the induction to proceed, we need to show that the number of ordered independent sets of size $2\delta + 1$ in G is at most $(n - \delta)^{\underline{2\delta+1}}$.

We partition ordered independent sets according to whether the first vertex is in $V_{>\delta}$ or in $V_{=\delta}$. In the first case (first vertex in $V_{>\delta}$) there are at most

$$\begin{aligned} h(n-(\delta+2))(n-(\delta+3))\cdots(n-(3\delta+1)) &= \frac{h}{n}\left(n(n-(\delta+2))^{\underline{2\delta}}\right) \\ &< \frac{h}{n}(n-\delta)^{\underline{2\delta+1}} \end{aligned} \tag{3.6}$$

ordered independent sets of size $2\delta + 1$, since once the first vertex has been chosen there are at most $n - (\delta + 2)$ choices for the second vertex, then at most $n - (\delta + 3)$ choices for the third, and so on.

In the second case (first vertex in $V_{=\delta}$) there are at most

$$\ell(n-(\delta+1))(n-(\delta+2))\cdots(n-2\delta)$$

ways to choose the first $\delta + 1$ vertices in the ordered independent set. The key observation now is that since G is vertex-critical there can be at most $\delta - 1$ vertices distinct from v with the same neighborhood as v, where v is the first vertex of the ordered independent set. It follows that one of choices 2 through δ has a neighbor w outside of $N(v)$. Since w cannot be included in the independent set, there are at most

$$(n-(2\delta+2))(n-(2\delta+3))\cdots(n-(3\delta+1))$$

choices for the final δ vertices. Combining these bounds, there are at most

$$\frac{\ell}{n}\left(n\frac{(n-(\delta+1))^{\underline{2\delta+1}}}{n-(2\delta+1)}\right) < \frac{\ell}{n}(n-\delta)^{\underline{2\delta+1}}$$

ordered independent sets of size $2\delta+1$ that begin with a vertex from $V_{=\delta}$. Combining with (3.6) we get $i_{2\delta+1}(G) < \frac{(n-\delta)^{\underline{2\delta+1}}}{(2\delta+1)!}$, as required.

This completes the proof that $i_t(G) \leq i_t(K_{\delta,n-\delta})$ for all $t \geq 2\delta+1$ and $G \in \mathcal{G}(n,\delta)$. We now turn to considering the cases where equality holds in the range $n \geq 3\delta+1$ and $2\delta+1 \leq t \leq n-\delta$. For $t = 2\delta+1$ and $n = 3\delta+1$, we clearly have equality iff G is obtained from $K_{\delta,2\delta+1}$ by adding some edges inside the partite set of size δ. For larger n, we prove by induction that equality can be achieved only for a graph of this form. If a graph G is not edge-critical, we delete edges until we obtain a graph G' which is edge-critical, using Lemma 3.2.4 to get $i_t(G) \leq i_t(G')$. If G' is critical, then the discussion in this section shows that we cannot achieve equality. If G' is not vertex-critical, Lemma 3.2.3 and our induction hypothesis shows that we only achieve equality for G' if there is $v \in V(G')$ with $G'-v$ obtained from $K_{\delta,n-\delta-1}$ by adding some edges inside the partite set of size δ, $G'-v-N(v)$ empty, and $d(v) = \delta$. First, notice that the cases where $G'-v \neq K_{\delta,n-\delta-1}$ imply that G' is not edge-critical, so in fact equality can only occur when $G'-v = K_{\delta,n-\delta-1}$. Since $d(v) = \delta$ the neighborhood of v cannot include all of the partite set of size $n-1-\delta$. If it fails to include a vertex of the partite set of size δ, there must be an edge in $G-v-N(v)$; so in fact, $N(v)$ is exactly the partite set of size δ and $G' = K_{\delta,n-\delta}$. From here it is evident that equality can only occur for G obtained from $K_{\delta,n-\delta}$ by adding some edges inside the partite set of size δ.

Now for each fixed $n \geq 3\delta+1$, we conclude from Lemma 3.2.1 that for $2\delta+1 \leq t \leq n-\delta$ we *cannot* have equality unless G is obtained from $K_{\delta,n-\delta}$ by adding some edges inside the partite set of size δ; and since the equality is trivial in these cases, the proof is complete.

The ideas introduced here to bound the number of ordered independent sets in a critical graph can be modified to give a result that covers a slightly larger range of t, at the expense of requiring n to be a little larger. Specifically we have the following.

Lemma 3.4.1. *For all $\delta \geq 3$, $t \geq \delta + 1$, $n \geq 3.2\delta$ and vertex-critical $G \in \mathcal{G}(n, \delta)$, we have $i_t(G) < i_t(K_{\delta,n-\delta})$. For $\delta = 3$ and $t = 4$ we get the same conclusion for vertex-critical $G \in \mathcal{G}(n, 3)$ with $n \geq 8$.*

Remark. The constant 3.2 has not been optimized here, but rather chosen for convenience.

Proof. By Lemma 3.2.1 it is enough to consider $t = \delta + 1$. The argument breaks into two cases, depending on whether G has at most $\delta - 2$ vertices with degree larger than m (a parameter to be specified later), or at least $\delta - 1$. The intuition is that in the former case, after an initial vertex v has been chosen for an ordered independent set, many choices for the second vertex should have at least two neighbors outside of $N(v)$, which reduces subsequent options, whereas in the latter case, an initial choice of one of the at least $\delta - 1$ vertices with large degree should lead to few ordered independent sets.

First suppose that G has at most $\delta - 2$ vertices with degree larger than m. Just as in (3.6), a simple upper bound on the number of ordered independent sets of size t whose first vertex is in $V_{>\delta}$ is

$$\frac{h}{n}\left(n(n-(\delta+2))(n-(\delta+3))\cdots(n-(2\delta+1))\right) < \frac{h}{n}(n-\delta)^{\underline{\delta+1}}. \qquad (3.7)$$

There are ℓ choices for the first vertex v of an ordered independent set that begins with a vertex from $V_{=\delta}$. For each such v, we consider the number of extensions to an ordered independent set of size $\delta+1$. This is at most

$$x(n-(\delta+2))^{\underline{\delta-1}}+y(n-(\delta+3))^{\underline{\delta-1}}+z(n-(\delta+4))^{\underline{\delta-1}} \tag{3.8}$$

where x is the number of vertices in $V(G)\setminus(\{v\}\cup N(v))$ that have no neighbors outside $N(v)$, y is the number with one neighbor outside $N(v)$, and z is the number with at least 2 neighbors outside $N(v)$. Note that $x+y+z=n-\delta-1$, and that by vertex-criticality $x\leq\delta-1$.

Let u_1 and u_2 be the two lowest degree neighbors of v. By vertex-criticality and our assumption on the number of vertices with degree greater than m, the sum of the degrees of u_1 and u_2 is at most $\delta+m$. Each vertex counted by y is adjacent to either u_1 or u_2, so counting edges out of u_1 and u_2 there are at most $m+\delta-2x-2$ such vertices.

For fixed x we obtain an upper bound on (3.8) by taking y as large as possible, so we should take $y=m+\delta-2x-2$ and $z=n-m-2\delta+x+1$. With these choices of y and z, a little calculus shows us that we obtain an upper bound by taking x as large as possible, that is, $x=\delta-1$. This leads to an upper bound on the number of ordered independent sets of size t whose first vertex is in $V_{=\delta}$ of

$$\ell\left(\begin{array}{l}(\delta-1)(n-(\delta+2))^{\underline{\delta-1}}+\\(m-\delta)(n-(\delta+3))^{\underline{\delta-1}}+\\(n-m-\delta)(n-(\delta+4))^{\underline{\delta-1}}\end{array}\right).$$

Combining with (3.7) we see that are done (for the case G has at most $\delta-2$ vertices with degree larger than m) as long as we can show that the expression above is strictly less than $\ell(n-\delta)^{\underline{\delta+1}}/n$, or equivalently that

$$n\left(\begin{array}{l}(\delta-1)(n-(\delta+2))(n-(\delta+3))+\\(m-\delta)(n-(\delta+3))(n-(2\delta+1))+\\(n-m-\delta)(n-(2\delta+1))(n-(2\delta+2))\end{array}\right) < (n-\delta)^{\underline{4}}\ . \tag{3.9}$$

We will return to this presently; but first we consider the case where G has at least $\delta-1$ vertices with degree larger than m. An ordered independent set of size $\delta+1$ in this case either begins with one of $\delta-1$ vertices of largest degree, in which case there are strictly fewer than $(n-m-1)^{\underline{\delta}}$ extensions, or it begins with one of the remaining $n-\delta+1$ vertices. For each such vertex v in this second case, the second vertex chosen is either one of the $k=k(v)\leq\delta-1$ vertices that have the same neighborhood as v, in which case there are at most $(n-(\delta+2))^{\underline{\delta-1}}$ extensions, or it is one of the $n-d(v)-1-k$ vertices that have a neighbor that is not a neighbor of v, in which case there are at most $(n-(\delta+3))^{\underline{\delta-1}}$ extensions. We get an upper bound on the total number of extensions in this second case (starting with a vertex not among the $\delta-1$ of largest degree) by taking k as large as possible and $d(v)$ as small as possible; this leads to a strict upper bound on the number of ordered independent sets of size $\delta+1$ in the case G has at least $\delta-1$ vertices with degree larger than m of

$$(\delta-1)(n-m-1)^{\underline{\delta}}+(n-\delta+1)\left(\begin{array}{l}(\delta-1)(n-(\delta+2))^{\underline{\delta-1}}+\\(n-2\delta)(n-(\delta+3))^{\underline{\delta-1}}\end{array}\right).$$

We wish to show that this is at most $(n-\delta)^{\underline{\delta+1}}$. As long as $m\geq\delta$ we have $n-m-i\leq n-\delta-i$, and so what we want is implied by

$$\left(\begin{array}{l}(\delta-1)(n-m-1)(n-m-2)+\\(n-\delta+1)(\delta-1)(n-(\delta+2))+\\(n-\delta+1)(n-2\delta)(n-(2\delta+1))\end{array}\right)\leq(n-\delta)^{\underline{3}}. \tag{3.10}$$

Setting $m = n/2$, we find that for $\delta \geq 3$, both (3.9) and (3.10) hold for all $n \geq 3.2\delta$. Indeed, in both cases at $n = 3.2\delta$ the right-hand side minus the left-hand side is a polynomial in δ (a quartic in the first case and a cubic in the second) that is easily seen to be positive for all $\delta \geq 3$; and in both cases we can check that for each fixed $\delta \geq 3$, when viewed as a function of n the right-hand side minus the left-hand side has positive derivative for all $n \geq 3.2\delta$. This completes the proof of the first statement. It is an easy check that both (3.9) and (3.10) hold for all $n \geq 8$ in the case $\delta = 3$, completing the proof of the lemma. $\square$

3.5 Proof of Theorem 3.1.4 part 2 ($\delta = 3$)

Recall that we are trying to show that for $\delta = 3$, $t \geq 3$ and $G \in \mathcal{G}(n, 3)$, we have $i_t(G) \leq i_t(K_{3,n-3})$, and that for $n \geq 6$ and $t = 3$ we have equality iff $G = K_{3,n-3}$, while for $n \geq 7$ and $4 \leq t \leq n-3$ we have equality iff G is obtained from $K_{3,n-3}$ by adding some edges inside the partite set of size 3.

For $t = 4$ and $n \geq 7$ we prove the result (including the characterization of uniqueness) by induction on n, with the base case $n = 7$ trivial. For $n \geq 8$, Lemma 3.4.1 gives strict inequality for all vertex-critical G, so we may assume that we are working with a G which is non-vertex-critical. Lemma 3.2.3 now gives the inequality $i_4(G) \leq i_4(K_{3,n-3})$, and the characterization of cases of inequality goes through exactly as it did for Theorem 3.1.4 parts 1 and 3. The result for larger t (including the characterization of uniqueness) now follows from Lemma 3.2.1.

For $t = 3$, we also argue by induction on n, with the base case $n = 6$ trivial. For $n \geq 7$, if G is not vertex-critical then the inequality $i_3(G) \leq i_3(K_{3,n-3})$ follows from Lemma 3.2.3, and the fact that there is equality in this case only for $G = K_{3,n-3}$ follows exactly as it did in the proofs of Theorem 3.1.4 parts 1 and 3. So we may assume that G is vertex-critical. We will also assume that G is edge-critical (this assumption is justified because in what follows we will show $i_3(G) < i_3(K_{3,n-3})$, and restoring the edges removed to achieve edge-criticality maintains the strictness of the inequality). Our study of critical 3-regular graphs will be based on a case analysis that adds ever more structure to the G under consideration. A useful preliminary observation is the following.

Lemma 3.5.1. *Fix $\delta = 3$. If a critical graph G has a vertex w of degree $n-3$ or greater, then $i_3(G) < i_3(K_{3,n-3})$.*

Proof. If $d(w) > n-3$ then there are no independent sets of size 3 containing w, and by Theorem 3.1.4 part 1 the number of independent sets of size 3 in $G - w$ (a graph of minimum degree 2) is at most $\binom{n-3}{3} < i_3(K_{3,n-3})$. If $d(w) = n-3$ and the two non-neighbors of w are adjacent, then we get the same bound. If they are not adjacent (so there is one independent set of size 3 containing w) and $G - w$ is not extremal among minimum degree 2 graphs for the count of independent sets of size 3, then we also get the same bound, since now $i_3(G - w) \leq \binom{n-3}{3} - 1$. If $G - w$ is extremal it is either $K_{2,n-3}$ or $K'_{2,n-3}$, and in either case w must be adjacent to everything in the partite set of size $n-3$ (to ensure that G has minimum degree 3), and then, since the non-neighbors of w are non-adjacent, it must be that $G = K_{3,n-3}$, a contradiction since we are assuming that G is critical. □

3.5.1 Regular G

If G is 3-regular then we have $i_3(G) < \binom{n-3}{3} + 1$. We see this by considering ordered independent sets of size 3. Given an initial vertex v, we extend to an ordered independent set of size 3 by adding ordered non-edges from $V \setminus (N(v) \cup \{v\})$. Since G is 3-regular there are $3n$ ordered edges in total, with at most 18 of them adjacent either to v or to something in $N(v)$. This means that the number of ordered independent sets of size 3 in G is at most

$$n((n-4)(n-5) - (3n-18)) < (n-3)(n-4)(n-5) + 6$$

with the inequality valid as long as $n \geq 7$. So from here on we may assume that G is not 3-regular, or equivalently that $V_{>3} \neq \emptyset$.

Remark. The argument above generalizes to show that δ-regular graphs have at most $\binom{n-\delta}{3} + \binom{\delta}{3}$ independent sets of size 3, with equality only possible when $n = 2\delta$.

Let $v \in V(G)$ have a neighbor in $V_{>\delta}$. By criticality $d(v) = 3$. Let w_1, w_2, and w_3 be the neighbors of v, listed in decreasing order of degree, so $d(w_1) = d$, $d(w_2) = x$ and $d(w_3) = 3$ satisfy $3 \leq x \leq d \leq n-4$ (the last inequality by Lemma 3.5.1) as well as $d > 3$ (see Figure 3.3).

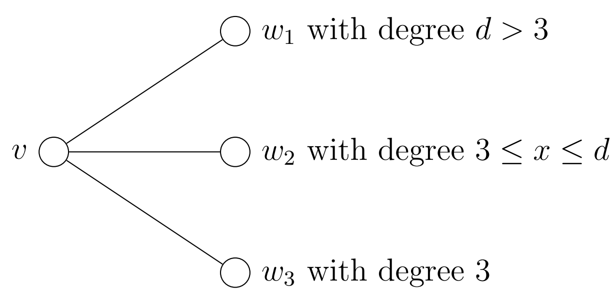

Figure 3.3. The generic situation from the end of Section 3.5.1 on.

3.5.2 No edge between w_3 and w_2

We now precede by a case analysis that depends on the value of x as well as on the set of edges present among the w_i's. The first case we consider is $w_3 \nsim w_2$. In this case we give upper bounds on the number of independent sets of size 3 which contain v and the number which do not. There are $\binom{n-4}{2} - |E(Y)|$ independent sets of size 3 which include v, where $Y = V \setminus (N(v) \cup \{v\})$. We lower bound $|E(Y)|$ by lower bounding the sum of the degrees in Y and then subtracting off the number of edges from Y to $\{v\} \cup N(v)$. This gives

$$|E(Y)| \geq \frac{3(n-4) - 2 - (d-1) - (x-1)}{2} = \frac{3(n-4) - x - d}{2}. \tag{3.11}$$

To bound the number of independent sets of size 3 which don't include v, we begin by forming G' from G by deleting v and (to restore minimum degree 3) adding an edge between w_3 and w_2 (we will later account for independent sets that contain both w_2 and w_3). The number of independent sets of size 3 in G' is, by induction, at most $i_3(K_{3,n-4})$. But in fact, we may assume that the count is strictly smaller than this. To see this, note that if we get exactly $i_3(K_{3,n-4})$ then by induction $G' = K_{3,n-4}$. For $n = 7$ this forces G to have a vertex of degree 4 and so $i_3(G) < i_3(K_{3,4})$ by Lemma 3.5.1. For $n > 7$, w_3 must be in the partite set of size $n-4$ in G' (to have degree 3) so since $w_2 \sim w_3$ (in G'), w_2 must be in the partite set of size 3. To avoid creating a vertex of degree $n-3$ in G, w_1 must be in the partite set of size $n-4$. But then all other vertices in the partite set of size $n-4$ only have neighbors of degree $n-4$ (in G), contradicting criticality.

So we may now assume that the number of independent sets of size 3 in G which do not include v is at most

$$\binom{n-4}{3} + (n-x-2), \tag{3.12}$$

the extra $n-x-2$ being an upper bound on the number of independent sets of size 3 that include both w_3 and w_2. Combining (3.11) and (3.12) we find that in this case

$$i_3(G) \leq \binom{n-4}{2} - \frac{3(n-4) - x - d}{2} + \binom{n-4}{3} + (n-x-2). \tag{3.13}$$

As long as $d < n + x - 6$ this is strictly smaller that $i_3(K_{3,n-3})$. Since $x \geq 3$ and $d < n-3$, this completes the case $w_3 \nsim w_2$.

3.5.3 Edge between w_3 and w_2, no edge between w_3 and w_1, degree of w_2 large

The next case we consider is $w_3 \sim w_2$, $w_3 \nsim w_1$, and $x > 3$. In this case we can run an almost identical the argument to that of Section 3.5.2, this time adding the edge from w_1 to w_3 when counting the number of independent sets of size 3 that don't include v. We add 1 to the right-hand side of (3.11) (to account for the fact that there is now only one edge from w_3 to Y instead of 2, and only $x-2$ from w_2 to Y instead of $x-1$) and replace (3.12) with $\binom{n-4}{3} + 1 + (n-d-2)$ (the 1 since in this case we do not need strict inequality in the induction step). Upper bounding $-d$ in this latter expression by $-x$, we get the same inequality as (3.13).

3.5.4 Edge between w_3 and w_2, edge between w_3 and w_1, degree of w_2 large

Next we consider the case $w_3 \sim w_2$, $w_3 \sim w_1$, and $x > 3$. Here we must have $w_1 \nsim w_2$, since otherwise G would not be edge-critical. The situation is illustrated in Figure 3.4. To bound $i_3(G)$, we consider v and w_3. Arguing as in Section 3.5.2

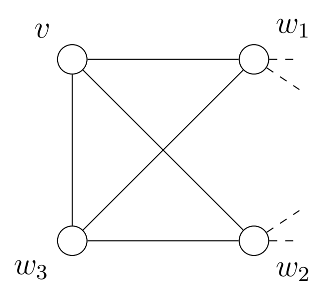

Figure 3.4. The situation in Section 3.5.4.

(around (3.11)), the number of independent sets including one of v, w_3 is at most

$$2\left(\binom{n-4}{2} - \frac{3(n-4)-(d-2)-(x-2)}{2}\right)$$

To obtain an upper bound on the number of independent sets including neither v nor w_3, we delete both vertices, add an edge from w_1 to w_2 (to restore minimum degree 3) and use induction to get a bound of

$$\binom{n-5}{3} + 1 + (n-d-2)$$

(where the $n-d-2$ bounds the number of independent sets containing both w_1 and w_2). Since $x \leq n-2$ the sum of these two bound is strictly smaller than $i_3(K_{3,n-3})$.

3.5.5 None of the above

If there is no v of degree 3 that puts us into one of the previous cases, then every v of degree 3 that has a neighbor w_1 of degree strictly greater than 3 may be assumed to have two others of degree 3, w_2 and w_3 say, with vw_2w_3 a triangle (see Figure 3.5).

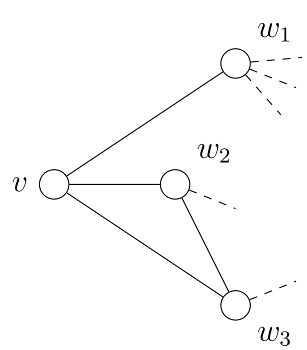

Figure 3.5. The situation in Section 3.5.5.

Since every neighbor of a vertex of degree greater than 3 has degree exactly 3 (by criticality) it follows that for every w_1 of degree greater than 3, every neighbor of w_1 is a vertex of a triangle all of whose vertices have degree 3. We claim that two of these triangles must be vertex disjoint. Indeed, if w_1 has two neighbors a and b with $a \sim b$ then the triangles associated with a and b must be the same, and by considering degrees we see that the triangle associated with any other neighbor of w_1 must be vertex disjoint from it. If a and b are not adjacent and their associated triangles have no vertex in common, then we are done; but if they have a vertex in common then (again by considering degrees) they must have two vertices in common, and the triangle associated with any other neighbor of w_1 must be vertex disjoint from both.

By suitable relabeling, we may therefore assume that G has distinct vertices w_1 (of degree greater than 3) and x, y_2, y_3, v, w_2 and w_3 (all of degree 3), with x and v adjacent to w_1, and with xy_2y_3 and vw_2w_3 forming triangles (see Figure 3.6). By considering degrees, we may also assume that the w_i's and y_i's are ordered so that $w_i \nsim y_i$ for $i = 1, 2$.

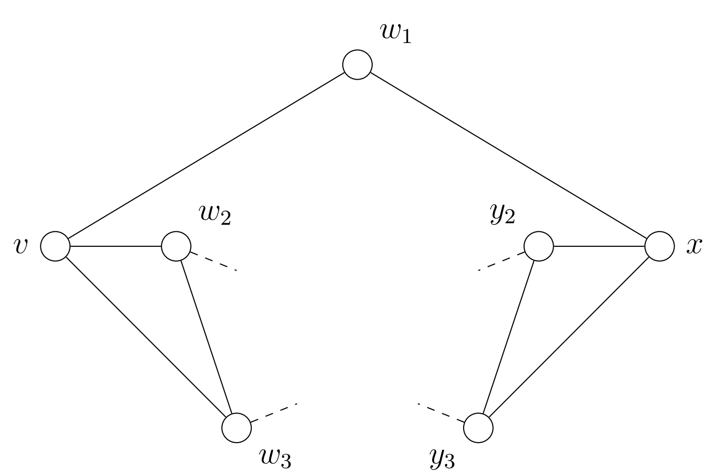

Figure 3.6. The forced structure in Section 3.5.5, before modification.

From G we create G' by removing the edges w_2w_3 and y_2y_3, and adding the edges w_2y_2 and w_3y_3 (see Figure 3.7). We will argue that $i_3(G) \leq i_3(G')$; but then by the argument of Section 3.5.2 we have $i_3(G') < i_3(K_{3,n-3})$, and the proof will be complete.

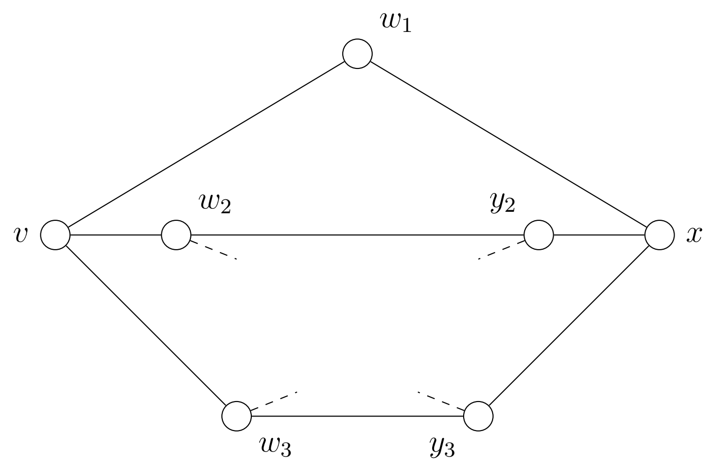

Figure 3.7. The forced structure in Section 3.5.5, after modification (i.e. in G').

Independent sets of size 3 in G partition into $I_{w_2y_2}$ (those containing both w_2 and y_2, and so neither of y_3, w_3), $I_{w_3y_3}$ (containing both w_3 and y_3), and I_{rest}, the rest. Independent sets of size 3 in G' partition into $I'_{w_2w_3}$, $I'_{y_2y_3}$, and I'_{rest}. We have $|I_{\text{rest}}| = |I'_{\text{rest}}|$ (in fact $I_{\text{rest}} = I'_{\text{rest}}$). We will show $i_3(G) \leq i_3(G')$ by exhibiting an injection from $I_{w_2y_2}$ into $I'_{w_2w_3}$ and one from $I_{w_3y_3}$ into $I'_{y_2y_3}$.

If it happens that for every independent set $\{w_2, y_2, a\}$ in G, the set $\{w_2, w_3, a\}$ is also an independent set in G', then we have a simple injection from $I_{w_2y_2}$ into $I'_{w_2w_3}$. There is only one way it can happen that $\{w_2, y_2, a'\}$ is an independent set in G but

$\{w_2, w_3, a'\}$ is not one in G'; this is when a' is the neighbor of w_3 that is not v or w_2. If $\{w_2, y_2, a'\}$ is indeed an independent set in G in this case, then letting b' be the neighbor of y_2 that is not x or y_3, we find that $\{w_2, w_3, b'\}$ is an independent set in G', but $\{w_2, y_2, b'\}$ is not one in G. So in this case we get an injection from $I_{w_2y_2}$ into $I'_{w_2w_3}$ by sending $\{w_2, y_2, a\}$ to $\{w_2, w_3, a\}$ for all $a \neq a'$, and sending $\{w_2, y_2, a'\}$ to $\{w_2, w_3, b'\}$. The injection from $I_{w_3y_3}$ into $I'_{y_2y_3}$ is almost identical and we omit the details.

CHAPTER 4

EXTREMAL H-COLORINGS

4.1 Introduction and statement of results

Fix a graph H with no isolated vertices (we will assume this for all H under consideration in this chapter). A natural extremal question to ask is the following: for a given family of graphs $\mathcal{G}$, which graphs G in $\mathcal{G}$ maximize $\hom(G, H)$?

If we assume that all graphs in $\mathcal{G}$ have n vertices, then there are several cases where this question has a trivial answer. First, if $H = K_q^{\text{loop}}$, the fully looped complete graph on q vertices, then every map $f : V(G) \to V(H)$ is an H-coloring (and so $\hom(G, K_q^{\text{loop}}) = q^n$). Second, if the empty graph $\overline{K}_n$ is contained in $\mathcal{G}$, then again every map $f : V(\overline{K}_n) \to V(H)$ is an H-coloring (i.e. $\hom(\overline{K}_n, H) = |V(H)|^n$). Motivated by this second trivial case, it is interesting to consider families $\mathcal{G}$ for which each $G \in \mathcal{G}$ has many edges.

For the family of n-vertex m-edge graphs, this question was first posed for $H = K_q$ around 1986, independently, by Linial and Wilf. Lazebnik provided an answer for $q = 2$ [46], but for general q there is still not a complete answer. However, much progress has been made (see [49] and the references therein). Recently, Cutler and Radcliffe answered this question for $H = H_{\text{ind}}$, $H = H_{\text{WR}}$, and another class of H [15, 16]. Many of the results in this family generally require a different set of extremal graphs for each choice of H.

Another interesting family to consider is the family of n-vertex d-regular graphs. Here, Kahn [39] used entropy methods to show that every n-vertex d-regular *bipartite* graph G satisfies $\hom(G, H_{\text{ind}}) \leq \hom(K_{d,d}, H_{\text{ind}})^{\frac{n}{2d}}$. Notice that when $2d|n$ this bound is achieved by $\frac{n}{2d}K_{d,d}$, the disjoint union of $n/2d$ copies of $K_{d,d}$. Galvin and Tetali [32] generalized this entropy argument, showing that for *any* H and any n-vertex d-regular bipartite G,

$$\hom(G, H) \leq \hom(K_{d,d}, H)^{\frac{n}{2d}}. \tag{4.1}$$

Kahn conjectured that (4.1) should hold for $H = H_{\text{ind}}$ for all (not necessarily bipartite) G, and Zhao [67] resolved this conjecture affirmatively, deducing the general result from the bipartite case. Interestingly, (4.1) does not hold for general H when biparticity is dropped, as there are examples of n, d, and H for which $\frac{n}{d+1}K_{d+1}$, the disjoint union of $n/(d+1)$ copies of the complete graph K_{d+1}, maximizes the number of H-colorings of graphs in this family. (For example, take H to be the disjoint union of two looped vertices; here $\log_2(\hom(G, H))$ equals the number of components of G.) Galvin proposes the following conjecture [28].

Conjecture 4.1.1 (Galvin, 2013 [28])**.** *Let G be an n-vertex d-regular graph. Then, for any H,*

$$\hom(G, H) \leq \max\{\hom(K_{d+1}, H)^{\frac{n}{d+1}}, \hom(K_{d,d}, H)^{\frac{n}{2d}}\}.$$

When $2d(d+1)|n$, this bound is achieved by either $\frac{n}{2d}K_{d,d}$ or $\frac{n}{d+1}K_{d+1}$. Evidence for this conjecture is given by Zhao [67, 68], who provided a large class of H for which $\hom(G, H) \leq \hom(K_{d,d}, H)^{\frac{n}{2d}}$. Galvin [28, 29] provides further results for various H (including triples (n, d, H) for which $\hom(G, H) \leq \hom(K_{d+1}, H)^{\frac{n}{d+1}}$) and asymptotic evidence for the conjecture.

It is clear that Conjecture 4.1.1 is true when $d = 1$, since the graph consisting of $n/2$ disjoint copies of an edge is the only 1-regular graph on n vertices. We prove the conjecture for $d = 2$ and also characterize the cases of equality.

Theorem 4.1.2. *Let G be an n-vertex 2-regular graph. Then, for any H,*

$$\hom(G,H) \leq \max\{\hom(C_3,H)^{\frac{n}{3}}, \hom(C_4,H)^{\frac{n}{4}}\}.$$

If $H \neq K_q^{loop}$, the only graphs achieving equality are $G = \frac{n}{3}C_3$ (when $\hom(C_3,H)^{\frac{n}{3}} > \hom(C_4,H)^{\frac{n}{4}}$), $G = \frac{n}{4}C_4$ (when $\hom(C_3,H)^{\frac{n}{3}} < \hom(C_4,H)^{\frac{n}{4}}$), or the disjoint union of copies of C_3 and copies of C_4 (when $\hom(C_3,H)^{\frac{n}{3}} = \hom(C_4,H)^{\frac{n}{4}}$).

It is possible for each of the equality conditions in Theorem 4.1.2 to occur. The first two sitations arise when H is a disjoint union of two looped vertices and $H = K_2$, respectively. For the third situation, we utilize that if G is connected and H is the disjoint union of H_1 and H_2, then $\hom(G,H) = \hom(G,H_1) + \hom(G,H_2)$. Letting H be the disjoint union of 8 copies of a single looped vertex and and 4 copies of K_2 gives $\hom(C_3,H)^{\frac{1}{3}} = \hom(C_4,H)^{\frac{1}{4}} = 2$. We prove Theorem 4.1.2 in Section 4.4 by analyzing H-colorings of cycles via the trace of the adjacency matrix for H.

Another natural and related family to study is $\mathcal{G}(n,\delta)$, the set of all n-vertex graphs with minimum degree δ (see e.g. Chapter 3). Our question here becomes: for a given H, which $G \in \mathcal{G}(n,\delta)$ maximizes $\hom(G,H)$? Since removing edges increases the number of H-colorings, it is tempting to believe that the answer to this question will be a graph that is δ-regular (or close to δ-regular). This in fact is not the case, even for $H = H_{\text{ind}}$. The following result appears in [27].

Theorem 4.1.3 (Galvin, 2011 [27])**.** *For $\delta \geq 1$, $n \geq 8\delta^2$, and $G \in \mathcal{G}(n,\delta)$, we have $\hom(G,H_{ind}) \leq \hom(K_{\delta,n-\delta},H_{ind})$, with equality only for $G = K_{\delta,n-\delta}$.*

With Conjecture 4.1.1 and the results of Theorem 4.1.3 in mind, the following conjecture is natural.

Conjecture 4.1.4. *Fix $\delta \geq 1$ and H. There exists a constant $c(\delta,H)$ (depending on*

δ and H) such that for $n \geq c(\delta, H)$ and $G \in \mathcal{G}(n, \delta)$,

$$\hom(G, H) \leq \max\{\hom(K_{\delta+1}, H)^{\frac{n}{\delta+1}}, \hom(K_{\delta,\delta}, H)^{\frac{n}{2\delta}}, \hom(K_{\delta,n-\delta}, H)\}.$$

This conjecture stands in marked contrast to the situation for the family of n-vertex m-edge graphs, where each choice of H seems to create a different set of extremal graphs. Here, we conjecture that for *any* H, one of exactly three situations can occur. For $2\delta(\delta+1)|n$ and n large, this represents the best possible conjecture, since for H consisting of a disjoint union of two looped vertices, $H = K_2$, and $H = H_{\text{ind}}$, the number of H-colorings of a graph $G \in \mathcal{G}(n, \delta)$ is maximized by $G = \frac{n}{\delta+1}K_{\delta+1}$, $G = \frac{n}{2\delta}K_{\delta,\delta}$, and $G = K_{\delta,n-\delta}$, respectively.

The purpose of this chapter is to make progress toward Conjecture 4.1.4. We first fully resolve the conjecture for $\delta = 1$ and $\delta = 2$, and characterize the graphs that achieve equality. Before we formally state these theorems, we highlight the degree conventions and notations that we will follow for the remainder of the chapter.

Convention. For a vertex v, let $d(v)$ denote the degree of v, where loops count *once* toward the degree. While δ will always refer to the minimum degree of a graph G, in this chapter Δ will always denote the maximum degree of a graph H (unless explicity stated otherwise).

Theorem 4.1.5. ($\delta = 1$). *Fix H, $n \geq 2$ and $G \in \mathcal{G}(n, 1)$.*

1. *Suppose that $H \neq K_\Delta^{loop}$ satisfies $\sum_{v \in V(H)} d(v) \geq \Delta^2$. Then* $\hom(G, H) \leq \hom(K_2, H)^{\frac{n}{2}}$, *with equality only for $G = \frac{n}{2}K_2$.*
2. *Suppose that H satisfies $\sum_{v \in V(H)} d(v) < \Delta^2$, and let $n_0 = n_0(H)$ be the smallest integer in $\{3, 4, \ldots\}$ satisfying $\sum_{v \in V(H)} d(v) < \left(\sum_{v \in V(H)} d(v)^{n_0-1}\right)^{\frac{2}{n_0}}$.*
   1. *If $2 \leq n < n_0$, then* $\hom(G, H) \leq \hom(K_2, H)^{\frac{n}{2}}$, *with equality only for $G = \frac{n}{2}K_2$ [unless $n = n_0 - 1$ and $\sum_{v \in V(H)} d(v) = \left(\sum_{v \in V(H)} d(v)^{n_0-2}\right)^{\frac{2}{n_0-1}}$, in which case $G = K_{1,n-1}$ also achieves equality].*

2. *If $n \geq n_0$, then* $\hom(G,H) \leq \hom(K_{1,n-1},H)$, *with equality only for* $G = K_{1,n-1}$.

Theorem 4.1.6. **($\delta = 2$).** *Fix H.*

1. *Suppose that $H \neq K_\Delta^{loop}$ satisfies* $\max\{\hom(C_3,H)^{\frac{1}{3}}, \hom(C_4,H)^{\frac{1}{4}}\} \geq \Delta$. *Then for all $n \geq 3$ and $G \in \mathcal{G}(n,2)$,* $\hom(G,H) \leq \max\{\hom(C_3,H)^{\frac{n}{3}}, \hom(C_4,H)^{\frac{n}{4}}\}$, *with equality only for $G = \frac{n}{3}C_3$ (when* $\hom(C_3,H)^{\frac{n}{3}} > \hom(C_4,H)^{\frac{n}{4}}$*), $G = \frac{n}{4}C_4$ (when* $\hom(C_3,H)^{\frac{n}{3}} < \hom(C_4,H)^{\frac{n}{4}}$*), or the disjoint union of copies of C_3 and copies of C_4 (when* $\hom(C_3,H)^{\frac{n}{3}} = \hom(C_4,H)^{\frac{n}{4}}$*).*

2. *Suppose that H satisfies* $\max\{\hom(C_3,H)^{\frac{1}{3}}, \hom(C_4,H)^{\frac{1}{4}}\} < \Delta$. *Then there exists a constant c_H such that for $n > c_H$ and $G \in \mathcal{G}(n,2)$,* $\hom(G,H) \leq \hom(K_{2,n-2},H)$, *with equality only for $G = K_{2,n-2}$.*

Theorems 4.1.5 and 4.1.6 are easily seen to resolve Conjecture 4.1.4 when $\delta = 1$ and $\delta = 2$, respectively. Notice that if G' is obtained from G by deleting some edges from G, then $\hom(G,H) \leq \hom(G',H)$. Because of this, their proofs focus on G which are edge-critical for δ (recall that we will be using the notations and definitions from Section 2.2 in this chapter).

The edge-critical graphs in $\mathcal{G}(n,1)$ are disjoint unions of stars, and the proof of Theorem 4.1.5 critically uses this fact. Theorem 4.1.6 relies on a nice structural decomposition of edge-critical graphs in $\mathcal{G}(n,2)$ (in particular Corollary 2.2.2) and also uses Theorem 4.1.2. The global structure of edge-critical graphs in $\mathcal{G}(n,\delta)$ for $\delta \geq 3$ is not very well understood (note that Section 3.5 exploits the local structure of critical graphs when $\delta = 3$).

We also make some progress in the general δ case of Conjecture 4.1.4 by providing a large class of H for which $K_{\delta,n-\delta}$ maximizes the number of H-colorings.

Theorem 4.1.7. *Fix δ and H. Suppose that H satisfies*

$$\sum_{v \in V(H)} d(v) < \Delta^2. \tag{4.2}$$

Then there exists a constant c_H such that for all $n \geq (c_H)^\delta$ and $G \in \mathcal{G}(n, \delta)$, $\hom(G, H) \leq \hom(K_{\delta,n-\delta}, H)$, *with equality only for $G = K_{\delta,n-\delta}$.*

If H has the property that all vertices of degree Δ share the same Δ neighbors, then the same result holds for all $n \geq c_H\delta^2$.

We prove Theorem 4.1.7 in Section 4.2 by partitioning $\mathcal{G}(n, \delta)$ based on the size of a maximal matching. The following corollary to Theorem 4.1.7 warrants special attention, and is immediate.

Corollary 4.1.8. *Suppose that $H \neq K_\Delta^{loop}$ has a looped dominating vertex, or that H satisfies (4.2) and H has a unique vertex of degree Δ. Then there exists a constant c_H such that for all $n \geq c_H\delta^2$ and $G \in \mathcal{G}(n, \delta)$,* $\hom(G, H) \leq \hom(K_{\delta,n-\delta}, H)$, *with equality only for $G = K_{\delta,n-\delta}$.*

The graphs $H = H_{\text{ind}}$ and $H = H_{\text{WR}}$ satisfy the conditions of Corollary 4.1.8, so in particular we provide an alternate proof of Galvin's result for $H = H_{\text{ind}}$ [27]. Another important graph H which satisfies the conditions of Corollary 4.1.8 is the *k-state hard-core constraint graph* $H(k)$ ($k \geq 1$), the graph with vertex set $\{0, 1, \ldots, k\}$ and edge $i \sim_{H(k)} j$ if $i + j \leq k$ (note that $H(1) = H_{\text{ind}}$). This graph naturally occurs in the study of multicast communications networks, and has been considered in e.g. [31, 57].

Notice that the condition on H in Theorem 4.1.7 is necessary but not sufficient for $\hom(G, H) \leq \hom(K_{\delta,n-\delta}, H)$ for all $G \in \mathcal{G}(n, \delta)$. Indeed, if H is a path on 3 vertices with a loop on one endpoint of the path, then $\sum_{v \in V(H)} d(v) = 5$ while $\Delta = 2$. However, for large enough n and $G \in \mathcal{G}(n, 2)$, $\hom(G, H) \leq \hom(K_{2,n-2}, H)$, as can be seen by computing $\hom(C_3, H)$, $\hom(C_4, H)$, and applying Theorem 4.1.6.

It is also interesting to consider a maximum degree condition in addition to a minimal degree condition (see e.g. [1, 33, 39]). Let $\mathcal{G}(n, \delta, D)$ denote the set of graphs on n vertices with minimum degree δ and maximum degree at most D. Which graphs

$G \in \mathcal{G}(n, \delta, D)$ maximize $\hom(G, H)$? For a fixed δ, this question is interesting for the H with the property that $\hom(G, H) \leq \hom(K_{\delta,n-\delta}, H)$ for all $G \in \mathcal{G}(n, \delta)$, as $K_{\delta,n-\delta} \in \mathcal{G}(n, \delta, D)$ only when $D \geq n - \delta$.

For $\delta = 1$ and any $D \geq 1$, we provide an answer.

Theorem 4.1.9. *Fix H and $D \geq 1$. For any $G \in \mathcal{G}(n, 1, D)$,*

$$\hom(G, H) \leq \max\{\hom(K_2, H)^{\frac{n}{2}}, \hom(K_{1,D}, H)^{\frac{n}{D+1}}\},$$

with the cases of equality as in Theorem 4.1.5.

The proof of Theorem 4.1.9 is given in Section 4.3, and again utilizes the fact that edge-critical graphs for $\delta = 1$ are disjoint unions of stars.

Our results generalize naturally to weighted H-colorings, where for weight set Λ each $f \in \mathrm{Hom}(G, H)$ is given weight $w_\Lambda(f)$ (as defined in Section 2.1). Although the proofs of the weighted versions come with almost no extra effort, for the clarity of presentation we will not do this here (the necessary changes are described in [19]).

4.2 Proof of Theorem 4.1.7

Suppose that we have a graph H satisfying

$$\sum_{v \in V(H)} d(v) < \Delta^2, \tag{4.3}$$

and let G be a graph with minimum degree δ. Let M be the edge set of a matching of maximum size in G and I the set of unmatched vertices.

We first derive some structural properties of our graph G based on M. Since M is maximal, I forms an independent set. Furthermore, suppose that x_1 and x_2 in $V(G)$ are matched in M. If x_1 has at least two edges into I, then x_2 cannot be adjacent to any vertex in I, as this would create an augmenting path (a path which starts and ends at distinct unmatched vertices and alternates edges in and out of M) of length 3 and therefore a matching of larger size. In summary:

At most one vertex in an edge of the matching M can have degree at least 2 into I, and if one has degree at least 2 into I then the other has degree 0 into I. (4.4)

For each edge in M, put the endpoint with the largest degree into I in a set $J \subset V(G)$, and put the other endpoint in a set $K \subset V(G)$. (If the degrees are equal, make an arbitrary choice.) A schematic picture of G is shown in Figure 4.2; there are at most $|K| = |M|$ total edges between I and K.

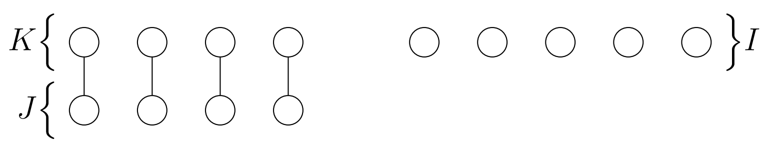

Figure 4.1. The relevant structure for G.

Also, if there are more than $|M|$ vertices in I that are adjacent to both endpoints of some edge in M, then by the pigeonhole principle there are distinct $y_1, y_2 \in I$ that are adjacent to both endpoints of some fixed edge in M. This would force both endpoints of an edge in M to have degree at least 2 into I, contradicting (4.4). Therefore we have:

There are at most $|M|$ vertices in I adjacent to both endpoints of some edge in M. (4.5)

In particular, suppose $n \geq 3\delta - 2$. Then if $|M| < \delta$ we have $|I| \geq \delta$. Since each $x \in I$ has at least δ neighbors to M, each $x \in I$ is adjacent to both endpoints of some edge in M. Since $|I| \geq \delta > |M|$, this contradicts (4.5). Therefore if $n \geq 3\delta - 2$ we have $|M| \geq \delta$ for all $G \in \mathcal{G}(n, \delta)$. We will first analyze the graphs where $|M| = \delta$ and then the graphs where $|M| > \delta$.

Case 1: Suppose that $|M| = \delta$ and $n > 3\delta$, so by (4.5) at most δ vertices in I are adjacent to both endpoints of some edge in M. Then there is at least one vertex in I that is adjacent to exactly one endpoint of each edge in M. However, this shows that *no* vertex in I can be adjacent to both endpoints of any edge in M, since any vertex in I adjacent to both vertices of an edge in M would force one endpoint of M to have degree at least 2 into I and the other endpoint of M to have degree at least 1 into I (contradicting (4.4)). It follows that each vertex in I must be adjacent to each vertex in J, and so by (4.4) there are no edges between I and K.

Now suppose $k_1, k_2 \in K$ with $k_1 \sim k_2$. Then there exist distinct $j_1, j_2 \in J$ with $k_1 \sim_M j_1$ and $k_2 \sim_M j_2$. Letting i_1 and i_2 denote any two distinct vertices in I (and recalling that everything in I is adjacent to everything in J), $i_1 \sim j_1 \sim k_1 \sim k_2 \sim j_2 \sim i_2$ is an augmenting path of length 5, which contradicts the maximality of M. Therefore K is an independent set and so $K \cup I$ is an independent set. Since G has minimum degree δ, every vertex in $K \cup I$ is adjacent to every vertex in J, and so G must be the complete bipartite graph $K_{\delta,n-\delta}$ with some edges added to the size δ partition class.

We now show that adding any edge to the size δ partition class in $K_{\delta,n-\delta}$ will strictly decrease the number of H-colorings. Since H cannot contain K_Δ^{loop} (by (4.3)), there are two (possibly non-distinct) non-adjacent neighbors of a vertex in H with degree Δ. If any edge is added to the size δ partition class in $K_{\delta,n-\delta}$, then it is impossible for any H-coloring to color the endpoints of that edge with the non-adjacent vertices in H, but such a coloring is possible in $K_{\delta,n-\delta}$. Since any H-coloring of $K_{\delta,n-\delta}$ with an edge added is an H-coloring of $K_{\delta,n-\delta}$, this shows that the number of H-colorings strictly decreases whenever an edge is added to $K_{\delta,n-\delta}$

In summary, we have shown that if G satisfies $n > 3\delta$, $|M| \leq \delta$, and $G \neq K_{\delta,n-\delta}$ then $\hom(G,H) < \hom(K_{\delta,n-\delta},H)$.

Case 2: Now suppose that $|M| = k \geq \delta + 1$. We will show that for large enough n we have $\hom(G,H) < \hom(K_{\delta,n-\delta},H)$, which will complete the proof.

Let $S(\delta,H)$ denote the vectors in $V(H)^\delta$ with the property that the elements of the vector have Δ common neighbors, and let $s(\delta,H) = |S(\delta,H)|$. (Note that $S(\delta,H) \neq \emptyset$, since if $v \in V(H)$ with $d(v) = \Delta$ then $(v,v,\ldots,v) \in S(\delta,H)$.) We obtain a lower bound on $\hom(K_{\delta,n-\delta},H)$ by coloring the size δ partition class using an element of $S(\delta,H)$, and then independently coloring the vertices in the size $n-\delta$ partition class using the Δ common neighbors. This gives

$$\hom(K_{\delta,n-\delta},H) \geq s(\delta,H)\Delta^{n-\delta}.$$

We will show that for n large and $k \geq \delta+1$, $\hom(G,H) < s(\delta,H)\Delta^{n-\delta}$.

Our initial coloring scheme will be to color J arbitrarily first, then K, then I, keeping track of an upper bound on the number of choices we have for the color at each vertex. If a vertex in J is colored with $v \in V(H)$, its neighbor in M has at most $d(v)$ choices for a color. Since each vertex in I is adjacent to some vertex in $J \cup K$, there are at most Δ choices for the color of each vertex in I. This gives

$$\hom(G,H) \leq \left(\sum_{v\in V(H)} d(v)\right)^k \Delta^{n-2k} = \Delta^n \left(\frac{\sum_{v\in V(H)} d(v)}{\Delta^2}\right)^k.$$

Recalling that H satisfies (4.3), if $k > \delta \log \Delta / \log\left(\frac{\Delta^2}{\sum_{v\in V(H)} d(v)}\right) = C_H\delta$ this upper bound is smaller than $\Delta^{n-\delta}$. So we may further assume that $\delta + 1 \leq k \leq C_H\delta$.

Let $I' \subset I$ be the set of vertices in I with neighbors exclusively in J, so by (4.4) we have $|I'| \geq n-3k$. Since each vertex in I' has at least δ neighbors in J, we imagine each $x \in I'$ picking a subset of size δ from J (from among the $\binom{k}{\delta}$ possibilities). By the pigeonhole principle there is a set $J_1 \subset J$ with $|J_1| = \delta$ and at least $(n-3k)/\binom{k}{\delta}$ vertices in I' adjacent to each vertex in J_1. See Figure 4.2.

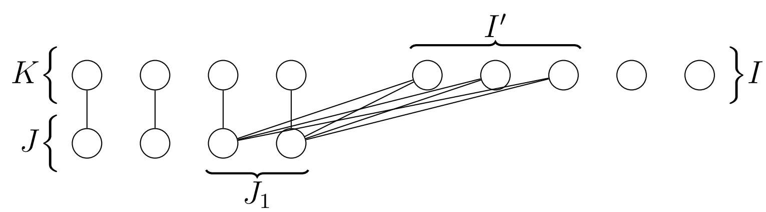

Figure 4.2. Vertices in I' adjacent to every vertex in J_1.

We'll partition the H-colorings of G based on whether the colors on J_1 form a vector in $S(\delta, H)$ or not. If they do, then we next color $J \setminus J_1$, then K, and then I, giving at most

$$s(\delta, H) \cdot \left(\sum_{v \in V(H)} d(v) \right)^{k-\delta} \cdot \Delta^{\delta} \cdot \Delta^{n-2k}$$

H-colorings of G of this type.

If the colors on J_1 do not form a vector in $S(\delta, H)$, then we have at least $(n-3k)/\binom{k}{\delta}$ vertices in I (namely those in I') that have at most $\Delta - 1$ choices for their color (here we're using that all edges are present between I' and J_1). Utilizing only this restriction, coloring $J \setminus J_1$, then K, then I gives at most

$$\left(\sum_{v \in V(H)} d(v) \right)^{k} \Delta^{n-2k} \left(\frac{\Delta - 1}{\Delta} \right)^{\frac{n-3k}{\binom{k}{\delta}}}$$

H-colorings of G of this type. Therefore, using $k \le C_H\delta$ and $\binom{a}{b} \le \left(\frac{ea}{b}\right)^b$, we have

$$\begin{aligned}
\hom(G,H) &\le s(\delta,H)\cdot\left(\sum_{v\in V(H)} d(v)\right)^{k-\delta} \Delta^{n-2k+\delta} \\
&\quad + \left(\sum_{v\in V(H)} d(v)\right)^{k} \Delta^{n-2k}\left(\frac{\Delta-1}{\Delta}\right)^{\frac{n-3k}{\binom{k}{\delta}}} \\
&\le s(\delta,H)\cdot\left(\frac{\sum_{v\in V(H)} d(v)}{\Delta^2}\right)^{k-\delta}\Delta^{n-\delta} \\
&\quad + \left(\frac{\sum_{v\in V(H)} d(v)}{\Delta^2}\right)^{k}\Delta^{n}\left(\frac{\Delta-1}{\Delta}\right)^{\frac{n-3C_H\delta}{(e\cdot C_H)^\delta}}
\end{aligned}$$

so that

$$\hom(G,H) \le s(\delta,H)\Delta^{n-\delta}\left(\frac{\sum_{v\in V(H)} d(v)}{\Delta^2}\right)^{k-\delta}(1+r_1(\delta,H))$$

where

$$r_1(\delta,H) = \frac{1}{s(\delta,H)}\left(\frac{\sum_{v\in V(H)} d(v)}{\Delta}\right)^{\delta}\left(\frac{\Delta-1}{\Delta}\right)^{\frac{n-3C_H\delta}{(e\cdot C_H)^\delta}}.$$

For $\delta+1 \le k \le C_H\delta$ and $n \ge (c_H)^\delta$, this is smaller than $s(\delta,H)\Delta^{n-\delta}$.

We sharpen the bounds on n when all of the vertices of H with degree Δ have identical neighborhoods. (Notice that this only requires a new argument for the range $\delta+1 \le k \le C_H\delta$.) Recall from Section 2.2 that $V_{=\Delta}(H)$ denotes the set of degree Δ vertices, and so by assumption each vertex in $V_{=\Delta}(H)$ has the same Δ neighbors (and also $s(\delta,H) = |V_{=\Delta}(H)|^\delta$). Our strategy is to find a set of δ vertices in J with large degree to I individually instead of finding those with a large common neighborhood in I.

Let $J = \{x_1,\dots,x_k\}$ and let a_t denote the number of edges from x_t to I for each t. Without loss of generality, assume $a_1 \ge a_2 \ge \cdots \ge a_k$. Since I is an independent set of size $n-2k$, there are at least $\delta(n-2k)$ edges from I to $J\cup K$. Since the degree to I of each vertex in K is at most 1 and each $a_t \leq n-2k$, we have

$$\delta(n-2k)-k \leq \sum_{t=1}^{k} a_t \leq (k-\delta+1)a_\delta + (\delta-1)(n-2k),$$

since $(n-2k)-k$ is a lower bound on the number of edges from I to J, $a_\delta + a_{\delta+1} + \cdots + a_k \leq (k-\delta+1)a_\delta$ (by the ordering of the a_i's), and $a_{\delta-1} \leq \cdots \leq a_1 \leq n-2k$. This gives

$$a_\delta \geq \frac{n-3k}{k-\delta+1}.$$

Now set $J_2 = \{x_1, \dots, x_\delta\}$. We first upper bound the number of H-colorings of G that color each vertex in J_2 with a color from $V_{=\Delta}(H)$. By coloring $J \setminus J_2$ arbitrarily, then coloring K, then coloring I, we have at most

$$s(\delta,H)\left(\sum_{v\in V(H)} d(v)\right)^{k-\delta} \Delta^{n-2k+\delta} = s(\delta,H)\Delta^{n-\delta}\left(\frac{\sum_{v\in V(H)} d(v)}{\Delta^2}\right)^{k-\delta} \tag{4.6}$$

H-colorings of G of this type.

The number of H-colorings of G that have some vertex of J_2 colored from $V(H) \setminus V_{=\Delta}(H)$ can be given an upper bound through similar means. Here, at least $\frac{n-3k}{k-\delta+1}$ vertices in I will have at most $\Delta - 1$ choices of a color for each coloring of $J \cup K$. Using $k \leq C_H\delta$, we have at most

$$\begin{aligned}
&\left(\sum_{v\in V(H)} d(v)\right)^{k} \Delta^{n-2k}\left(\frac{\Delta-1}{\Delta}\right)^{\frac{n-3k}{k-\delta+1}} \\
&\qquad \leq \Delta^{n-\delta}\left(\frac{\sum_{v\in V(H)} d(v)}{\Delta^2}\right)^{k}\Delta^{\delta}\left(\frac{\Delta-1}{\Delta}\right)^{\frac{n-3C_H\delta}{C_H\delta-\delta+1}}
\end{aligned} \tag{4.7}$$

H-colorings of G of this type. Combining (4.6) and (4.7) we find that

$$\hom(G,H) \leq s(\delta,H)\Delta^{n-\delta}\left(\frac{\sum_{v\in V(H)} d(v)}{\Delta^2}\right)^{k-\delta}(1+r_2(\delta,H))$$

where

$$r_2(\delta, H) = \frac{1}{s(\delta, H)} \left(\frac{\sum_{v \in V(H)} d(v)}{\Delta} \right)^{\delta} \left(\frac{\Delta - 1}{\Delta} \right)^{\frac{n - 3C_H\delta}{C_H\delta - \delta + 1}}$$

For $\delta + 1 \leq k \leq C_H\delta$ and $n > c_H\delta^2$, this is smaller than $s(\delta, H)\Delta^{n-\delta}$.

4.3 Proof of Theorems 4.1.5 and 4.1.9 ($\delta = 1$)

We begin with the proof of Theorem 4.1.5. Recall that we will assume that $|V(H)| = q$, and furthermore we will assume that G is edge-critical, so G has no edge between two vertices of degree larger than one. (This will give us the inequalities desired; we will address the uniqueness statements in the theorems separately.) In particular, G is the disjoint union of stars, so we can write $G = \cup_i K_{1,n_i-1}$, where $\sum_i n_i = n$. See Figure 4.3.

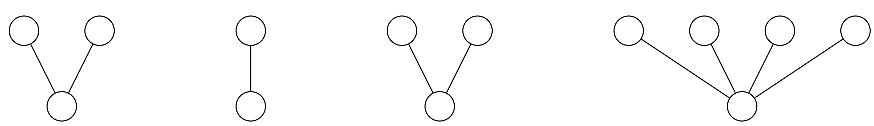

Figure 4.3. A graph G that is the disjoint union of stars; here we may take $n_1 = 2$, $n_2 = 1$, $n_3 = 2$, and $n_4 = 4$.

Since the stars K_{1,n_i-1} are disjoint and can therefore be colored independently, we have

$$\hom(G, H) = \prod_i \hom(K_{1,n_i-1}, H) = \prod_i \hom(K_{1,n_i-1}, H)^{\frac{n_i}{n_i}}.$$

If x is an integer value in $[2, n]$ that maximizes $\hom(K_{1,x-1}, H)^{\frac{1}{x}}$, then

$$\hom(G, H) \leq \prod_i \hom(K_{1,x-1}, H)^{\frac{n_i}{x}} = \hom(K_{1,x-1}, H)^{\frac{n}{x}}, \tag{4.8}$$

with equality occurring (when $x|n$) for $\frac{n}{x}K_{1,x-1}$. Because of this, it will be useful to know the integer value(s) of $x \geq 2$ that maximize $\hom(K_{1,x-1}, H)^{\frac{1}{x}}$.

First we derive a formula for $\hom(K_{1,x-1}, H)^{\frac{1}{x}}$ for each integer $x \geq 2$. Notice that all H-colorings of $K_{1,x-1}$ can be obtained by coloring the center of the star with any $v \in V(H)$ and then coloring the leaves (independently) with any neighbor of v, so

$$\hom(K_{1,x-1}, H) = \sum_{v \in V(H)} d(v)^{x-1}. \tag{4.9}$$

Now (4.9) holds for each integer $x \geq 2$, and so it will be useful to know the maximum value of

$$\left(\sum_{v \in V(H)} d(v)^{x-1} \right)^{\frac{1}{x}} \tag{4.10}$$

over all integers $x \geq 2$. In fact, we will study (4.10) in a slightly more general setting; for the remainder of this proof we will analyze (4.10) over all *real* numbers $x \geq 2$.

Recall that we are assuming that H has no isolated vertices, so for all $v \in V(H)$ we have $1 \leq d(v) \leq \Delta$. Since there exists a $w \in V(H)$ with $d(w) = \Delta$, we have

$$\left(\sum_{v \in V(H)} d(v)^{x-1} \right)^{\frac{1}{x}} \to \Delta \qquad \text{as } x \to \infty. \tag{4.11}$$

To obtain more information, for a fixed real $x \geq 2$ let $a = a(x, H) \in \mathbb{R}$ be such that

$$d(v_1)^{x-1} + \cdots + d(v_q)^{x-1} = a^x.$$

Since $1 \leq d(v) \leq \Delta$ for all $v \in V(H)$, for any $\varepsilon > 0$ we have

$$d(v_1)^{x-1+\varepsilon} + \cdots + d(v_q)^{x-1+\varepsilon} \leq \Delta^\varepsilon \left(d(v_1)^{x-1} + \cdots + d(v_q)^{x-1}\right) = \Delta^\varepsilon a^x, \qquad (4.12)$$

with strict inequality if $d(v_i) < \Delta$ for some i. Therefore for any $\varepsilon > 0$ (4.12) gives

$$a > \Delta \quad \Longrightarrow \quad \left(\sum_{v \in V(H)} d(v)^{x-1+\varepsilon}\right)^{\frac{1}{x+\varepsilon}} < \left(\sum_{v \in V(H)} d(v)^{x-1}\right)^{\frac{1}{x}}. \qquad (4.13)$$

If $a = \Delta$ and $d(v_i) = \Delta$ for all i, then (4.12) gives $q\Delta^{x-1} = \Delta^x$ so $H = K_q^{\text{loop}}$. If $a = \Delta$ and $d(v_i) < \Delta$ for some i, then for any $\varepsilon > 0$ (4.12) gives

$$\left(\sum_{v \in V(H)} d(v)^{x-1+\varepsilon}\right)^{\frac{1}{x+\varepsilon}} < \Delta = \left(\sum_{v \in V(H)} d(v)^{x-1}\right)^{\frac{1}{x}}. \qquad (4.14)$$

Finally, for any $\varepsilon > 0$, if $a < \Delta$ then $\Delta^\varepsilon a^x < \Delta^{x+\varepsilon}$, and so (4.12) gives

$$a < \Delta \quad \Longrightarrow \quad \left(\sum_{v \in V(H)} d(v)^{x-1+\varepsilon}\right)^{\frac{1}{x+\varepsilon}} < \Delta. \qquad (4.15)$$

This already provides a substantial amount of information, fully analyzing the graphs H where $\sum_{v \in V(H)} d(v) \geq \Delta^2$ (here we focus on $x = 2$ and so the condition on H means $a = a(2, H) \geq \Delta$). Indeed, for $H \neq K_q^{\text{loop}}$, (4.13) and (4.14) applied at $x = 2$ imply

$$\left(\sum_{v \in V(H)} d(v)^{y-1}\right)^{\frac{1}{y}} < \left(\sum_{v \in V(H)} d(v)\right)^{\frac{1}{2}}$$

for any $y > 2$, which implies $\hom(K_{1,x-1}, H)^{\frac{1}{x}} < \hom(K_{1,1}, H)^{\frac{1}{2}}$ for any integer $x > 2$.

For the graphs H satisfying $\sum_{v \in V(H)} d(v) < \Delta^2$, we may already obtain a statement for large n (using (4.11) and also (4.15) at $x = 2$), but with an additional argument we can obtain a statement for all n. We utilize the following lemma, which we will prove momentarily.

Lemma 4.3.1. *The function* $\left(\sum_{v\in V(H)} d(v)^{x-1}\right)^{\frac{1}{x}}$ *has at most one local maximum or minimum.*

If we assume Lemma 4.3.1, then (4.11), (4.13), (4.14), and (4.15) show that for $H \neq K_q^{\text{loop}}$ the function $\left(\sum_{v\in V(H)} d(v)^{x-1}\right)^{\frac{1}{x}}$ is either decreasing to Δ on $(2,\infty)$, increasing to Δ on $(2,\infty)$, or decreasing on $(2,x_0)$ and increasing to Δ on (x_0,∞) for some $x_0 > 2$. See Figure 4.4 for the possible behaviors of $\left(\sum_{v\in V(H)} d(v)^{x-1}\right)^{\frac{1}{x}}$. So if $H \neq K_q^{\text{loop}}$, then (4.8) shows that Theorem 4.1.5 holds for any edge-critical $G \in \mathcal{G}(n,1)$. This implies that the upper bounds given in Theorem 4.1.5 hold for any $G \in \mathcal{G}(n,1)$.

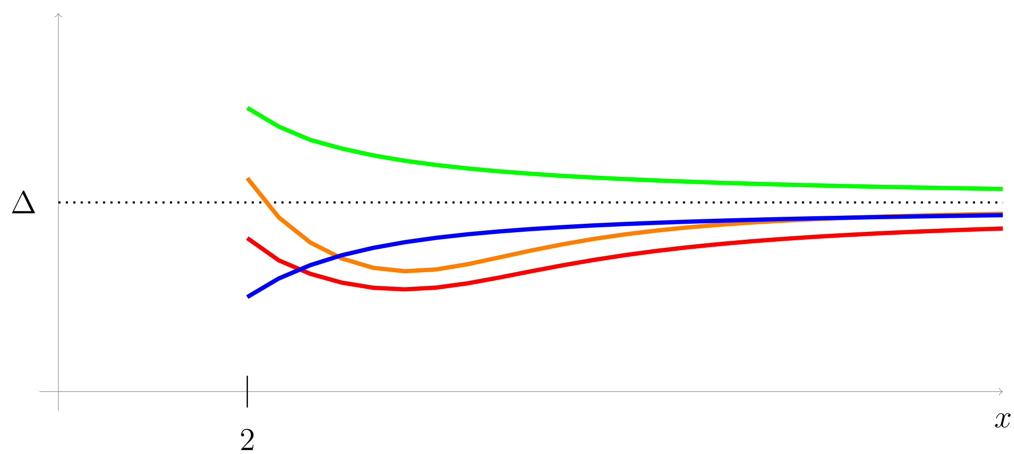

Figure 4.4. The possible behaviors of the function $\left(\sum_{v\in V(H)} d(v)^{x-1}\right)^{\frac{1}{x}}$ for $H \neq K_\Delta^{\text{loop}}$.

Finally, we need to argue that for $H \neq K_q^{\text{loop}}$, the edge-critical graphs in $\mathcal{G}(n,1)$ which achieve equality are the only possible graphs in $\mathcal{G}(n,1)$ which achieve equality.

It suffices to consider the addition of a single edge to one of the graphs achieving equality and showing that the number of H-colorings decreases in this case.

By considering the neighbors of a vertex $v \in V(H)$ with $d(v) = \Delta$, adding any edge to a disjoint union of stars strictly lowers the number of H-colorings unless H contains K_Δ^{loop} (a slight modification of the argument given in Case 1 of Section 4.2 will work, realizing that we need to consider both edges joining vertices in the same component and also edges joining vertices in different components). If H does contain K_Δ^{loop} and $H \neq K_\Delta^{\text{loop}}$, then H contains some other component and furthermore H satisfies part 1 of the theorem. Since $\frac{n}{2}K_2$ is the unique edge-critical graph achieving equality for this H and H has at least 2 components, adding any edge to $\frac{n}{2}K_2$ (which will necessarily join together two components of $\frac{n}{2}K_2$) will lower the number of H-colorings in this case as well. This completes the proof of Theorem 4.1.5.

Proof of Lemma 4.3.1: This lemma is a corollary of the following proposition about L^p norms, which is a special case of Lemma 1.11.5 in [63] (or, equivalently, Lemma 2 in Terence Tao's blog post 245C, Notes 1: Interpolation of L^p spaces). Recall that we assume H has no isolated vertices.

Proposition 4.3.2. *Define a measure μ on $V(H)$ by $\mu(v) = \frac{1}{d(v)}$, and let $g : V(H) \to \mathbb{R}$ be given by $g(v) = d(v)$. Then the function defined by $\frac{1}{x} \mapsto ||g||_{L^x(V(H))} = \left(\sum_{v \in V(H)} d(v)^{x-1}\right)^{\frac{1}{x}}$ is log-convex for $x \in (2, \infty)$.*

Recall that a log-convex function can has most one local maximum or local minimum. The composition of the reciprocal map and the map given in Proposition 4.3.2 is the function defined by $x \mapsto \left(\sum_{v \in V(H)} d(v)^{x-1}\right)^{\frac{1}{x}}$. Since the reciprocal is strictly monotone and therefore preserves local extremal values, Lemma 4.3.1 follows.

Lastly, we prove Theorem 4.1.9. Notice that we can still delete any edge from a graph $G \in \mathcal{G}(n, \delta, D)$ and remain in $\mathcal{G}(n, \delta, D)$ as long the edge deletion does not lower the minimum degree. Therefore the proof of Theorem 4.1.5 also proves Theorem 4.1.9

by restricting the function $\left(\sum_{v \in V(H)} d(v)^{x-1}\right)^{\frac{1}{x}}$ to values in $[2, D+1]$.

4.4 Proof of Theorem 4.1.2

Recall that we will assume that that $|V(H)| = q$, and we begin with a few remarks about the number of H-colorings of a cycle C_k. Let A denote the adjacency matrix of H. Then for $k \geq 3$, $\hom(C_k, H) = \operatorname{Tr} A^k$; indeed, the diagonal entry (ii) in A^k counts the number of H-colorings of the path on $k+1$ vertices P_{k+1} that color both endpoints with color i, and by identifying the endpoints we obtain a coloring of C_k with one fixed vertex having color i. Therefore if $\lambda_1, \lambda_2, \ldots, \lambda_q$ are the eigenvalues of A, then

$$\hom(C_k, H) = \lambda_1^k + \cdots + \lambda_q^k. \tag{4.16}$$

It is possible to obtain results using ideas based on Proposition 4.3.2 (with some additional observations); we provide an alternate proof. Without loss of generality, assume that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_q$. Notice that $\lambda_1 > 0$ and $\lambda_1 \geq |\lambda_q|$; this follows from the Perron-Frobenius theorem (see e.g. [61]), but is also readily seen since otherwise (4.16) would imply that $\hom(C_k, H) < 0$ for large odd k.

First we address the inequality, and deal with the cases of equality at the end. Suppose $k \geq 4$ is even and let $b = b(k, H) \geq \lambda_1$ be such that

$$\lambda_1^k + \lambda_2^k + \cdots + \lambda_q^k = b^k.$$

Then

$$\lambda_1^{k+2} + \cdots + \lambda_q^{k+2} \leq \lambda_1^2(\lambda_1^k + \cdots + \lambda_q^k) \leq b^2(\lambda_1^k + \cdots + \lambda_q^k) = b^{k+2},$$

with equality only for $b = \lambda_1$, so

$$(\lambda_1^{k+2} + \cdots + \lambda_q^{k+2})^{\frac{1}{k+2}} \leq (\lambda_1^k + \cdots + \lambda_q^k)^{\frac{1}{k}}, \tag{4.17}$$

which implies that $\hom(C_k, H)^{\frac{1}{k}} \leq \hom(C_4, H)^{\frac{1}{4}}$ for even $k \geq 6$.

Suppose next that $k \geq 5$ is odd, and so as above we have $\lambda_1^{k-1} + \cdots + \lambda_q^{k-1} = b^{k-1}$. Then

$$\lambda_1^k + \cdots + \lambda_q^k \leq |\lambda_1|\lambda_1^{k-1} + \cdots + |\lambda_q|\lambda_q^{k-1} \leq \lambda_1 b^{k-1} \leq b^k, \tag{4.18}$$

with equality only for $b = \lambda_1$, which implies that $\hom(C_k, H)^{\frac{1}{k}} \leq \hom(C_{k-1}, H)^{\frac{1}{k-1}}$.

Summarizing the above, the function $\hom(C_k, H)^{\frac{1}{k}}$ is non-increasing from every even $k \geq 4$ to both $k+1$ and $k+2$ and so, for $k \geq 5$, $\hom(C_k, H)^{\frac{1}{k}} \leq \hom(C_4, H)^{\frac{1}{4}}$. Therefore for all $k \geq 3$,

$$\hom(C_k, H)^{\frac{1}{k}} \leq \max\{\hom(C_3, H)^{\frac{1}{3}}, \hom(C_4, H)^{\frac{1}{4}}\}.$$

Now, if G is any 2-regular graph, then G is the disjoint union of cycles C_{k_i}. So if $\hom(C_4, H)^{\frac{1}{4}} \geq \hom(C_3, H)^{\frac{1}{3}}$,

$$\hom(G, H) = \prod_i \hom(C_{k_i}, H) \leq \prod_i \hom(C_4, H)^{\frac{k_i}{4}} = \hom(C_4, H)^{\frac{n}{4}},$$

with a similar statement holding if $\hom(C_3, H)^{\frac{1}{3}} \geq \hom(C_4, H)^{\frac{1}{4}}$.

Finally, we deal with the cases of equality. There is equality in (4.17) and (4.18) only when $b = \lambda_1$ (and so $\lambda_2 = \cdots = \lambda_q = 0$). Recall that A is symmetric and so has distinct eigenvectors associated to each λ_i, so in this case A has rank 1 and therefore all rows of A are scalar multiples of any other row. If any entry $A_{(ij)} = 0$, then some column and row of A is the 0 vector, which corresponds to an isolated vertex in H. Since we assume H has no isolated vertices, A must be the matrix of all 1's and so $H = K_q^{\text{loop}}$. Therefore, for $H \neq K_q^{\text{loop}}$ we have $\hom(C_k, H)^{\frac{1}{k}} < \max\{\hom(C_3, H)^{\frac{1}{3}}, \hom(C_4, H)^{\frac{1}{4}}\}$ whenever $k \geq 5$. The statement in the theorem about equality is now evident.

4.5 Proof of Theorem 4.1.6 ($\delta = 2$)

4.5.1 Preliminary remarks

We first gather together a number of observations that we'll use in the proof. We will use both Lemma 2.2.1 and Corollary 2.2.2. In fact, Lemma 2.2.1 is enough to prove the case when $\max\{\hom(C_3, H)^{\frac{n}{3}}, \hom(C_4, H)^{\frac{n}{4}}\} \geq \Delta^n$ (without a characterization of uniqueness); we will delay the details of this until Section 4.5.2. A graph H which satisfies

$$\max\{\hom(C_3, H)^{\frac{1}{3}}, \hom(C_4, H)^{\frac{1}{4}}\} < \Delta \tag{4.19}$$

requires a few more observations.

Lemma 4.5.1. *Suppose that the endpoints of P_4 are mapped to H. Then there are at most Δ^2 extensions to an H-coloring of P_4. If H does not contain K_Δ^{loop} or $K_{\Delta,\Delta}$ as a component, then there are strictly fewer than Δ^2 extensions to an H-coloring of P_4.*

Proof. The first statement is obvious, since P_4 is connected and the maximum degree of H is Δ. Suppose there are Δ^2 extensions to an H-coloring of P_4. Let $w_1 \sim_{P_4} w_2 \sim_{P_4} w_3 \sim_{P_4} w_4$ denote the vertices of P_4, and let w_1 and w_4 be given colors v_1 and v_4 in H, respectively. We color w_2 first and then w_3, conditioning on whether v_1 is looped or not.

Suppose that v_1 is unlooped in H. Clearly $d(v_1) = \Delta$ and each neighbor of v_1 also has degree Δ. Since some of the Δ^2 paths from v_1 map w_3 to v_1, it must be the case that $v_4 \sim_H v_1$. Furthermore, if w_2 maps to v_2 (necessarily $v_2 \sim_H v_1$), then every neighbor of v_2 is adjacent to v_4. Since v_2 can be any neighbor of v_1, this implies that $K_{\Delta,\Delta}$ is the component of H containing v_1 and v_4.

A similar analysis for looped v_1 shows that K_Δ^{loop} is the component of H containing v_1 and v_4. □

Corollary 4.5.2. *Suppose that H satisfies (4.19). If $k \geq 4$ and the endpoints of P_k are mapped to H, then there are strictly fewer than Δ^{k-2} extensions to an H-coloring of P_k.*

Proof. Notice that $\hom(C_4, K_{\Delta,\Delta})^{\frac{1}{4}} \geq \Delta$ and $\hom(C_4, K_\Delta^{\text{loop}})^{\frac{1}{4}} = \Delta$. Color, beginning from one endpoint, until there are two uncolored vertices left. Then apply Lemma 4.5.1. □

We can strengthen Corollary 4.5.2 when k is large.

Lemma 4.5.3. *Suppose that H satisfies (4.19). Then there exists a constant l_H (depending on H) such that if $k \geq l_H$ and the endpoints of P_k are mapped to H, then there are strictly fewer than $\frac{1}{|V(H)|^2}\Delta^{k-4}$ extensions to an H-coloring of P_k.*

Proof. Notice that a path must be mapped to a connected component of H; focus on that component. If A is the adjacency matrix of that component, then the number of H-colorings of P_k with endpoints colored i and j is $A^k_{(ij)}$. If λ_1 denotes the largest eigenvalue of A, then by the Perron-Frobenius Theorem (see for example [61, Theorem 1.5]) there exists a strictly positive vector $\mathbf{x}$ such that $A^k\mathbf{x} = \lambda_1^k\mathbf{x}$ for all $k \geq 1$. By considering the row of A containing $\max_{i,j} A^k_{(ij)}$, we see that there is a constant c such that $\max_{i,j} A^k_{(ij)} \leq c\lambda_1^k$ (we can take $c = \max_j x_j / \min_j x_j$, where $\mathbf{x} = (x_j)$). Since $\lambda_1 \leq (\sum_i \lambda_i^4)^{\frac{1}{4}} = \hom(C_4, H)^{\frac{1}{4}} < \Delta$ implies $\lambda < \Delta$, this proves the lemma.

□

4.5.2 The proof

We are now ready to prove Theorem 4.1.6. We assume that G is edge-critical until we discuss the cases of equality in the upper bound. First, suppose that H satisfies

$$\max\{\hom(C_3,H)^{\frac{1}{3}},\hom(C_4,H)^{\frac{1}{4}}\}\geq\Delta. \tag{4.20}$$

Using induction on n, we'll show that for any $G\in\mathcal{G}(n,2)$,

$$\hom(G,H)\leq\max\{\hom(C_3,H)^{\frac{n}{3}},\hom(C_4,H)^{\frac{n}{4}}\}.$$

The base case $n=3$ is trivial.

For the inductive step, assume first that $\hom(C_3,H)^{\frac{1}{3}}\leq\hom(C_4,H)^{\frac{1}{4}}$. If all components of G are cycles, then we're finished by Theorem 4.1.2. If some component of G is not a cycle, then by Lemma 2.2.1 we can partition $V(G)$ into $Y_1\cup Y_2$, with $1\leq|Y_1|\leq n-3$ and Y_1 connected to Y_2. We imagine first coloring Y_2 and then extending this to Y_1. By induction, there are at most $\hom(C_4,H)^{\frac{n-|Y_1|}{4}}$ H-colorings of Y_2. But since Y_1 is connected to Y_2, for every fixed H-coloring of Y_2, each vertex in Y_1 has at most Δ choices for a color. Therefore,

$$\hom(G,H)\leq\Delta^{|Y_1|}\hom(C_4,H)^{\frac{n-|Y_1|}{4}}\leq\hom(C_4,H)^{\frac{n}{4}}. \tag{4.21}$$

The case when $\hom(C_3,H)^{\frac{1}{3}}\geq\hom(C_4,H)^{\frac{1}{4}}$ is similar.

With the upper bound established in this case, we turn to the cases of equality. First suppose that H satisfies $\max\{\hom(C_3,H)^{\frac{1}{3}},\hom(C_4,H)^{\frac{1}{4}}\}>\Delta$. Then (4.21) is strict, which implies that equality can only be obtained for the disjoint union of cycles and hence Theorem 4.1.2 provides the cases of equality among edge-critical graphs.

Now suppose $\max\{\hom(C_3,H)^{\frac{1}{3}},\hom(C_4,H)^{\frac{1}{4}}\}=\Delta$ (which implies that H cannot have $K_{\Delta,\Delta}$ as a component). Notice that equality is achieved for any G with $\hom(G,H) = \Delta^n$, and suppose that $H \neq K_\Delta^{\text{loop}}$ (so since $\hom(C_4,H)^{\frac{1}{4}} \leq \Delta$, H cannot contain K_Δ^{loop} as a component). By Theorem 4.1.2, the construction of G in Corollary 2.2.2 must start with disjoint copies of C_3 and/or C_4. Corollary 4.5.2 implies that only paths of length 1 may be added to these cycles, and to achieve the bound of Δ^n, *every* coloring of these cycles must provide Δ choices for the color of the vertex in the path of length 1. We outline the possible situations which occur when adding a path of length 1 to the cycles in Figure 4.5; the vertex labeled v must have Δ choices for a color regardless of how the adjacent cycles are colored.

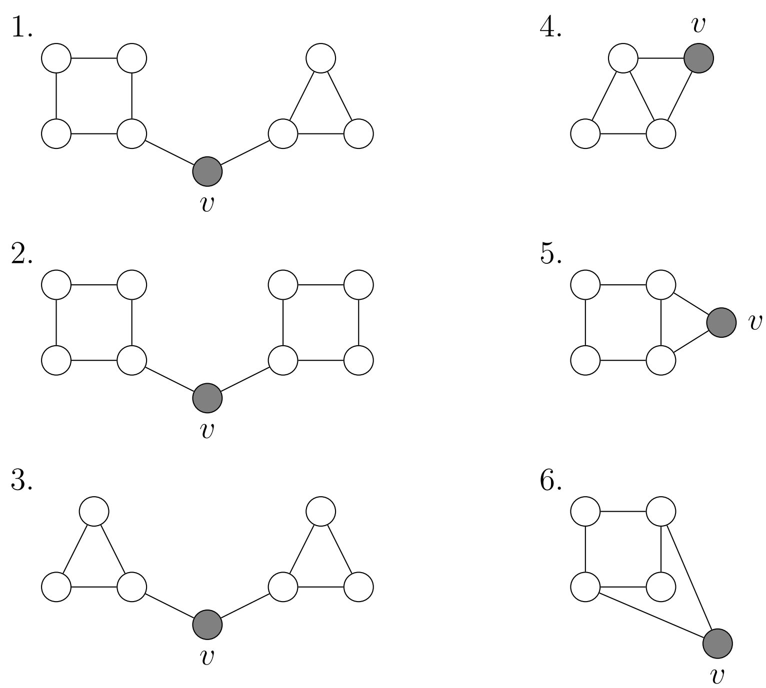

Figure 4.5. The possible situations which occur when a path of length 1 (labeled v) is added to the cycles.

We'll prove that in the situation of Case 5 from Figure 4.5, having Δ choices for the color on v for each coloring of C_4 forces H to contain K_Δ^{loop} as a component. Suppose that the neighbors of v have colors i and j. We can assume that $i \neq j$, since C_4 can map its partition classes to the endpoints of any edge in H. Since v has Δ possibilities for its color, necessarily i and j must each have Δ neighbors and furthermore those Δ neighbors must be simultaneously neighbors of both i and j. In particular, since i and j are on adjacent vertices of C_4, we have $i \sim_H j$ and so i and j must be possible colors for v. This means that i and j must be looped and that all neighbors of i are also neighbors of j. But if k is any other neighbor of i, then a similar argument (replacing j by k) shows that k is looped and is adjacent to all other neighbors of i. Therefore the component containing i is K_Δ^{loop}, which contradicts our assumption that K_Δ^{loop} is not a component of H.

A routine but tedious analysis of the other cases shows that having Δ choices for the color of the vertex in a path of length 1 always forces H to contain K_Δ^{loop} as a component. Therefore, equality can only occur when G is a disjoint union of cycles, so Theorem 4.1.2 again characterizes the cases of equality among edge-critical graphs.

Finally we need to show that edge-critical graphs are the only graphs achieving equality. Arguing as in Case 1 of Section 4.2, adding any edge to a C_4 will strictly lower the number of H-colorings unless H contains K_Δ^{loop}. The cases of adding an edge between two disjoint cycles is similar, and so adding any edge to a graph achieving equality will strictly lower the number of H-colorings unless H is the disjoint union of some number of fully looped complete graphs. If H is of this form and $H \neq K_\Delta^{\text{loop}}$, then $\hom(C_3, H)^{\frac{1}{3}} > \hom(C_4, H)^{\frac{1}{4}}$, and so in fact adding any edge to $\frac{n}{3}C_3$ will strictly lower the number of H-colorings since the disjoint copies of C_3 can be colored using different components of H, but the copies of C_3 joined by an edge must all be colored by a single component of H.

Now suppose that H satisfies

$$\max\{\hom(C_3,H)^{\frac{1}{3}}, \hom(C_4,H)^{\frac{1}{4}}\} < \Delta. \tag{4.22}$$

Recall that from Theorem 4.1.2 we have

$$\hom(C_k,H)^{1/k} \leq \max\{\hom(C_3,H)^{1/3}, \hom(C_4,H)^{1/4}\},$$

which we bound (for simplicity) by

$$\hom(C_k,H) \leq (\Delta^4 - 1)^{k/4} \qquad \text{for } k \geq 3. \tag{4.23}$$

As in the proof of Theorem 4.1.7, we will let $S(2,H)$ denotes the vectors in $V(H)^2$ with the property that the elements of the vector have Δ common neighbors, and $s(2,H) = |S(2,H)|$. Notice that $\hom(K_{2,n-2}) \geq s(2,H)\Delta^{n-2}$.

Suppose again that G is edge-critical. We'll utilize the construction of G from Corollary 2.2.2 to produce all H-colorings of G by coloring the disjoint cycles first and then coloring the paths. If there are more than k vertices in the disjoint cycles, then by (4.23) we have $\hom(G,H) \leq (\Delta^4-1)^{k/4}\Delta^{n-k}$. Therefore, we may assume that there are at most c_1 vertices in disjoint cycles. (All constants in the remainder of this proof will depend on H but will be independent of n.)

After coloring the cycles, we look at the paths that are added iteratively. If any path has length longer than some constant l, then by Lemma 4.5.3 we have $\hom(G,H) < \Delta^{n-2}$. Since a path of length k, for $2 \leq k \leq l$, has at most $\Delta^k - 1$ extensions to an H-coloring by Corollary 4.5.2, if there are more than c_2 such paths then

$$\hom(G,H) < \prod_{i=1}^{c_2} (\Delta^{k_i} - 1)\Delta^{n - \sum_{i=1}^{O(1)} k_i} < \Delta^{n-2}.$$

So, we may assume that the decomposition of G from Corollary 2.2.2 has fewer than c_3 vertices in either disjoint cycles or paths of length $2 \leq k \leq l$, and no paths of length $k > l$. Therefore, the decomposition has at least $n - c_3$ vertices in paths of length 1. Furthermore, each path of length 1 must be attached to two of the at most c_3 vertices composing the disjoint cycles and the paths of length at least 2. By the pigeonhole priciple there exists a $c_4 > 0$ and two vertices in G with at least $c_4 n$ paths of length 1 joining them.

We have shown that every edge-critical graph G which does not have two vertices with at least $c_4 n$ paths of length 1 joining them has $\hom(G, H) < \Delta^{n-2}$, Furthermore, by Lemma 4.5.3 we have the same bound on $\hom(G, H)$ if G has a path of length longer than l. We now deal with the remaining edge-critical graphs G.

Let w_1 and w_2 denote the vertices in G joined by at least $c_4 n$ paths of length 1. Recall from Section 4.2 that $S(2, H)$ is the set of vectors in $V(H)^2$ with the property that the elements of the vector have Δ common neighbors, and $s(2, H) = |S(2, H)|$. Suppose first that the colors on w_1 and w_2 are an element of $S(2, H)$. If G is different from $K_{2,n-2}$, then w_1, w_2, and the at least $c_4 n$ paths of length 1 between them do not form all of G. But then G must contain either a cycle which does not include w_1 or w_2, or a path of length k (for $2 \leq k \leq l_H$) from w_i to w_j for some $i, j \in \{1, 2\}$. By first coloring w_1 and w_2, then any remaining disjoint cycles, and finally the remaining vertices, Corollary 4.5.2 and (4.23) imply that there exists a $c_5 < 1$ such that there are at most

$$s(2, H) c_5 \Delta^{n-2} \tag{4.24}$$

H-colorings of G of this type.

Now suppose that the colors on w_1 and w_2 are not an element of $S(2, H)$. By first coloring w_1 and w_2, then any remaining disjoint cycles, and finally the remaining vertices, we have at most

$$|V(H)|^2\Delta^{n-c_4n-2}(\Delta-1)^{c_4n} \tag{4.25}$$

H-colorings of G of this type.

Combining (4.24) and (4.25) gives

$$\hom(G,H) \le |V(H)|^2\Delta^{n-2}\left(\frac{\Delta-1}{\Delta}\right)^{c_4n} + s(2,H)c_5\Delta^{n-2} < s(2,H)\Delta^{n-2},$$

with the last inequality holding for large enough n.

We've shown that the only edge-critical graph which achieves equality is $K_{2,n-2}$. We repeat an argument given from Case 1 in Section 4.2 to prove that this is the only graph which achieves equality. Since H cannot contain K_Δ^{loop} (by (4.22)), there are two (possibly non-distinct) non-adjacent neighbors of a vertex $v \in V(H)$ with degree Δ. If any edge is added to $K_{2,n-2}$ (necessarily within a partition class), then it is impossible for any H-coloring to color the endpoints of that edge with the non-adjacent vertices, but such a coloring is possible in $K_{2,n-2}$. Therefore adding any edges to $K_{2,n-2}$ produces a graph G with $\hom(G,H) < \hom(K_{2,n-2},H)$.

CHAPTER 5

H-COLORING BIPARTITE GRAPHS

5.1 Introduction and statement of results

Fix graphs G, H, a set of positive weights $\Lambda = \{\lambda_i : i \in V(H)\}$ indexed by the vertices of H, and the associated probability distribution p_Λ (as defined in Section 2.3). The question to be addressed in this chapter is the following. What can be said about an f that is drawn from $\text{Hom}(G, H)$ according to the distribution p_Λ? Specifically, for each $f \in \text{Hom}(G, H)$ and $k \in V(H)$ set

$$s(k, f) = \frac{|f^{-1}(k)|}{|V(G)|},$$

the proportion of vertices receiving color k, and

$$\bar{p}_\Lambda(k) = \frac{1}{|V(G)|} \sum_{v \in V(G)} p_\Lambda(f(v) = k) \ \ (= E_\Lambda\left(s(k, f)\right)).$$

The aim of this chapter is to give fairly precise estimates for $\bar{p}_\Lambda(k)$ and the distribution of $s(k, f)$ for f chosen according to p_Λ, when G is bipartite and either regular or sufficiently close to regular. We will prove the results for regular bipartite G, and will discuss the necessary modifications for close to regular bipartite G in Section 5.4; proofs of these latter results can be found in [20].

The point of departure for this work is a result of Kahn on the hard-core model. When $H = H_{\text{ind}}$ with $V(H_{\text{ind}}) = \{0, 1\}$ and $E(H_{\text{ind}}) = \{00, 01\}$, the set of vertices of G mapped to 1 forms an independent set in G, and $\text{Hom}(G, H_{\text{ind}})$ can be identified with $\mathcal{I}(G)$, the set of independent sets in G. For each $\lambda > 0$, the *hard-core* model on G is the probability distribution $\mathrm{hc}(\lambda)$ on $\mathcal{I}(G)$ that assigns to each $I \in \mathcal{I}(G)$ a probability proportional to $\lambda^{|I|}$. One of the oldest and most studied spin models in statistical physics, this is a simple mathematical model of the occupation of space (represented by G) by particles of non-negligible size. The model can easily be realized as a spin model with distribution p_Λ given by assigning weights $\lambda_0 = 1$ and $\lambda_1 = \lambda$ to the vertices of H_{ind}. See Figure 5.1.

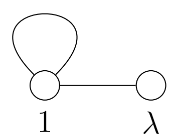

Figure 5.1. The graph H_{ind} with weighting $\lambda_0 = 1$ and $\lambda_1 = \lambda$.

Kahn [39] studied this model on a regular bipartite graph G. He proved that for all fixed $\lambda > 0$, the model exhibits a *phase coexistence* in the sense that if G has equipartition $\mathcal{E} \cup \mathcal{O}$ then most $\mathrm{hc}(\lambda)$ independent sets tend to come either mostly from $\mathcal{E}$ or mostly from $\mathcal{O}$, in the sense that the size of an independent set chosen according to $\mathrm{hc}(\lambda)$ is concentrated close to $\lambda/(2(1+\lambda))$, which is exactly the expected size of an independent set chosen according to the distribution that half the time picks a $\mathrm{hc}(\lambda)$ independent set from $\mathcal{E}$ and half the time picks from $\mathcal{O}$. The following theorem ([39, Theorem 1.4 & Corollary 1.5]) formalizes this.

Theorem 5.1.1 (Kahn, 2001 [39]). *Let $\lambda > 0$ be fixed. There are positive constants c_1, c_2, c_3 and c_4 (depending on λ) such that for every d-regular bipartite graph G on*

n vertices, the following two statements hold. Firstly, for every $\varepsilon \geq c_1/\sqrt{d}$, if I is chosen from $\mathcal{I}(G)$ according to the distribution $\mathrm{hc}(\lambda)$ *then*

$$\Pr\left(\left||I| - \frac{\lambda n}{2(1+\lambda)}\right| \geq \varepsilon n\right) \leq c_2 \varepsilon^{-1} 2^{-c_3\varepsilon^2 n}.$$

Secondly,

$$\left|\frac{E(|I|)}{n} - \frac{\lambda}{2(1+\lambda)}\right| \leq c_4 \zeta$$

where

$$\zeta = \max\left\{\frac{1}{\sqrt{d}}, \sqrt{\frac{\log n}{n}}\right\}. \tag{5.1}$$

In particular, a uniformly chosen independent set ($\lambda = 1$) from a regular bipartite graph consists, with high probability, of close to one quarter of the vertices. While this corollary may seem more natural than the formulation of Theorem 5.1.1, it is worth noting that in order to prove the theorem in the special case of $\lambda = 1$ it is necessary (at least using the entropy methods of [39]) to pass to the more general weighted model first. Similarly, it might seem more natural in the present chapter to focus on the structure of uniform H-colorings, but we are unable to obtain any results without introducing weights.

From (5.1) we see that Theorem 5.1.1 only gives a concentration result when we consider families of graphs with d going to infinity. This is not just an artifact of the proof. For families of graphs with d fixed (and only n going to infinity), the behavior of $E(|I|)/n$ depends very much on the particular choice of family. As an example, consider the case $d = 2$. If G_n is the disjoint union of $n/4$ copies of the cycle C_4, and I is chosen uniformly from $\mathcal{I}(G)$, then $E(|I|)/n$ is easily seen to be concentrated close to $2/7$. If, however, G_n is the disjoint union of $n/6$ copies of the cycle C_6, then $E(|I|)/n$ is concentrated close to $5/18$. For this reason we implicitly assume throughout that d is going to infinity.

We now set up some notation that allows us to state our main result, which is an extension of Theorem 5.1.1 to arbitrary weighted H-colorings. This notation established in Section 2.3 will be used here and it is a good idea to review that section before continuing on with the remainder of this chapter. From the remainder of this chapter, whenever H and Λ are mentioned, it will be assumed that H is a finite graph without multiple edges but perhaps with loops, and that Λ is a set of positive weights indexed by the vertices of H. Recall that for $A, B \subseteq V(H)$ we write $A \sim B$ if for all $u \in A$ and $v \in B$ we have $u \sim_H v$, and set

$$\eta_\Lambda(H) = \max\{\lambda_A \lambda_B : A \sim B\}$$

where $\lambda_\bullet = \sum_{i \in \bullet} \lambda_i$. (We will abuse notation and write λ_H for $\lambda_{V(H)}$.) Also recall that

$$\mathcal{M}_\Lambda(H) = \{(A, B) \in V(H)^2 : A \sim B,\ \lambda_A \lambda_B = \eta_\Lambda(H)\}.$$

Next define

$$a_\Lambda^+(k) = \frac{\max\left\{\lambda_A \lambda_k \mathbf{1}_{\{k \in B\}} + \lambda_B \lambda_k \mathbf{1}_{\{k \in A\}} : (A, B) \in \mathcal{M}_\Lambda(H)\right\}}{2\eta_\Lambda(H)}$$

and define $a_\Lambda^-(k)$ similarly, with max replaced by min. (After the statement of Theorem 5.1.4, we will give some explicit examples to illuminate these definitions.) We make a few remarks regarding the definitions of $a_\Lambda^+(k)$ and $a_\Lambda^-(k)$. Note that if k does not appear in any $(A, B) \in \mathcal{M}_\Lambda(H)$ then $a_\Lambda^+(k) = 0$ and that if there is a pair $(A, B) \in \mathcal{M}_\Lambda(H)$ in which k does not appear then $a_\Lambda^-(k) = 0$. Note also that $a_\Lambda^-(k) \leq a_\Lambda^+(k)$. Finally, note that $a_\Lambda^+(k)$ and $a_\Lambda^-(k)$ both take the form

$$\frac{\lambda_k \mathbf{1}_{\{k \in A\}}}{2\lambda_A} + \frac{\lambda_k \mathbf{1}_{\{k \in B\}}}{2\lambda_B}$$

for some $(A, B) \in \mathcal{M}_\Lambda(H)$. We may interpret this quantity as the expected proportion of vertices mapped to k in a pure-(A,B) coloring chosen according to p_Λ, i.e., a p_Λ-chosen H-coloring subject to the condition that all vertices from one partition class of G get mapped to A and all from the other class get mapped to B.

Finally, for every $\varepsilon > 0$ and $k \in V(H)$ define

$$I_k(\varepsilon) = [0, a_\Lambda^-(k) - \varepsilon) \cup (a_\Lambda^+(k) + \varepsilon, 1].$$

Before stating our main result, we motivate it by considering weighted H-colorings of $K_{d,d}$, the complete bipartite graph with d vertices in each partition class, for some fixed H and Λ. The adjacency structure of $K_{d,d}$ ensures that all H-colorings are pure-(A,B) for some (A,B) with $A \sim B$, and that moreover all but a vanishing proportion (in d) of $Z_\Lambda(K_{d,d}, H)$ comes from pure-(A,B) colorings for some $(A,B) \in \mathcal{M}_\Lambda(H)$. It follows that for each $k \in V(H)$, in an H-coloring chosen according to p_Λ we have that with probability $1 - o(1)$ the proportion of vertices of $K_{d,d}$ mapped to k will be between $a_\Lambda^-(k) - o(1)$ and $a_\Lambda^+(k) + o(1)$. Our main result, which we now state, asserts that this property of $K_{d,d}$ is essentially shared by all d-regular graphs.

Theorem 5.1.2. *Fix H and Λ. There are positive constants c_1, c_2, c_3 and c_4 (depending on H and Λ) such that for every d-regular bipartite graph G on n vertices, the following two statements hold. Firstly, for every $\varepsilon \geq c_1/\sqrt{d}$ and $k \in V(H)$ we have*

$$p_\Lambda\left(s(k,f) \in I_k(\varepsilon)\right) \leq c_2 \varepsilon^{-1} 2^{-c_3 \varepsilon^2 n}. \tag{5.2}$$

Secondly, for each $k \in V(H)$ we have

$$\bar{p}_\Lambda(k) \in \left[a_\Lambda^-(k) - c_4\zeta, a_\Lambda^+(k) + c_4\zeta\right] \tag{5.3}$$

where ζ is as defined in (5.1).

In other words, for regular bipartite G the distribution p_Λ is concentrated on H-

colorings for which, for every $k \in V(H)$, the proportion of vertices mapped to k is roughly between $a_\Lambda^-(k)$ and $a_\Lambda^+(k)$.

We prove Theorem 5.1.2 in Section 5.2; the proof goes along the following lines. We upper bound the contribution to $Z_\Lambda(G,H)$ from those $f \in \mathrm{Hom}(G,H)$ with $|f^{-1}(k)|/n = \gamma \geq a^+(k) + \varepsilon$ by $Z_{\Lambda(k,\delta)}(G,H)/(1+\delta)^{\gamma n}$ for some suitably small $\delta > 0$ (where $\Lambda(k,\delta)$ is obtained from Λ by multiplying λ_k by $1+\delta$ and leaving all other λ_i unchanged). We in turn upper bound $Z_{\Lambda(k,\delta)}(G,H)$ using a result of Galvin and Tetali [32] to the effect that for all H and Λ and all d-regular bipartite graphs G on n vertices we have

$$Z_\Lambda(G,H) \leq Z_\Lambda(K_{d,d},H)^{\frac{n}{2d}} \tag{5.4}$$

(where recall $K_{d,d}$ is the complete bipartite graph with d vertices in each partition class). We upper bound $Z_{\Lambda(k,\delta)}(K_{d,d},H)$ in terms of $\eta_{\Lambda(k,\delta)}(H)$, and in the end we get, using our choice of $a_\Lambda^+(k)$ and for some sufficiently small δ, an upper bound on the contribution that is significantly smaller than a trivial lower bound on $Z_\Lambda(G,H)$, showing that those $f \in \mathrm{Hom}(G,H)$ with $|f^{-1}(k)|/n \geq a^+(k) + \varepsilon$ do not contribute greatly to the partition function. The same strategy works for $|f^{-1}(k)|/n$ falling significantly below $a^-(k)$. The details are given in Section 5.2.

When $a_\Lambda^-(k) = a_\Lambda^+(k)$ for all k, we obtain a single vector around which $(s(k,f) : k \in V(H))$ is concentrated for f chosen according to p_Λ.

Corollary 5.1.3. *Fix H and Λ. Suppose that for all $k \in V(H)$ there is an $a_\Lambda(k)$ such that $a_\Lambda^-(k) = a_\Lambda^+(k) = a_\Lambda(k)$. Then there are positive constants c_1, c_2, c_3 and c_4 (depending on H and Λ) such that for every d-regular, bipartite graph G on n vertices the following two statements hold. Firstly, for $\varepsilon \geq c_1/\sqrt{d}$ we have*

$$p_\Lambda\left(\left|\left|(s(k,f))_{k\in V(H)} - (a_\Lambda(k))_{k\in V(H)}\right|\right|_\infty \geq \varepsilon\right) \leq c_2\varepsilon^{-1}2^{-c_3\varepsilon^2 n}.$$

Secondly, we have

$$\left|\left|(\bar{p}_\Lambda(k))_{k\in V(H)} - (a_\Lambda(k))_{k\in V(H)}\right|\right|_\infty \leq c_4\zeta$$

with ζ as in (5.1).

A situation in which Corollary 5.1.3 applies is when either $\mathcal{M}_\Lambda(H) = \{(A, A)\}$ or $\mathcal{M}_\Lambda(H) = \{(A, B), (B, A)\}$ (for some $A \neq B$). This is in a sense the generic situation. Indeed, for every H, if the weights λ_i are chosen from any continuous distribution supported on $\{x \in \mathbb{R}^{|V(H)|} : x > 0\}$, then with probability 1 we will have $\mathcal{M}_\Lambda(H)$ of the form described. As we will see in Example C below, Corollary 5.1.3 also applies in some other natural situations.

The gap between $a^-_\Lambda(k)$ and $a^+_\Lambda(k)$ (if there is one) cannot be closed in general, as the first part of the following theorem shows.

Theorem 5.1.4. *Fix H and Λ. There is a family $\{G_d\}_{d=1}^\infty$ of d-regular bipartite graphs, a function $g(d) = o(1)$ and a positive constant c (depending on H and Λ) such that for each $k \in V(H)$,*

$$\left.\begin{array}{l} p_\Lambda\left(\left|s(k,f) - a^+_\Lambda(k)\right| \leq g(d)\right) \\ p_\Lambda\left(\left|s(k,f) - a^-_\Lambda(k)\right| \leq g(d)\right) \end{array}\right\} \geq c - g(d).$$

There is also a family $\{G'_d\}_{d=1}^\infty$ of d-regular bipartite graphs, a function $g(d) = o(1)$ and (for each $k \in V(H)$) an $a_\Lambda(k)$ satisfying $a^-_\Lambda(k) \leq a_\Lambda(k) \leq a^+_\Lambda(k)$ such that for each k,

$$p_\Lambda\left(|s(k,f) - a_\Lambda(k)| \leq g(d)\right) \geq 1 - g(d)$$

and

$$|\bar{p}_\Lambda(k) - a_\Lambda(k)| \leq g(d).$$

We prove Theorem 5.1.4 in Section 5.3. The graphs G_d we exhibit will be suitably chosen random regular graphs, and we will use the expansion of these graphs to show that all but $o(1)$ of p_Λ is concentrated on pure-(A,B) colorings for $(A,B) \in \mathcal{M}_\Lambda(H)$. The graphs G'_d will be disjoint unions of complete bipartite graphs on $2d$ vertices. Basic concentration estimates together with the independence of the components will give the claimed result.

We now explore the consequences of Theorem 5.1.2 for some specific choices of H and Λ.

Example A (Hard-core model) Let $H = H_{\text{ind}}$ be as described earlier, with $\lambda_0 = 1$ and $\lambda_1 = \lambda$. We have seen that an element of $\text{Hom}(G, H_{\text{ind}})$ chosen according to p_Λ is a configuration in the hard-core model on G with activity λ. With these choices we have $\mathcal{M}_\Lambda(H_{\text{ind}}) = \{(\{0\},\{0,1\}),(\{0,1\},\{0\})\}$ and

$$a^-_\Lambda(1) = a^+_\Lambda(1) = \frac{\lambda}{2(1+\lambda)}$$

and so Theorem 5.1.2 indeed generalizes Theorem 5.1.1, as claimed.

Example B (Multistate hard-core model) Let $H = H_k$ be the graph on vertex set $\{0,\dots,k\}$ with $i \sim_{H_k} j$ if and only if $i + j \le k$, and $\lambda_i = \lambda^i$ for some fixed $\lambda > 0$. An element of $\text{Hom}(G, H_k)$ chosen according to p_Λ is exactly a configuration of the *multistate hard-core* (or *multicast communications*) model on G with activity λ. This model allows multiple particles (up to and including k) at each site, with the restriction that there are no more than k particles in total across each edge. A generalization of the hard-core model (the case $k = 1$), it has been studied in a variety of contexts: in communications [57], statistical physics [51] and combinatorics [31]. For k even the unique pair $(A,A) \in \mathcal{M}_\Lambda(H_k)$ has $A = \{1,\dots,k/2\}$, while for k odd, say $k = 2\ell+1$, we have $\mathcal{M}_\Lambda(H_k) = \{(A,B),(B,A)\}$ with $A = \{1,\dots,\ell\}$ and $B = \{1,\dots,\ell+1\}$. In either case Corollary 5.1.3 shows that for this model $(s(k,f) : k \in V(H))$ is concentrated close to a single value for f chosen according to

p_Λ.

Example C (Uniform proper q-colorings) Let $H = K_q$, the complete graph on q vertices, and $\Lambda = (1, \dots, 1)$. An element of $\mathrm{Hom}(G, K_q)$ chosen according to p_Λ corresponds to a uniform proper q-coloring of G. In this case elements of $\mathcal{M}_\Lambda(K_q)$ consist of all partitions of $V(K_q)$ into two classes as near equal in size as possible, and an easy calculation gives that for all colors k

$$a_\Lambda^-(k) = \frac{1}{2\lceil q/2 \rceil} \quad \text{and} \quad a_\Lambda^+(k) = \frac{1}{2\lfloor q/2 \rfloor}$$

so that in particular $a_\Lambda^-(k) = a_\Lambda^+(k) = 1/q$ for q even, and we get the following corollary of Theorem 5.1.2.

Corollary 5.1.5. *Fix $q \in \mathbb{N}$. There are positive constants c_1, c_2 and c_3 (depending on q) such that for every d-regular, bipartite graph G on n vertices, the following statements hold. If χ is a uniformly chosen q-coloring of G and $\varepsilon \geq c_1/\sqrt{d}$ then for q even*

$$\Pr\left(\exists k \in V(H) : \left|\frac{|\chi^{-1}(k)|}{n} - \frac{1}{q}\right| \geq \varepsilon\right) \leq c_2 \varepsilon^{-1} 2^{-c_3 \varepsilon^2 n},$$

and for q odd

$$\left.\begin{array}{l} \Pr\left(\exists k \in V(H) : \frac{|\chi^{-1}(k)|}{n} \leq \frac{1}{q+1} - \varepsilon\right) \\ \Pr\left(\exists k \in V(H) : \frac{|\chi^{-1}(k)|}{n} \geq \frac{1}{q-1} + \varepsilon\right) \end{array}\right\} < c_2 \varepsilon^{-1} 2^{-c_3 \varepsilon^2 n}.$$

Recall that a coloring is said to be *equitable* if the number of vertices in any two color classes differ by at most 1; that is, the partition of vertices among the color classes is as uniform as possible. So for even q, Corollary 5.1.5 states that almost all proper q-colorings of a regular bipartite graph are "almost equitable". Of course, by the symmetry of K_q we have $E(|\chi^{-1}(k)|) = n/q$ for all k in this case.

5.2 Proof of Theorem 5.1.2

We use (5.4) to upper bound $Z_\Lambda(G,H)$. For each of the at most $4^{|V(H)|}$ ordered pairs $A \sim B$ of subsets of H, the contribution to $Z_\Lambda(K_{d,d},H)$ from those f with one partition class of $K_{d,d}$ mapped onto A and the other partition class mapped onto B is at most

$$(\lambda_A \lambda_B)^d$$

and so

$$Z_\Lambda(G,H) \leq Z_\Lambda(K_{d,d},H)^{\frac{n}{2d}} \leq \eta_\Lambda(H)^{\frac{n}{2}} 4^{\frac{n|V(H)|}{2d}} = \eta_\Lambda(H)^{\frac{n}{2}} C^{\frac{n}{d}}, \tag{5.5}$$

where C is a positive constant depending only on H.

On the other hand, we get a lower bound (with any $\lambda_A\lambda_B = \eta_\Lambda(H)$) of

$$Z_\Lambda(G,H) \geq \eta_\Lambda(H)^{\frac{n}{2}}. \tag{5.6}$$

We now use (5.5) and (5.6) to prove (5.2). Fix $k \in V(H)$ and an integer n_k satisfying $0 \leq n_k \leq n$ and

$$\frac{n_k}{n} \in [0, a_\Lambda^-(k) - \varepsilon) \cup (a_\Lambda^+(k) + \varepsilon, 1] \ \ (= I_k(\varepsilon))\,.$$

Write $c_k(n_k)$ for the contribution to $Z_\Lambda(G,H)$ from those $f \in \mathrm{Hom}(G,H)$ with $|f^{-1}(k)| = n_k$. We aim to obtain an upper bound on $c_k(n_k)$ (via (5.5)) which is substantially lower than the lower bound (5.6), indicating that this term does not contribute greatly to $Z_\Lambda(G,H)$.

We begin by considering n_k for which

$$\gamma := \frac{n_k}{n} = a_\Lambda^+(k) + \varepsilon'$$

for some ε' satisfying $\varepsilon \leq \varepsilon' \leq 1 - a_\Lambda^+(k)$. For any $\delta > 0$ let $\Lambda(k,\delta)$ be obtained from

Λ by replacing λ_k with $(1+\delta)\lambda_k$ and leaving all other λ_i's unchanged (and let $\lambda_{(k,\delta),A}$ denote $\sum_{j\in A}\lambda_j$ with λ_j in $\Lambda(k,\delta)$). By (5.5) we have

$$\begin{aligned}(1+\delta)^{n_k}c_k(n_k) &\leq Z_{\Lambda(k,\delta)}(G,H) \\ &\leq \eta_{\Lambda(k,\delta)}(H)^{\frac{n}{2}}C^{\frac{n}{d}}. \end{aligned} \tag{5.7}$$

Before proceeding, we need to understand $\eta_{\Lambda(k,\delta)}(H)$. Viewed as a function of δ, the quantity $\lambda_{(k,\delta),A}\lambda_{(k,\delta),B}$ (for $(A,B)\in\mathcal{M}_\Lambda(H)$) is of the form $a+b\delta+c\delta^2$ where $a=\eta_\Lambda(H)$, $b=\lambda_A\lambda_k\mathbf{1}_{\{k\in B\}}+\lambda_B\lambda_k\mathbf{1}_{\{k\in A\}}$ and $c=\lambda_k^2\mathbf{1}_{\{k\in A\cap B\}}$. From this formulation we can easily identify that set $\emptyset\neq\mathcal{S}^+_\Lambda(k,H)\subseteq\mathcal{M}_\Lambda(H)$ with the property that for all $\delta>0$, all $(A,B)\in\mathcal{M}_\Lambda(H)$ and all $(A',B')\in\mathcal{S}^+_\Lambda(k,H)$ we have $\lambda_{(k,\delta),A'}\lambda_{(k,\delta),B'}\geq\lambda_{(k,\delta),A}\lambda_{(k,\delta),B}$: $\mathcal{S}^+_\Lambda(k,H)$ consists of all those $(A',B')\in\mathcal{M}_\Lambda(H)$ for which b is maximum and (subject to this condition) c is maximum. This latter condition simply means that if some of the pairs that maximize b have $c>0$ we only take those pairs, and if they all have $c=0$ we take all pairs.

It is easily seen that there is a sufficiently small $\delta^+_k>0$ (depending on H and Λ) with the property that for all $0<\delta<\delta^+_k$ and $(A',B')\in\mathcal{S}^+_\Lambda(k,H)$ we have $(A',B')\in\mathcal{M}_{\Lambda(k,\delta)}(H)$. Choose one such, (A^+,B^+), arbitrarily. Note that by construction

$$a^+_\Lambda(k)=\frac{(\lambda_{A^+})\lambda_k\mathbf{1}_{\{k\in B^+\}}+(\lambda_{B^+})\lambda_k\mathbf{1}_{\{k\in A^+\}}}{2\eta_\Lambda(H)}=\frac{\lambda_k\mathbf{1}_{\{k\in A^+\}}}{2\lambda_{A^+}}+\frac{\lambda_k\mathbf{1}_{\{k\in B^+\}}}{2\lambda_{B^+}}.$$

Now combining (5.6) and (5.7) and choosing $\delta<\delta^+_k$ we have

$$\begin{aligned}p_\Lambda(|f^{-1}(k)|=n_k) &= \frac{c_k(n_k)}{Z_\Lambda(G,H)} \\ &\leq C^{\frac{n}{d}}\left(\frac{(\lambda_{(k,\delta),A^+})(\lambda_{(k,\delta),B^+})}{(\lambda_{A^+})(\lambda_{B^+})(1+\delta)^{2(a^+_\Lambda(k)+\varepsilon')}}\right)^{\frac{n}{2}}. \end{aligned} \tag{5.8}$$

Our aim is to show that there is a positive constant c (depending on H and Λ) such that for all $0 < \varepsilon' \leq 1 - a_\Lambda^+(k)$ we can find a $0 < \delta < \delta_k^+$ for which

$$\frac{(\lambda_{(k,\delta),A^+})(\lambda_{(k,\delta),B^+})}{(\lambda_{A^+})(\lambda_{B^+})(1+\delta)^{2(a_\Lambda^+(k)+\varepsilon')}} \leq 2^{-c\varepsilon'^2}. \tag{5.9}$$

Combining this with (5.8) we see that if $\varepsilon > c/\sqrt{d}$ for some suitably large positive constant c (depending on Λ and H) then for all $\varepsilon < \varepsilon' \leq 1 - a_\Lambda^+(k)$ for which $a^+(k)n + \varepsilon' n$ is an integer we have

$$p_\Lambda\left(|f^{-1}(k)| = a^+(k)n + \varepsilon' n\right) \leq 2^{-c'\varepsilon'^2 n}$$

for a suitable positive c', and so

$$\begin{aligned} p_\Lambda\left(|f^{-1}(k)| \geq a^+(k)n + \varepsilon n\right) &\leq \sum_{\ell \geq \varepsilon n} 2^{-\frac{c'\ell^2}{n}} \\ &\leq 2^{-c'\varepsilon^2 n} \sum_{\ell \geq 0} 2^{-2\ell c'\varepsilon} \\ &\leq c''\varepsilon^{-1} 2^{-c'\varepsilon^2 n} \end{aligned} \tag{5.10}$$

for suitably large c'' (depending on c'). An almost identical argument (the details of which we leave to the reader) yields

$$p_\Lambda\left(|f^{-1}(k)| \leq a^-(k)n - \varepsilon n\right) \leq c''\varepsilon^{-1} 2^{-c'\varepsilon^2 n} \tag{5.11}$$

for $\varepsilon > c/\sqrt{d}$. Combining (5.10) and (5.11) gives (5.2).

We now turn to (5.9). Observe that it is enough to prove (5.9) for all $0 < \varepsilon' \leq \varepsilon_0$, where $\varepsilon_0 \leq 1 - a_\Lambda^+(k)$ may be any constant (perhaps depending on H and Λ). Indeed, for any $\varepsilon' \geq \varepsilon_0$ we know that there is a choice of $\delta < \delta_k^+$ for which

$$\begin{aligned} \frac{(\lambda_{(k,\delta)A^+})(\lambda_{(k,\delta),B^+})}{(\lambda_{A^+})(\lambda_{B^+})(1+\delta)^{2(a_\Lambda^+(k)+\varepsilon')}} &\leq \frac{(\lambda_{(k,\delta),A^+})(\lambda_{(k,\delta)B^+})}{(\lambda_{A^+})(\lambda_{B^+})(1+\delta)^{2(a_\Lambda^+(k)+\varepsilon_0)}} \\ &\leq 2^{-c\varepsilon_0^2}. \end{aligned}$$

Setting $c' = c\varepsilon_0^2$ we have $2^{-c\varepsilon_0^2} \le 2^{-c'\varepsilon'^2}$ for $\varepsilon' \ge \varepsilon_0$ and $2^{-c\varepsilon'^2} \le 2^{-c'\varepsilon'^2}$ for $\varepsilon' < \varepsilon_0$, so we may replace c with c' to obtain the result for the full range of ε'. From now on we will assume that $\varepsilon' < \varepsilon_0$, for a certain ε_0 that will be specified later.

Setting

$$\gamma_A = \frac{\lambda_k \mathbf{1}_{\{k\in A^+\}}}{2\lambda_{A^+}}, \quad \gamma_B = \frac{\lambda_k \mathbf{1}_{\{k\in B^+\}}}{2\lambda_{B^+}}$$

(so $a_\Lambda^+(k) = \gamma_A + \gamma_B$) the left-hand side of (5.9) becomes

$$\frac{(\lambda_{A^+}) + \delta\lambda_k \mathbf{1}_{\{k\in A^+\}}}{(1+\delta)^{2\gamma_A+\varepsilon'}(\lambda_{A^+})} \times \frac{(\lambda_{B^+}) + \delta\lambda_k \mathbf{1}_{\{k\in B^+\}}}{(1+\delta)^{2\gamma_B+\varepsilon'}(\lambda_{B^+})}. \tag{5.12}$$

If either $A^+ = \{k\}$ or $k \notin A^+$ then the first term of (5.12) is $(1+\delta)^{-\varepsilon'}$ so that in this case we have that for any $0 < \varepsilon' \le 1$ there is a small enough $\delta > 0$ with

$$\frac{(\lambda_{A^+}) + \delta\lambda_k \mathbf{1}_{\{k\in A^+\}}}{(1+\delta)^{2\gamma_A+\varepsilon'}(\lambda_{A^+})} \le 2^{-c\varepsilon'} \le 2^{-c\varepsilon'^2},$$

where c is a positive constant depending on H and Λ (the last inequality using $\varepsilon' \le 1$). If $k \in A^+$ and $|A^+| > 1$ then the first term of (5.12) takes the form

$$\begin{aligned}
\frac{(\lambda_{A^+}) + \delta\lambda_k}{(1+\delta)^{2\gamma_A+\varepsilon'}(\lambda_{A^+})} &\le \frac{1+\delta(\lambda_k/\lambda_{A^+})}{1+\delta(2\gamma_A+\varepsilon')} \\
&= 1 - \frac{\delta\varepsilon'}{1+\delta(\lambda_k/(\lambda_{A^+})+\varepsilon')} \\
&\le 1 - \frac{\delta\varepsilon'}{3},
\end{aligned} \tag{5.13}$$

with (5.13) valid for sufficiently small ε'. Now taking $\delta = \varepsilon'$ (having chosen ε_0 small enough that this choice is allowed, and that (5.13) holds), we get a bound of $2^{-c\varepsilon'^2}$ on the first term of (5.12), where c is a positive constant depending on H and Λ only.

Repeating this analysis for the second term of (5.12), we obtain (5.9) and thus (5.2).

Applying (5.2) with $\varepsilon = c\sqrt{(\log n)/n}$ (if $(\log n)/n > 1/d$) and $\varepsilon = c/\sqrt{d}$ (otherwise), where $c \geq c_1$ satisfies $c^2 c_3 \geq 1$, we easily obtain (5.3), based on the observation that in both cases

$$E_\Lambda(s(k,f)) \leq (a_\Lambda^+(k) + \varepsilon)\left(1 - c_2\varepsilon^{-1}2^{-c_3\varepsilon^2 n}\right) + c_2\varepsilon^{-1}2^{-c_3\varepsilon^2 n}$$

with a similar lower bound involving $a_\Lambda^-(k)$.

5.3 Proof of Theorem 5.1.4

The graph G_d will be a random d-regular bipartite graph on $n = c^{d/\log d}$ vertices (where $c > 1$ will depend on the particular H and Λ under consideration). A standard method of constructing such a graph is as follows. We begin with a set of size nd consisting of $nd/2$ type I vertices $\{u_{ij} : 1 \leq i \leq n/2, 1 \leq j \leq d\}$ and $nd/2$ type II vertices $\{v_{ij} : 1 \leq i \leq n/2, 1 \leq j \leq d\}$. We then choose a uniformly random perfect matching from the type I vertices to the type II vertices, and turn this into a d-regular bipartite multigraph on n vertices with bipartition classes $\mathcal{E} = \{u_1, \ldots, u_{n/2}\}, \mathcal{O} = \{v_1, \ldots, v_{n/2}\}$ by, for each $i = 1, \ldots, n/2$, identifying $u_{i,1}, \ldots, u_{i,d}$ with u_i and $v_{i,1}, \ldots, v_{i,d}$ with v_i. Finally, we condition on the result being a simple graph. This process generates a d-regular bipartite graph on n vertices with bipartition classes $\mathcal{E}$, $\mathcal{O}$, uniformly (see for example [66]).

O'Neil [53] showed that the probability that the multigraph produced by this process is simple is (for large enough d) at least $e^{-d^2/3}$. It follows that if we establish that the multigraph produced (before conditioning on being simple) has a certain property with probability at least $1 - e^{-d^2}$ (say), then there is a simple d-regular graph with that property.

We want to establish that for large enough d the multigraph has a number of desirable expansion properties. First, we want to show that for each $C \log d \leq j \leq 3n(\log d)/d$ (for some constant $C > 0$, depending on c), every subset of $\mathcal{E}$ of size j and every subset of $\mathcal{O}$ of size j has at least αj distinct neighbors where $\alpha = d/(C \log d)$. For a particular such j, the probability that the graph fails to have this property is (by a union bound) at most

$$\begin{aligned}
2\binom{n/2}{j}\binom{n/2}{\alpha j}\frac{(\alpha j d)^{\underline{jd}}}{(nd/2)^{\underline{jd}}} &\leq \left(\frac{en}{2\alpha j}\right)^{2\alpha j}\left(\frac{2\alpha j}{n}\right)^{jd} \\
&= e^{\frac{2jd}{C\log d}}\left(\frac{2jd}{Cn\log d}\right)^{jd-\frac{2jd}{C\log d}} \\
&\leq e^{\frac{2jd}{C\log d}}\left(\frac{2jd}{Cn\log d}\right)^{jd/2}
\end{aligned}$$

(for large enough d, depending on C) with the first inequality using $\binom{n}{r} \leq (en/r)^r$. For $j \geq d \log d$ we bound $2jd/(Cn \log d) \leq 1/2$ (valid for $C \geq 12$) so that for large enough d (depending on C)

$$e^{\frac{2jd}{C\log d}}\left(\frac{2jd}{Cn\log d}\right)^{jd/2} \leq 1.4^{-jd} \leq e^{-2d^2}.$$

For $j \leq d \log d$ we instead bound $(2dj)/(Cn \log d) \leq d^2/n$ (valid for $C \geq 2$). We now have

$$e^{\frac{2jd}{C\log d}}\left(\frac{2jd}{Cn\log d}\right)^{jd/2} \leq \exp\left\{2jd\log d - \frac{jd^2\log c}{2\log d}\right\} \leq \exp\left\{\frac{-jd^2\log c}{3\log d}\right\}$$

(again for large d, recalling $n = c^{d/\log d}$), which is at most e^{-2d^2} for $j \geq C \log d$ for suitable C depending on c. Since there are at most $n = c^{d/\log d}$ choices for j, the probability that the graph fails to have the desired property for some j is at most e^{-d^2}. If the process results in a simple graph, then we trivially get the same expansion for subsets of $\mathcal{E}$ or $\mathcal{O}$ of size at most $C \log d$, since for $1 \leq j \leq C \log d$ there is a trivial lower bound of d on the neighborhood size of a set of size j, and we have $d \geq jd/(C \log d)$ for j in this range.

Next we establish that the graph has the property that for every subset A of $\mathcal{E}$ of size $3n(\log d)/d$ and every subset B of $\mathcal{O}$ of size $3n(\log d)/d$, there is an edge joining a vertex of A to a vertex of B. By a union bound, the probability that the multigraph fails to have the property is at most

$$\binom{n/2}{\beta n}^2 \frac{(nd/2 - \beta nd)^{\underline{\beta nd}}}{(nd/2)^{\underline{\beta nd}}} \leq \exp\left\{2\beta n \log(e/(2\beta)) - 2\beta^2 dn\right\}$$

where $\beta = 3(\log d)/d$. With $n = c^{d/\log d}$, this is at most e^{-d^2} for large enough d (depending on c). We have shown the following.

Lemma 5.3.1. *Fix $c > 1$. There are $d_0 \geq 1$ and positive C, both depending on c, such that for all $d \geq d_0$ there is a d-regular, bipartite graph G_d on $n = c^{d/\log d}$ vertices with bipartition classes $\mathcal{E}$ and $\mathcal{O}$ satisfying the following:*

1. *Every subset of $\mathcal{E}$ or $\mathcal{O}$ of size j, with $1 \leq j \leq 3n(\log d)/d$, has at least $jd/(C \log d)$ neighbors.*
2. *Every pair of subsets each of size $3n(\log d)/d$, one from $\mathcal{E}$ and one from $\mathcal{O}$, have an edge between them.*

We now fix such a G_d and study $Z_\Lambda(G_d, H)$. Given $f \in \mathrm{Hom}(G_d, H)$ set

$$\mathcal{E}(f) = \{k \in V(H) : |f^{-1}(k) \cap \mathcal{E}| \geq 3n(\log d)/d\}$$

and

$$\mathcal{O}(f) = \{k \in V(H) : |f^{-1}(k) \cap \mathcal{O}| \geq 3n(\log d)/d\}.$$

Clearly both $\mathcal{E}(f)$ and $\mathcal{O}(f)$ are non-empty, and by Lemma 5.3.1, we have $\mathcal{E}(f) \sim \mathcal{O}(f)$ (that is, everything in $\mathcal{E}(f)$ is adjacent to everything in $\mathcal{O}(f)$). So we can partition $\mathrm{Hom}(G_d, H)$ into classes indexed by pairs (A, B) with $A \sim B$. Write $\mathcal{C}(A, B)$ for the class corresponding to (A, B). We want to establish that for $(A, B) \in \mathcal{M}_\Lambda(H)$ we have

$$\sum_{f \in \mathcal{C}(A,B)} w_\Lambda(f) = (1 + o(1))\eta_\Lambda(H)^{n/2} \tag{5.14}$$

while for all other (A, B) we have

$$\sum_{f \in \mathcal{C}(A,B)} w_\Lambda(f) = o\left(\eta_\Lambda(H)^{n/2}\right), \tag{5.15}$$

where all asymptotic terms are (unless stated otherwise) as $d \to \infty$. From this we see that

$$Z_\Lambda(G_d, H) = |\mathcal{M}_\Lambda(H)|(1 + o(1))\eta_\Lambda(H)^{n/2},$$

and that all but a vanishing proportion of $Z_\Lambda(G_d, H)$ comes from pure-(A, B) colorings (with $(A, B) \in \mathcal{M}_\Lambda(H)$) in which $\mathcal{E}$ is mapped to A and $\mathcal{O}$ to B, with each such (A, B) contributing equally to $Z_\Lambda(G_d, H)$; this is enough to give the first part of Theorem 5.1.4. Indeed, fix $(A, B) \in \mathcal{M}_\Lambda(H)$. A proportion $(1 + o(1))/|\mathcal{M}_\Lambda|$ of $Z_\Lambda(G_d, H)$ is obtained by independently coloring $\mathcal{E}$ from A and $\mathcal{O}$ from B according to the given weights. Fix $k \in A$. We claim that with very high probability, a proportion very close to λ_k/λ_A of $\mathcal{E}$ gets mapped to k. Set $p = \lambda_k/\lambda_A$ and $m = n/2$. The number U_k of vertices of $\mathcal{E}$ mapped to k is a binomial random variable with parameters m and p. So by Tchebychev's inequality,

$$\Pr\left(|U_k - pm| \geq \log m\sqrt{mp(1-p)}\right) \leq \frac{1}{\log^2 m}.$$

This shows that the proportion of vertices mapped to k in a pure-(A, B) coloring is very close to

$$\frac{\lambda_k \mathbf{1}_{\{k \in A\}}}{2\lambda_A} + \frac{\lambda_k \mathbf{1}_{\{k \in B\}}}{2\lambda_B}$$

with high probability. Applying this with $(A, B) = (A^+, B^+)$ and $(A, B) = (A^-, B^-)$, the first part of Theorem 5.1.4 follows.

The lower bound in (5.14) is obtained by considering pure-(A, B) colorings with $\mathcal{E}$ mapped to A and $\mathcal{O}$ to B. To establish (5.15) and the upper bound in (5.14), fix $0 \leq j \leq 3n(\log d)/d$, let $q = |V(H)|$, and assume that d is large. We consider the contribution to $\sum_{f\in\mathcal{C}(A,B)} w_\Lambda(f)$ from those $f \in \mathcal{C}(A,B)$ in which, for each $k \notin A\cup B$, we have at most j vertices mapped to k, and we have at least one $k' \notin A \cup B$ whose preimage has size j. To bound the contribution from these f, we first bound the number of ways of locating the vertices that are mapped to k for each $k \notin A \cup B$ by $\left(\sum_{i\leq j}\binom{n}{i}\right)^q$. The contribution to the sum of the weights from these exceptional vertices is at most $(\lambda_H)^{qj}$. For the contribution from the remaining vertices, we deal separately with the cases $(A,B) \in \mathcal{M}_\Lambda(H)$ and $(A,B) \notin \mathcal{M}_\Lambda(H)$. For $(A,B) \notin \mathcal{M}_\Lambda(H)$, we simply upper bound the contribution by $(\lambda_A\lambda_B)^{n/2}$, leading to

$$\begin{aligned}\sum_{f\in\mathcal{C}(A,B)} w_\Lambda(f) &\leq (\lambda_A\lambda_B)^{\frac{n}{2}}\left(\sum_{i\leq j}\binom{n}{i}\right)^q (\lambda_H)^{qj} \\ &= o\left(\eta_\Lambda(H)^{n/2}\right),\end{aligned}$$

as required. For $(A,B) \in \mathcal{M}_\Lambda(H)$, consider a k' that has preimage size j. We claim that there are at least $jd/(2C\log d)$ vertices which, in the specification of f, need to be mapped to $A\cup B$ and which are adjacent to at least one of the j vertices mapped to k'. Indeed, by Lemma 5.3.1, the neighborhood size of the j vertices mapped to k' is at least $jd/(C\log d)$, and at most qj vertices have been mapped to vertices from outside $A\cup B$, so there are at least $jd/(C\log d) - qj > jd/(2C\log d)$ vertices that are adjacent to a vertex mapped to k' and need to be mapped to vertices from $A \cup B$. Since k' cannot be adjacent to everything in A, nor can it be adjacent to everything in B (else we would not have $(A,B) \in \mathcal{M}_\Lambda(H)$), our choice on these at least $jd/(2C\log d)$ vertices is restricted to a proper subset of $A\cup B$; the contribution we get from the remaining vertices (those mapped to $A\cup B$) is therefore at most

$$\frac{(\lambda_A\lambda_B)^{\frac{n}{2}}}{(1+\varepsilon)^{\frac{jd}{2C\log d}}}$$

where $\varepsilon > 0$ (depending on H and Λ) can be chosen uniformly for all A, B. Combining these observations we get that

$$\sum_{f \in \mathcal{C}(A,B)} w_\Lambda(f) \leq \eta_\Lambda(H)^{\frac{n}{2}} \frac{\left(\sum_{i \leq j} \binom{n}{i}\right)^q (\lambda_H)^{qj}}{(1+\varepsilon)^{\frac{jd}{2C \log d}}}.$$

If $j = 0$, the right-hand side above is $(\lambda_A \lambda_B)^{n/2}$. For $j > 0$ it can be bounded above by

$$\eta_\Lambda(H)^{n/2} \left(\frac{1}{(1+\varepsilon')^{\frac{d}{\log d}}} \right)^j$$

for some $\varepsilon' > 0$ (depending on H and Λ) for all j in the range $1 \leq j \leq 3n(\log d)/d$, as long as c is sufficiently small (recall $n = c^{d/\log d}$). Summing over j gives the upper bound in (5.14).

We now turn to the second part of Theorem 5.1.4. We take G'_d to be the disjoint union of m copies of $K_{d,d}$ where $m = m(d) = \omega(1)$. Fix $k \in V(H)$. Let X be the number of vertices mapped to k in a p_Λ-chosen H-coloring of G'_d, and X_i the number mapped to k in the ith copy of $K_{d,d}$. Define $a_\Lambda(k)$ by $E(X_i) = 2da_\Lambda(k)$, and note that $\mathrm{Var}(X_i) \leq 4d^2$. Since $X = \sum_{i=1}^m X_i$ we have $E(X) = 2dma_\Lambda(k)$ and $\mathrm{Var}(X) \leq 4md^2$. By Tchebychev's inequality,

$$P(|X - 2dma_\Lambda(k)| > 2dm\varepsilon) = P(|X/2dm - a_\Lambda(k)| > \varepsilon) \leq 1/m\varepsilon^2.$$

So choosing $\varepsilon = o(1)$ with $m\varepsilon^2 = \omega(1)$ (for example, $\varepsilon = 1/m^{1/3}$), the probability that the proportion of vertices mapped to k in a p_Λ-chosen H-coloring of G'_d differs from $a_\Lambda(k)$ by more than $o(1)$ is at most $o(1)$. The claimed bound on $s(k, f)$ follows, as does the estimate of $\bar{p}_\Lambda(k)$.

5.4 Results for non-regular graphs

The condition that G be regular can be relaxed quite a bit; we simply require that G has not too many low degree vertices, that the sum of the degrees of high degree vertices is not too large, and that the difference between the sizes of the partition classes is not too great. Here we state the theorems and describe the main differences in the proofs; for details we refer the reader to [20].

Theorem 5.4.1. *Fix H and Λ. There are positive constants c_1, c_2, c_3 and c_4 (depending on H and Λ) such that the following statements hold. Let G be a bipartite graph on n vertices with bipartition classes $\mathcal{E}$ and $\mathcal{O}$ (with $|\mathcal{O}| \geq |\mathcal{E}|$). Let d be an arbitrary positive parameter. Let ε satisfy $\varepsilon \geq c_1\sqrt{h(G,d)}$ where*

$$h(G,d) = \frac{1}{d} + \frac{|\{v \in \mathcal{E} : d(v) < d\}|}{n} + \frac{|\mathcal{O}| - |\mathcal{E}|}{n} + \frac{1}{dn}\sum_{v \in \mathcal{O}}(d(v) - d)\mathbf{1}_{\{d(v) \geq d\}}.$$

Then for each $k \in V(H)$ we have (5.2), as well as (5.3) with now

$$\zeta = \max\left\{\sqrt{h(G,d)}, \sqrt{\frac{\log n}{n}}\right\}.$$

If G is d-regular then $h(G,d) = 1/d$ and so Theorem 5.4.1 is a generalization of Theorem 5.1.2. The proof of Theorem 5.4.1 follows the same lines as already described for Theorem 5.1.2, except that we now require a new upper bound on $Z_\Lambda(G,H)$. We modified the entropy-based proof of (5.4) to obtain the following, which is just what we need for Theorem 5.4.1.

Theorem 5.4.2. *Fix H and Λ, and suppose that $\lambda_i > 1$ for all $i \in V(H)$. Let G be any bipartite graph on bipartition classes $\mathcal{E}$ and $\mathcal{O}$, with $|\mathcal{O}| \geq |\mathcal{E}|$, and let d be an arbitrary positive parameter. Then*

$$Z_\Lambda(G,H) \leq (\lambda_H)^{|\{w \in \mathcal{E} : d(w) < d\}|} \prod_{v \in \mathcal{O}} Z_\Lambda(K_{d(v),d}, H)^{\frac{1}{d}}.$$

Note that if G is d-regular then Theorem 5.4.2 reduces to (5.4). Note also that the condition imposed on the λ_i by Theorem 5.4.2 is not restrictive: if Λ' is obtained from Λ by multiplying all $\lambda_i \in \Lambda$ by the same positive constant then $p_\Lambda(n_1(f) = \cdot) = p_{\Lambda'}(n_1(f) = \cdot)$ and so we may assume without loss of generality that $\min\{\lambda_i : i \in V(H)\} > 1$.

Theorem 5.4.1 is only of interest in situations where $h(G, d)$ can be shown to be small (as, for example, when G is d-regular). A natural situation where we can say something about $h(G, d)$ is in percolation. Given a graph G and a parameter $0 \leq p \leq 1$, let G_p be a random subgraph of G obtained by deleting each edge independently with probability $1 - p$ (so the probability that $G_p = H$ is $p^{|E(H)|}(1-p)^{|E(G)|-|E(H)|}$). A corollary of Theorem 5.4.1 is the following "phase transition" phenomenon for percolation on a regular bipartite graph. If G is a d-regular bipartite graph and p is much greater than $1/d$, then the typical appearance of a p_Λ-chosen H-coloring of G_p is similar to that of a p_Λ-chosen H-coloring of G, whereas if p is much smaller than $1/d$, then as long as there is some $k \in V(H)$ with $\lambda_k/\lambda_H \notin [a_\Lambda^-(k), a_\Lambda^+(k)]$, these two objects have different appearances.

Corollary 5.4.3. *Fix H and Λ. Let $f(d) = \omega(1)$. There is a function $g(d) = o(1)$ (depending on $f(d)$) such that if $\{G^d\}_{d=1}^\infty$ is a sequence of d-regular bipartite graphs and p satisfies $p \geq f(d)/d$, then with probability at least $1 - g(d)$ the graph G_p^d satisfies that for each $k \in V(H)$ we have*

$$p_\Lambda\left(s(k, f) \in I_k(g(d))\right) \leq g(d)$$

and

$$\bar{p}_\Lambda(k) \in \left[a_\Lambda^-(k) - g(d), a_\Lambda^+(k) + g(d)\right].$$

If on the other hand $p \leq 1/(f(d)d)$ then with probability at least $1 - g(d)$ we have *that for each* $k \in V(H)$*,*

$$p_\Lambda\left(\left|s(k,f) - \frac{\lambda_k}{\lambda_H}\right| \leq g(d)\right) \geq 1 - g(d)$$

and

$$\left|\bar{p}_\Lambda(k) - \frac{\lambda_k}{\lambda_H}\right| \leq g(d).$$

For the multicast model (Example B), for example, we have

$$a_\Lambda^-(0) = a_\Lambda^+(0) = \frac{1}{2\left(\sum_{i\leq\lfloor k/2\rfloor}\lambda^i\right)} + \frac{1}{2\left(\sum_{i\leq\lceil k/2\rceil}\lambda^i\right)} > \frac{1}{\sum_{i\leq k}\lambda^i}$$

and so Corollary 5.4.3 shows a phase transition for this model. For the uniform q-coloring model (Example C), on the other hand, Corollary 5.4.3 gives no information about what happens as p crosses $1/d$. Indeed, for uniform colorings we have $a^-(k) = \frac{1}{2\lceil q/2\rceil}$ and $a^+(k) = \frac{1}{2\lfloor q/2\rfloor}$ for all k, and $\lambda_k/\lambda_H = 1/q$. Therefore $\lambda_k/\lambda_H \in [a^-(k), a^+(k)]$ for all colors k.

CHAPTER 6

H-COLORING TORI

6.1 Introduction and statement of results

There have been numerous papers devoted to the study of the space of H-colorings of particular graphs and families of graphs, for various special instances of H. Some recent papers (see for example [7], [9], [24], [32], and also chapter 5) have taken a broader approach, treating the space of H-colorings for arbitrary H. This chapter also falls into this latter category.

Many of the graphs G on which it is natural (from a statistical physics viewpoint) to study $\mathrm{Hom}(G,H)$ are regular and bipartite. Examples include the infinite lattice $\mathbb{Z}^d$, the hexagonal lattice, and the Bethe lattice (regular tree). For this reason much attention has been focused on this special case, and in this chapter that is also where our focus lies. The notation established in Section 2.3 will be used here and it is a good idea to review that section before continuing on with the remainder of this chapter.

In [32], an entropy approach was taken to obtain nearly matching upper and lower bounds on $|\mathrm{Hom}(G,H)|$ for arbitrary H and d-regular bipartite G, specifically

$$\eta(H)^{\frac{|V(G)|}{2}} \leq |\mathrm{Hom}(G,H)| \leq \eta(H)^{\frac{|V(G)|}{2}} 2^{\frac{|V(G)||V(H)|}{2d}}. \tag{6.1}$$

In Chapter 5, this work was extended considerably. For all H and $k \in V(H)$, optimal numbers $a^+(k)$ and $a^-(k)$ are constructed with the following property: for each $\varepsilon > 0$, if f is uniformly chosen from $\mathrm{Hom}(G,H)$, then (for suitably large d) with high probability the proportion of vertices of G mapped to k is between $a^-(k) - \varepsilon$ and $a^+(k) + \varepsilon$.

Let G be a bipartite graph with fixed bipartition $\mathcal{E} \cup \mathcal{O}$. Recall that for $A, B \subseteq V(H)$ with $A \sim B$, a pure-(A, B) coloring is an $f \in \mathrm{Hom}(G, H)$ with $f(u) \in A$ for all $u \in \mathcal{E}$ and $f(v) \in B$ for all $v \in \mathcal{O}$. If G is regular and has n vertices, then the number of pure-(A, B) colorings of G is $(|A||B|)^{n/2}$. An intuition driving the results of [32] and Chapter 5 is that in a certain sense, most $f \in \mathrm{Hom}(G, H)$ are close to pure-(A, B) colorings for some (A, B) that maximizes $|A||B| \in \mathcal{M}(H)$.

Such an intuition cannot be formalized for *all* regular bipartite G — for example, by the independence of the coloring on different components of a disconnected graph, it is easy to see that the intuition cannot be true for a graph that consist of a large number of small components (see e.g. the family $\{G'_d\}_{d=1}^{\infty}$ from Theorem 5.1.4). If, however, we are working with connected graphs with reasonable expansion (meaning that each subset of vertices from one partition class has a reasonably large number of neighbors in the other class) then we might expect it to be true that most $f \in \mathrm{Hom}(G, H)$ are close to pure-(A, B) colorings for some (A, B). This is shown for random regular bipartite graphs, for example, in Theorem 5.1.4; the proof critically uses the excellent expansion of random graphs.

For other graphs with weaker but still good expansion we expect similar results. One family of graphs that is of particular interest, given the statistical physics interpretation of H-colorings, is the integer lattice $\mathbb{Z}^d$ with the usual nearest neighbor adjacency, together with its finite analog the discrete torus $\mathbb{Z}_m^d$, the graph obtained from an axis-parallel box in $\mathbb{Z}^d$ by identifying opposite faces. These graphs have been focus of study for particular homomorphism models (see e.g. [30] for independent sets and [8] for proper colorings), as well as for general H-colorings (see e.g. [7]).

In Chapter 5 information is given about the number of occurrences of each color in a uniformly chosen H-coloring of $\mathbb{Z}_m^d$, but no information is given about how the vertices of a particular color are distributed between $\mathcal{E}$ and $\mathcal{O}$. Some special cases of this problem have been previously addressed, as we now discuss.

In [43], in the course of deriving the asymptotic formula

$$|\mathrm{Hom}(Q_d, H_{\mathrm{ind}})| = (2\sqrt{e} + o(1))2^{2^{d-1}} \tag{6.2}$$

(as $d \to \infty$), Korshunov and Sapozhenko showed that if I is a uniformly chosen independent set from Q_d (that is, if I is the preimage of the unlooped vertex in a uniformly chosen f from $\mathrm{Hom}(Q_d, H_{\mathrm{ind}})$), then with high probability I has size close to $2^d/4$ and is contained almost entirely in a single partition class. Kahn [39] and Galvin [26] extended these results to the case of I chosen from the set of independent sets according to the hard-core distribution with parameter λ, that is, the distribution in which each set I is chosen with probability proportional to $\lambda^{|I|}$ for some $\lambda > 0$ (Korshunov and Sapozhenko's setting is $\lambda = 1$).

In [40], Kahn considered the set $\mathrm{Hom}(Q_d, \mathbb{Z})/\!\sim$ (where $\mathbb{Z}$ is given a graph structure by declaring consecutive integers to be adjacent, and $\sim$ is the equivalence relation defined by $h \sim g$ if and only if $h - g$ is a constant function). Answering a question of Benjamini, Häggström and Mossel [4], he showed that if f is a uniformly chosen element from this set (a "cube-indexed random walk"), then with high probability f takes on only constantly many values (independent of d). Extending this work, Galvin [23] showed that in fact f takes on only at most five (consecutive) values, that f is constant on all but $o(2^d)$ (actually, at most $g(d)$ for any $g(d) = \omega(1)$) vertices on one of the two bipartition classes of Q_d, and that on the other partition classes each of two values appear on $(1/4 - o(1))2^d$ of the vertices. Using a correspondence between $\mathrm{Hom}(Q_d, \mathbb{Z})/\!\sim$ and $\mathrm{Hom}(Q_d, K_3)$, the results of [23] also answer the question of the structure of a typical (uniformly chosen) proper 3-coloring of Q_d. In the process of showing

$$|\mathrm{Hom}(Q_d, K_3)| = (6e + o(1))2^{2^{d-1}} \tag{6.3}$$

it is shown in [23] that $\mathrm{Hom}(Q_d, K_3)$ may be partitioned into an exceptional subset of size $o(1)|\mathrm{Hom}(Q_d, K_3)|$, and six equal sized subsets, with the property that within each of these six subsets, all colorings are constant on all but $o(2^d)$ (again, actually at most $g(d)$ for any $g(d) = \omega(1)$) vertices on one of the two bipartition classes of Q_d, and on the other partition classes each of two colors appear on $(1/4 - o(1))2^d$ of the vertices. Peled [55] has recently extended these results on the 3-coloring and cube-indexed random walk models to more general tori.

One of the main purposes of this chapter is to extend these structural characterizations of $\mathrm{Hom}(Q_d, H_{\mathrm{ind}})$ and $\mathrm{Hom}(Q_d, K_3)$ to arbitrary H and from Q_d to $\mathbb{Z}_m^d$ for all even m. We also extend to the class of probability distributions p_Λ on $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ (as defined in Section 2.3). Because of a technical limitation of one step in our proof, all weights λ_i under consideration in this paper will be rational.

We will always think of d as the variable in our functions, with m, H and (when present) Λ some fixed parameters, and so all implicit constants depend only on m, H and Λ, but not on d. Where necessary we will always assume that d is large enough to support our assertions.

Recall some basic notation as defined in Section 2.3: for $S \subseteq \mathrm{Hom}(\mathbb{Z}_m^d, H)$ and $T \subseteq V(H)$ we write $w_\Lambda(S)$ for $\sum_{f \in S} w_\Lambda(f)$ and λ_T for $\sum_{k \in T} \lambda_k$. Set

$$\eta_\Lambda(H) = \max\{\lambda_A \lambda_B \colon A, B \subseteq V(H),\ A \sim B\}$$

and

$$\mathcal{M}_\Lambda(H) = \left\{(A, B) \in V(H)^2 \colon A \sim B,\ \lambda_A \lambda_B = \eta_\Lambda(H)\right\}.$$

We now state our first main result, a structural decomposition of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ (in the presence of weight-set Λ) into finitely many classes of similar-looking colorings.

Theorem 6.1.1. *Fix H, rational Λ and $m \geq 2$ even. There is a partition of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ into $|\mathcal{M}_\Lambda(H)| + 1$ classes as*

$$\mathrm{Hom}(\mathbb{Z}_m^d, H) = D_\Lambda(0) \cup \bigcup_{(A,B)\in\mathcal{M}_\Lambda(H)} D_\Lambda(A,B)$$

with the following properties.

1. *$w_\Lambda(D_\Lambda(0)) \leq 2^{-\Omega(d)} Z_\Lambda(\mathbb{Z}_m^d, H)$.*
2. *For each $(A,B) \in \mathcal{M}_\Lambda(H)$ and $f \in D_\Lambda(A,B)$, the number of vertices $v \in \mathcal{E}$ (resp. $\mathcal{O}$) with $f(v) \notin A$ (resp. $f(v) \notin B$) is at most $(m - \Omega(1))^d$, and moreover all but at most $(m-\Omega(1))^d$ vertices w of $\mathcal{O}$ (resp. $\mathcal{E}$) have the property that all colors from A (resp. B) appear on $N(w)$.*

We prove Theorem 6.1.1 in Section 6.3. This decomposition already gives significant information about the structure of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ and the distribution p_Λ on $\mathrm{Hom}(\mathbb{Z}_m^d, H)$. For the purpose of obtaining long-range influence results (see Section 6.2), we need a slightly stronger decomposition result that in addition quantifies the number of vertices of each color in an arbitrary element of each partition class as well as the sizes of the partition classes. In what follows we use $X = Y(1 \pm 2^{-\Omega(d)})$ to indicate $|X/Y - 1| \leq 2^{-\Omega(d)}$.

Theorem 6.1.2. *Fix H, rational Λ and $m \geq 2$ even. There is a partition of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ into $|\mathcal{M}_\Lambda(H)| + 1$ classes as*

$$\mathrm{Hom}(\mathbb{Z}_m^d, H) = C_\Lambda(0) \cup \bigcup_{(A,B)\in\mathcal{M}_\Lambda(H)} C_\Lambda(A,B)$$

with the following properties.

1. *$w_\Lambda(C_\Lambda(0)) \leq 2^{-\Omega(d)} Z_\Lambda(\mathbb{Z}_m^d, H)$.*
2. *For each $(A,B) \in \mathcal{M}_\Lambda(H)$, $f \in C_\Lambda(A,B)$, $k \in A$ and $\ell \in B$, the proportion of vertices of $\mathcal{E}$ (resp. $\mathcal{O}$) colored k (resp. ℓ) is within $2^{-\Omega(d)}$ of λ_k/λ_A (resp. λ_ℓ/λ_B).*

3. *If $A \neq B$ is such that $(A,B),(B,A) \in \mathcal{M}_\Lambda(H)$ then*

$$w_\Lambda(C_\Lambda(A,B)) = w_\Lambda(C_\Lambda(B,A))\left(1 \pm 2^{-\Omega(d)}\right).$$

4. *If $(A,B),(\tilde{A},\tilde{B}) \in \mathcal{M}_\Lambda(H)$ are such that $\varphi(A) = \tilde{A}$ and $\varphi(B) = \tilde{B}$ for some weight preserving automorphism φ of H, then*

$$w_\Lambda(C_\Lambda(A,B)) = w_\Lambda(C_\Lambda(\tilde{A},\tilde{B}))\left(1 \pm 2^{-\Omega(d)}\right).$$

5. *For each $(A,B) \in \mathcal{M}_\Lambda(H)$, $x \in \mathcal{E}$, $y \in \mathcal{O}$, $k \in A$ and $\ell \in B$,*

$$p_\Lambda(f(x) = k | f \in C_\Lambda(A,B)) = \frac{(1 \pm 2^{-\Omega(d)})\lambda_k}{\lambda_A}$$

   and

$$p_\Lambda(f(y) = \ell | f \in C_\Lambda(A,B)) = \frac{(1 \pm 2^{-\Omega(d)})\lambda_\ell}{\lambda_B}.$$

In Section 6.3 we derive Theorem 6.1.2 from Theorem 6.1.4. In Section 6.2 we discuss a long-range influence phenomenon that is implied by Theorem 6.1.2.

Theorem 6.1.2 does not make a general statement about the relative sizes of the $C_\Lambda(A,B)$'s, but there are two important situations in which we can conclude that the partition of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ guaranteed by Theorem 6.1.2 is an *approximate equipartition*.

Definition 6.1.3. Fix H, rational Λ and $m \geq 2$ even. An *approximate equipartition* of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ is a partition into $|\mathcal{M}_\Lambda(H)| + 1$ classes satisfying conditions (1) through (5) of Theorem 6.1.2, as well as the condition that for all $(A,B),(A',B') \in \mathcal{M}_\Lambda(H)$ we have

$$w_\Lambda(C_\Lambda(A,B)) = \left(1 \pm 2^{-\Omega(d)}\right) w_\Lambda(C_\Lambda(A',B')).$$

A corollary of statements 1 and 3 is that if $\mathcal{M}_\Lambda(H) = \{(A,B),(B,A)\}$ for some $A \neq B$ (as, for example, in the case $H = H_{\mathrm{ind}}$ for arbitrary Λ), then the Theorem

6.1.2 partition of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ is an approximate equipartition with

$$w_\Lambda\left(C_\Lambda(A,B)\right), \, w_\Lambda\left(C_\Lambda(B,A)\right) = Z_\Lambda(\mathbb{Z}_m^d, H)\left(\frac{1}{2} \pm 2^{-\Omega(d)}\right).$$

Furthermore, if $\mathcal{M}_\Lambda(H) = \{(A,A)\}$ for some A then the Theorem 6.1.2 partition of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ is trivially an approximate equipartition with $w_\Lambda\left(C_\Lambda(A,A)\right) = Z_\Lambda(\mathbb{Z}_m^d, H)\left(1 - 2^{-\Omega(d)}\right)$. These are in a sense the two generic situations, as for every H, if the weights λ_i are chosen from any continuous distribution supported on $\{x \in \mathbb{R}^{|V(H)|} : x > 0\}$ then with probability 1 we will have $\mathcal{M}_\Lambda(H)$ of one of the two forms described.

A corollary of statements 1 and 4 is that if $\mathcal{M}_\Lambda(H)$ is transitive, that is, if for each $(A,B), (\tilde{A}, \tilde{B}) \in \mathcal{M}_\Lambda(H)$ there is a weight preserving automorphism φ of H with $\varphi(A) = \tilde{A}$ and $\varphi(B) = \tilde{B}$, then the partition of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ is an approximate equipartition with

$$w_\Lambda\left(C_\Lambda(A,B)\right) = Z_\Lambda(\mathbb{Z}_m^d, H)\left(\frac{1}{|\mathcal{M}_\Lambda(H)|} \pm 2^{-\Omega(d)}\right).$$

This is far from a generic situation, but is the case for a number of very important examples, such as the uniform proper q-coloring model ($H = K_q$ and $\Lambda = (1, \ldots, 1)$), where it easily seen that

$$|\mathcal{M}_\Lambda(K_q)| = \begin{cases} \binom{q}{q/2} & \text{if } q \text{ even} \\ \binom{q}{(q-1)/2} + \binom{q}{(q+1)/2} & \text{if } q \text{ odd,} \end{cases}$$

or more concisely $|\mathcal{M}_\Lambda(K_q)| = (1 + \mathbf{1}_{\{q \text{ odd}\}})\binom{q}{[q/2]}$. (Note that $\mathcal{M}(K_q)$ consists of all pairs (A,B) with A and B disjoint, $A \cup B = V(K_q)$, and $|A|$, $|B|$ as near equal as possible). Another example of this behavior is the uniform Widom-Rowlinson model (H the complete looped path on three vertices, or equivalently the complete looped graph on $\{1,2,3\}$ with edge 13 removed). In this case we have $\mathcal{M}_\Lambda(H) = \{(A,A),(B,B)\}$ with $A=\{1,2\}$ and $B=\{2,3\}$.

The existence of these equipartitions is what drives our long-range influence results: Corollaries 6.2.2, 6.2.3 and 6.2.4 in Section 6.2. A representative result from that section is the following: in a proper q-coloring of Q_d chosen uniformly conditioned on a particular vertex $v \in \mathcal{E}$ being colored 1, the probability that another vertex $u \in \mathcal{E}$ is colored 1 is close to $2/q$, whereas the probability that a vertex $w \in \mathcal{O}$ is colored 1 is close to 0, regardless of the distances between u, v and w.

In general, we cannot say anything more about the relative (Λ-weighted) sizes of the $C_\Lambda(A,B)$ as various different types of behaviors are possible. We'll illustrate this now with two examples. A fact that we use in both examples is that for G connected and H consisting of components H_1 and H_2 we can identify $\mathrm{Hom}(G,H)$ with the disjoint union of $\mathrm{Hom}(G,H_1)$ and $\mathrm{Hom}(G,H_2)$.

First, consider H the disjoint union of H_{ind} and K_3 (note that $\eta(H_{\mathrm{ind}}) = \eta(K_3) = 2$) with $\Lambda = (1,\ldots,1)$. The results of [43] and [23] (see (6.2), (6.3) and the discussions around these equations) together imply that in any decomposition of $\mathrm{Hom}(Q_d,H)$ satisfying the conditions of Theorem 6.1.2, along with the exceptional class we have eight partition classes. Six of these correspond to the six elements of $\mathcal{M}(K_3)$, and these each have size $(1+o(1))e/(6e+2\sqrt{e})|\mathrm{Hom}(Q_d,H)| \approx .14|\mathrm{Hom}(Q_d,H)|$. The two remaining classes correspond to the two elements of $\mathcal{M}(H_{\mathrm{ind}})$ and each have size $(1+o(1))\sqrt{e}/(6e+2\sqrt{e})|\mathrm{Hom}(Q_d,H)| \approx .08|\mathrm{Hom}(Q_d,H)|$.

For an example with a different type of behavior, let H be the disjoint union of K_4^{loop} (the complete looped graph on four vertices) and K_8 (note that $\eta(K_8) = \eta(K_4^{\mathrm{loop}}) = 16$, with $\mathcal{M}(K_4^{\mathrm{loop}}) = (V(K_4^{\mathrm{loop}}), V(K_4^{\mathrm{loop}})))$, again with $\Lambda = (1,\ldots,1)$. It is immediate that $|\mathrm{Hom}(Q_d,K_4^{\mathrm{loop}})| = 16^{2^{d-1}}$ and that all colorings in this set are pure-$(V(K_4^{\mathrm{loop}}), V(K_4^{\mathrm{loop}}))$ colorings. It is also fairly straightforward to verify that $|\mathrm{Hom}(Q_d,K_8)| = \omega(16^{2^{d-1}})$. Indeed, consider proper 8-colorings of Q_d which are pure-(A, B) for some (A, B), except that there is one vertex from $\mathcal{E}$ that is colored from B. An easy count gives that there are $(1/2)(3/2)^d 16^{2^{d-1}}$ such colorings. This implies that in any decomposition of $\mathrm{Hom}(Q_d, H)$ satisfying the conditions of Theorem 6.1.2, along with the exceptional set we have $\binom{8}{4} + 1$ partition classes. The first $\binom{8}{4}$ of these classes correspond to the elements of $\mathcal{M}(K_8)$ and each have size $\Omega(|\mathrm{Hom}(Q_d, H)|)$, and the last class corresponds to the unique element of $\mathcal{M}(K_4^{\mathrm{loop}})$ and has size $o(|\mathrm{Hom}(Q_d, H)|)$.

The proof of Theorem 6.1.2 is based on the notion of an ideal edge. Let H and $f \in \mathrm{Hom}(\mathbb{Z}_m^d, H)$ be given. Say that an edge $e = uv \in E$ (with $u \in \mathcal{E}$) is *ideal* (with respect to f) if $f(N(u)) = B$ and $f(N(v)) = A$ for some $(A, B) \in \mathcal{M}(H)$. We will only be interested in the probability that a particular edge is not ideal with respect to f, when f is chosen uniformly from $\mathrm{Hom}(\mathbb{Z}_m^d, H)$. Note that by the symmetry of the torus, this probability is independent of the particular edge we choose. Our main technical result is the following.

Theorem 6.1.4. *Fix H, $m \geq 2$ even, and $e \in E$. If f is chosen uniformly from $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ then*

$$\Pr(e \textit{ is not ideal with respect to } f) \leq 2^{-\Omega(d)}.$$

The analogous result for $m = 2$ and $H = \mathbb{Z}$ (with two elements of $\mathrm{Hom}(Q_d, \mathbb{Z})$ identified if they differ by a constant) was proved by Kahn in [40], and our proof follows similar lines. A standard trick of comparing a weighted H-coloring model to a uniform H'-coloring model for a certain graph H' (depending on H an Λ) makes the generalization from uniform to arbitrary Λ relatively straightforward.

All of our results are for fixed m, and become interesting as d grows. It would be of great interest to obtain similar results for fixed d, as m grows (as Peled [55] has done in the case $H = K_3$), as this would allow us to say something about the space of

Gibbs measures for the probability distribution p_Λ on the infinite space $\mathrm{Hom}(\mathbb{Z}^d, H)$ (see for example [7], [9], for a discussion of Gibbs measures in the specific context of homomorphism models). Unfortunately, a careful examination of our proof of Theorem 6.1.4, keeping track of the dependency of the final constants on m, shows that at best we may take $m = c \log d$ for some absolute constant $c > 0$ if we wish to obtain useful results.

6.2 Long-range influence

Roughly speaking we say that a distribution p_Λ on $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ exhibits *long-range influence* if the distribution of p_Λ restricted to a single vertex x is sensitive to conditioning on the color of another vertex y, even in the limit as d and the distance from x to y go to infinity.

More formally, given a graph H, a weight set Λ and even m, we say that the Λ-weighted H-coloring model on $\mathbb{Z}_m^d$ exhibits long-range influence if there is a choice of $x, y \in V$ and $k, \ell \in V(H)$ (actually a sequence of choices, one for each d) with $\mathrm{dist}(x, y) = \omega(1)$ (where dist is usual graph distance) such that

$$\frac{p_\Lambda(f(x) = k | f(y) = \ell)}{p_\Lambda(f(x) = k)} \not\to 1 \text{ as } d \to \infty. \tag{6.4}$$

Theorem 6.1.2 strongly implies such a phenomenon, at least in the case where the partition of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ guaranteed by Theorem 6.1.2 is an approximate equipartition. The following is an immediate corollary of Theorem 6.1.2, and in particular statement 5 of that theorem.

Theorem 6.2.1. *Fix H, rational Λ and $m \geq 2$ even. Suppose that the Theorem 6.1.2 partition of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ is an approximate equipartition. Fix $k, \ell \in V(H)$. For all*

$x \in \mathcal{E}$ we have

$$p_\Lambda(f(x) = k) = \left(\frac{1}{|\mathcal{M}_\Lambda(H)|} \pm 2^{-\Omega(d)}\right) \sum_{(A,B)\in\mathcal{M}_\Lambda(H)\colon\, k\in A} \frac{\lambda_k}{\lambda_A}$$

(and by symmetry this is also true for $x \in \mathcal{O}$). On the other hand, if $x, y \in \mathcal{E}$ then

$$p_\Lambda(f(x) = k | f(y) = \ell) = \left(\frac{1}{|\mathcal{M}_\Lambda(H)|} \pm 2^{-\Omega(d)}\right) \sum_{(A,B)\in\mathcal{M}_\Lambda(H)\colon\, \ell, k\in A} \frac{\lambda_k}{\lambda_A}$$

and if $x \in \mathcal{E}$ and $y \in \mathcal{O}$ then

$$p_\Lambda(f(x) = k | f(y) = \ell) = \left(\frac{1}{|\mathcal{M}_\Lambda(H)|} \pm 2^{-\Omega(d)}\right) \sum_{(A,B)\in\mathcal{M}_\Lambda(H)\colon\, k\in A,\, \ell\in B} \frac{\lambda_k}{\lambda_A}.$$

By choosing k, ℓ appropriately, these three quantities can be made to be different (in the limit as $d \to \infty$). Rather than stating an unwieldy general proposition to this effect, we illustrate it with three examples. It will be helpful first to set up some notation. Fix m, H and Λ. For each $d \in \mathbb{N}$ and $x \in V$, we define the *occupation probability vector* $\vec{v}_d(x)$ by

$$\vec{v}_d(x) = (p_\Lambda(f(x) = k)\colon k \in V(H))\,.$$

(We suppress dependance on m, H and Λ to aid readability.) If the choice of f is conditioned on an event E we use $\vec{v}_d(x|E)$ to denote the conditional occupation probability vector, that is,

$$\vec{v}_d(x|E) = (p_\Lambda(f(x) = k|E)\colon k \in V(H))\,.$$

In what follows we use $d_\infty(\cdot, \cdot)$ for ℓ_∞ distance.

Our first example is the independent set model, that is, $H = H_{\text{ind}}$ where $V(H_{\text{ind}}) =$

$\{0,1\}$ and $E(H_{\text{ind}}) = \{00, 01\}$. We list 0 first in the occupation and conditional occupation probability vectors. Our weighting vector will assign weight 1 to 0 and rational weight λ to 1. (This is the *hard-core model* with fugacity λ, results on which from [39] have been discussed earlier.) Noting that $\mathcal{M}_\lambda(H_{\text{ind}}) = \{(A,B),(B,A)\}$ where $A = \{0,1\}$ and $B = \{0\}$, we have the following.

Corollary 6.2.2. *Fix $m \geq 2$ even and rational $\lambda > 0$. For all $x \in V$ we have*

$$d_\infty\left(\vec{v}_d(x), \left(\frac{2+\lambda}{2(1+\lambda)}, \frac{\lambda}{2(1+\lambda)}\right)\right) \leq 2^{-\Omega(d)}.$$

On the other hand, if $x, y \in \mathcal{E}$ then

$$d_\infty\left(\vec{v}_d(x|\{f(y)=1\}), \left(\frac{1}{1+\lambda}, \frac{\lambda}{1+\lambda}\right)\right) \leq 2^{-\Omega(d)}$$

and if $x \in \mathcal{E}$ and $y \in \mathcal{O}$ then

$$d_\infty\left(\vec{v}_d(x|\{f(y)=1\}), (1,0)\right) \leq 2^{-\Omega(d)}.$$

(This result was earlier proven in [26] for $m = 2$ and all λ (not necessarily rational) satisfying $\lambda > cd^{-1/3}\log d$ for some constant $c > 0$.)

Our second example is the uniform Widom-Rowlinson model. Here $H = H_{\text{WR}}$ is the graph on vertex set $\{1,2,3\}$ with all edges (and loops) present except the edge connecting 1 and 3. In the occupation and conditional occupation probability vectors we list the vertices in numerical order. Noting that $\mathcal{M}(H_{\text{WR}}) = \{(A,A),(B,B)\}$ where $A = \{1,2\}$ and $B = \{2,3\}$, we get the following via a routine calculation.

Corollary 6.2.3. *Fix $m \geq 2$ even. For all $x \in V$ we have*

$$d_\infty\left(\vec{v}_d(x), \left(\frac{1}{4}, \frac{1}{2}, \frac{1}{4}\right)\right) \leq 2^{-\Omega(d)}.$$

On the other hand, if $x, y \in \mathcal{E}$ then

$$d_\infty\left(\vec{v}_d(x|f(y)=1), \left(\frac{1}{2}, \frac{1}{2}, 0\right)\right) \leq 2^{-\Omega(d)}$$

while if $x \in \mathcal{E}$ and $y \in \mathcal{O}$ then

$$d_\infty\left(\vec{v}_d(x|f(y)=1), \left(0, \frac{1}{2}, \frac{1}{2}\right)\right) \leq 2^{-\Omega(d)}.$$

Our final example is the uniform proper q-coloring model ($H = K_q$ where $V(K_q) = \{1, \ldots, q\}$ and $E(K_q) = \{ij : i \neq j\}$, and $\Lambda = \vec{1}$). We list color 1 first in the occupation and conditional occupation probability vectors. By our earlier observation that $\mathcal{M}(H)$ consists of all pairs (A, B) with $A \cup B = \{1, \ldots, q\}$, $A \cap B = \emptyset$ and $|A| - |B| \in \{0, \pm 1\}$, we get the following via a routine calculation.

Corollary 6.2.4. *Fix $m \geq 2$ even and $q \in \mathbb{N}$. For all $x \in V$ we have*

$$\vec{v}_d(x) = \left(\frac{1}{q}, \ldots, \frac{1}{q}\right).$$

On the other hand, if $x, y \in \mathcal{E}$ then

$$d_\infty\left(\vec{v}_d(x|f(y)=1), \left(\frac{2}{q}, \frac{q-2}{q(q-1)}, \ldots, \frac{q-2}{q(q-1)}\right)\right) \leq 2^{-\Omega(d)}$$

and if $x \in \mathcal{E}$ and $y \in \mathcal{O}$ then

$$d_\infty\left(\vec{v}_d(x|f(y)=1), \left(0, \frac{1}{q-1}, \ldots, \frac{1}{q-1}\right)\right) \leq 2^{-\Omega(d)}.$$

The exact equality for $\vec{v}_d(x)$ here follows by symmetry. This corollary, in the special case $m = 2$ and $q = 3$, was proved in [25] (and is implicit in [23]).

In the uniform proper q-coloring model it is natural to allow q, the number of colors, to vary with d (see e.g. [11, 37, 38, 59]). We may define long-range influence in this case exactly as in (6.4), simply allowing H to also change with d.

The Dobrushin uniqueness theorem [18, 59] implies that we do not have long-range influence in the q-coloring model on $\mathbb{Z}_m^d$ when $q > 2d$ (in the case $m = 2$) or $q > 4d$ (in the case $m \geq 4$). On the other hand, Corollary 6.2.4 establishes that we do have long-range influence for all constant q.

We can in fact say a little bit more. Using a techniques described in this chapter (we will not outline the necessary changes here; a description may be found in [21]), we can prove the following.

Theorem 6.2.5. *Fix $m \geq 2$ even. If f is chosen uniformly from* $\mathrm{Hom}(\mathbb{Z}_m^d, K_q)$ *with* $q < (\log d)/(m+2)$*, then for any $x, y \in \mathcal{E}$ and $k \in \{1, \ldots, q\}$ we have*

$$\lim_{d \to \infty} \frac{\Pr(f(x) = k)}{\Pr(f(x) = k | f(y) = k)} = \frac{1}{2}.$$

6.3 Proofs of Theorems 6.1.1 and 6.1.2

We first note that if the weight set Λ' is obtained from Λ by multiplying each λ_k by the same constant, then the distributions p_Λ and $p_{\Lambda'}$ are identical. We may therefore assume without loss of generality that $\lambda_k \geq 1$ for all $k \in V(H)$.

Our main technical result (Theorem 6.1.4) considers uniformly chosen homomorphisms, so to apply it to homomorphisms chosen according to p_Λ we need to first relate p_Λ to uniform distribution on a graph $H(\Lambda)$ built from H and Λ. We use a technique introduced in [9].

Let $C = C(\Lambda)$ be the smallest integer such that $C\lambda_k$ is an integer for all $k \in V(H)$. For each k let S_k be an arbitrary set of size $C\lambda_k$, with the S_k's disjoint. We construct $H(\Lambda)$ on vertex set $\cup_{k \in V(H)} S_k$ by joining x and y if and only if $x \in S_k$ and $y \in S_\ell$ for some $k\ell \in E(H)$. Equivalently, $H(\Lambda)$ is obtained from H by replacing each vertex k by a set of size $C\lambda_k$, each edge by a complete bipartite graph and each loop by a complete looped graph; see Figure 6.1.

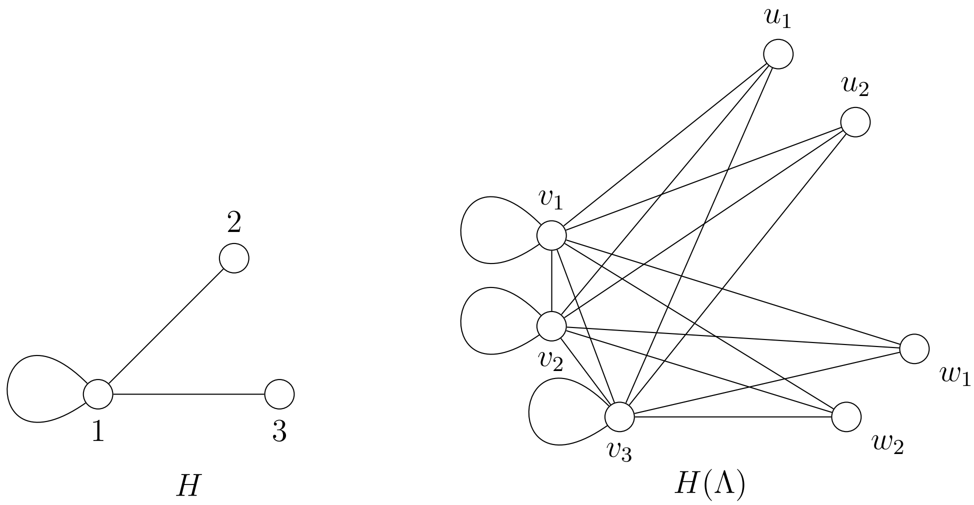

Figure 6.1. An example H and $H(\Lambda)$ with $\lambda_1 = 1$, $\lambda_2 = 3/2$ and $\lambda_3 = 1$, so $C = 2$. Here $S_1 = \{v_1, v_2, v_3\}$, $S_2 = \{u_1, u_2\}$, and $S_3 = \{w_1, w_2\}$.

For each $f \in \mathrm{Hom}(\mathbb{Z}_m^d, H)$ let A_f consist of those $g \in \mathrm{Hom}(\mathbb{Z}_m^d, H(\Lambda))$ with $g(v) \in S_{f(v)}$ for each $v \in V$. It is straightforward to verify that each A_f satisfies $|A_f| = C^{m^d} w_\Lambda(f)$, and that the A_f's form a partition of $\mathrm{Hom}(\mathbb{Z}_m^d, H(\Lambda))$. This implies that choosing an element g uniformly from $\mathrm{Hom}(\mathbb{Z}_m^d, H(\Lambda))$ and then letting $f \in \mathrm{Hom}(\mathbb{Z}_m^d, H)$ be such that $g \in A_f$ is equivalent to choosing f from $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ according to p_Λ.

Before continuing, we note the following easily established correspondence between $\mathcal{M}(H(\Lambda))$ and $\mathcal{M}_\Lambda(H)$:

$$
\begin{gathered}
|\mathcal{M}(H(\Lambda))| = |\mathcal{M}_\Lambda(H)| \\
\text{and} \\
(A', B') \in \mathcal{M}(H(\Lambda)) \text{ if and only if} \\
A' = \cup_{k\in A} S_k \text{ and } B' = \cup_{\ell \in B} S_\ell \text{ for some } (A,B) \in \mathcal{M}_\Lambda(H).
\end{gathered}
\tag{6.5}
$$

Now let g be chosen uniformly from $\mathrm{Hom}(\mathbb{Z}_m^d, H(\Lambda))$. By Theorem 6.1.4, the expected number of non-ideal edges of g is at most $(m - \Omega(1))^d$ and so by Markov's inequality there is a subset $\mathrm{Hom}'(\mathbb{Z}_m^d, H(\Lambda))$ of $\mathrm{Hom}(\mathbb{Z}_m^d, H(\Lambda))$ with

$$
|\mathrm{Hom}'(\mathbb{Z}_m^d, H(\Lambda))| \geq \left(1 - 2^{-\Omega(d)}\right) |\mathrm{Hom}(\mathbb{Z}_m^d, H(\Lambda))| \tag{6.6}
$$

and with each $g \in \mathrm{Hom}'(\mathbb{Z}_m^d, H(\Lambda))$ having at most $(m - \Omega(1))^d$ non-ideal edges.

We now need an isoperimetric bound on the discrete torus. The following result is due to Bollobás and Leader [6, Theorem 8].

Lemma 6.3.1. *Let $X \subseteq V$ satisfy $|X| \leq m^d/2$. The number of edges in E which have exactly one vertex in common with X is at least $|X|^{(d-1)/d}$.*

We will use the following corollary.

Corollary 6.3.2. *Let a satisfy $(ma)^{d/(d-1)} < 1/4$. If at most $m^d a$ edges are deleted from $\mathbb{Z}_m^d$ then the resulting graph has a component with at least $m^d(1 - (ma)^{d/(d-1)})$ vertices.*

Proof. Let $\mathcal{D}$ be the set of deleted edges, and let $C_1, C_2, \ldots, C_k$ be the components of the graph on vertex set V with edge set $E \setminus \mathcal{D}$, listed in order of increasing size (where size is measured by number of vertices). If $k = 1$, we are done. Otherwise, let $X = \cup_{i=1}^{\ell} C_i$ where ℓ is chosen as large as possible so that $|X| \leq m^d/2$. Since $\mathcal{D}$ includes all of the edges which have exactly one vertex in common with X, we have by Lemma 6.3.1

$$m^d a \geq |\mathcal{D}| \geq |X|^{\frac{d-1}{d}}$$

and so

$$|X| \leq m^d (ma)^{\frac{d}{d-1}} < m^d/4$$

(the final inequality by hypothesis). By the definition of ℓ, we have $|C_{\ell+1}| > m^d/4$. If $\ell = k-1$, we are done (since then $|C_k| \geq m^d(1-(ma)^{d/(d-1)})$). We complete the proof by arguing that we must have $\ell = k-1$. If not, let X' be the union of all the components other than $C_{\ell+1}$ and those in X. By the same argument as above (since $|X'| \leq m^d/2$) we have $|X'| < m^d/4 < |C_{\ell+1}|$. This is a contradiction, since by our ordering of the components X' is a union of components all at least as large as $C_{\ell+1}$. $\square$

Corollary 6.3.2 implies that for each $g \in \mathrm{Hom}'(\mathbb{Z}_m^d, H(\Lambda))$ there is a collection $\mathcal{F}$ of edges which spans a connected subgraph of $\mathbb{Z}_m^d$ on at least $m^d - (m-\Omega(1))^d$ vertices, and that all of these edges are ideal (note that in this application we have $a = 2^{-\Omega(d)}$ and so certainly $(ma)^{d/(d-1)} < 1/4$). By the connectivity of the subgraph induced by these edges, it follows that there is some $(A', B') \in \mathcal{M}(H(\Lambda))$ such that for each $uv \in \mathcal{F}$ with $u \in \mathcal{O}$, we have that $N(u)$ is colored from A' (and so in particular v is) and $N(v)$ is colored from B' (and so in particular u is). We may therefore decompose $\mathrm{Hom}'(\mathbb{Z}_m^d, H(\Lambda))$ as

$$\mathrm{Hom}'(\mathbb{Z}_m^d, H(\Lambda)) = \cup_{(A',B') \in \mathcal{M}(H(\Lambda))} D(A', B')$$

with the property that for each $g \in D(A', B')$ we can find a subset of V of size at least $m^d - (m-\Omega(1))^d$ with each vertex of this set colored from A' (resp. B') if it is in $\mathcal{E}$ (resp. $\mathcal{O}$), and moreover all but at most $(m-\Omega(1))^d$ vertices of $\mathcal{O}$ (resp. $\mathcal{E}$) have all of A' (resp. B') appearing on their neighborhoods.

We now pass to a partition of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$. For each $(A,B) \in \mathcal{M}_\Lambda(H)$, let $D_\Lambda(A,B)$ be the set of all $f \in \mathrm{Hom}(\mathbb{Z}_m^d, H)$ for which there is some $g \in A_f$ with $g \in D(A', B')$, where (A', B') is obtained from (A,B) by the correspondence described in (6.5). The $D_\Lambda(A,B)$'s are disjoint, for if $f \in D_\Lambda(A,B)$ (with corresponding $g \in D(A',B')$) and $\tilde{f} \in D_\Lambda(\tilde{A},\tilde{B})$ (with corresponding $\tilde{g} \in D(\tilde{A}', \tilde{B}')$) with $(A,B) \neq (\tilde{A},\tilde{B})$, the neighborhoods of the endvertices of any edge which is ideal for both g and $\tilde{g}$ witness that $f \neq \tilde{f}$.

Moreover, $D_\Lambda(A,B)$ inherits from $D(A',B')$ that for all $f \in D_\Lambda(A,B)$, the number of vertices $v \in \mathcal{E}$ (resp. $\mathcal{O}$) with $f(v) \notin A$ (resp. $f(v) \notin B$) is at most $(m-\Omega(1))^d$ (for concreteness, $(m-\kappa)^d$ for some $0 < \kappa < m$ that depends on H and Λ but may be chosen to be independent of (A,B)), and moreover all but at most $(m-\Omega(1))^d$ vertices w of $\mathcal{O}$ (resp. $\mathcal{E}$) have the property that all colors from A (resp. B) appear on $N(w)$.

Set $D_\Lambda(0) = \mathrm{Hom}(\mathbb{Z}_m^d, H) \setminus \cup_{(A,B)\in\mathcal{M}_\Lambda(H)} D_\Lambda(A,B)$. If $f \in D_\Lambda(0)$ then $A_f \subseteq \mathrm{Hom}(\mathbb{Z}_m^d, H(\Lambda)) \setminus \mathrm{Hom}'(\mathbb{Z}_m^d, H(\Lambda))$ and so by (6.6)

$$C^{m^d} w_\Lambda(D_\Lambda(0)) \leq 2^{-\Omega(d)} |\mathrm{Hom}(\mathbb{Z}_m^d, H(\Lambda))| = 2^{-\Omega(d)} C^{m^d} Z_\Lambda(\mathbb{Z}_m^d, H(\Lambda)).$$

This completes the proof of Theorem 6.1.1.

We now turn to Theorem 6.1.2. Our construction of the $C_\Lambda(A,B)$'s will be from scratch (and so in particular we will not refer to ideal edges); however, to establish the required properties of the $C_\Lambda(A,B)$'s we will relate them to the $D_\Lambda(A,B)$'s.

For each $(A,B) \in \mathcal{M}_\Lambda(H)$ we define a set $C_\Lambda(A,B)'$ as follows. First, for each $F_1 \subseteq \mathcal{E}$ and $F_2 \subseteq \mathcal{O}$ with $|F_1|+|F_2| \leq (m-\kappa)^d$ (with κ as described in the construction of $D_\Lambda(A,B)$ above), let $C_\Lambda^{(F_1,F_2)}(A,B)'$ include all $f \in \mathrm{Hom}(\mathbb{Z}_m^d, H)$ for which every vertex of $\mathcal{E} \setminus F_1$ is colored from A, every vertex from F_1 is colored from A^c, every vertex of $\mathcal{O} \setminus F_2$ is colored from B, and every vertex from F_2 is colored from B^c (note that for some choices of (F_1, F_2) we may have $C_\Lambda^{(F_1,F_2)}(A,B)' = \emptyset$). Next, set

$$C_\Lambda(A,B)' = \cup_{(F_1,F_2)} C_\Lambda^{(F_1,F_2)}(A,B)'.$$

By our upper bound on $|F_1| + |F_2|$, we have $D_\Lambda(A,B) \subseteq C_\Lambda(A,B)'$ for each (A,B). It is also clear that $|C_\Lambda(A,B)'| = |C_\Lambda(B,A)'|$ (because the mapping from $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ to itself, induced by any automorphism of $\mathbb{Z}_m^d$ that maps $\mathcal{E}$ to $\mathcal{O}$, maps $C_\Lambda(A,B)'$ to $C_\Lambda(B,A)'$ bijectively, and is weight-preserving), and (for a similar reason) that if $\varphi(A) = \tilde{A}$ and $\varphi(B) = \tilde{B}$ for some weight-preserving automorphism φ of H then $|C_\Lambda(A,B)'| = |C_\Lambda(\tilde{A},\tilde{B})'|$. We do not yet have a partition of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$, however, as the $C_\Lambda(A,B)'$'s are not necessarily disjoint.

Most of the rest of the proof is devoted to establishing the following two facts. First, for each $(A,B) \in \mathcal{M}_\Lambda(H)$, $x \in \mathcal{E}$, $y \in \mathcal{O}$, $k \in A$ and $\ell \in B$, if f is chosen from $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ according to p_Λ then

$$\begin{aligned} p_\Lambda(f(x) = k | f \in C_\Lambda(A,B)') &= \tfrac{\left(1+2^{-\Omega(d)}\right)\lambda_k}{\lambda_A}, \\ p_\Lambda(f(y) = \ell | f \in C_\Lambda(A,B)') &= \tfrac{\left(1+2^{-\Omega(d)}\right)\lambda_\ell}{\lambda_B}. \end{aligned} \tag{6.7}$$

For the second, say that $f \in C_\Lambda(A,B)'$ is *balanced* if for each $k \in A$ (resp. $\ell \in B$) the proportion of vertices of $\mathcal{E}$ (resp. $\mathcal{O}$) colored k (resp. ℓ) is within a multiplicative factor $1 \pm (1-\kappa/(4m))^d$ of λ_k/λ_A (resp. λ_ℓ/λ_B). For all $(A,B) \in \mathcal{M}_\Lambda(H)$ we have the following:

$$p_\Lambda\left(f \text{ is not balanced} | f \in C_\Lambda(A,B)'\right) \leq \exp\left\{-\left(m - \frac{\kappa}{2}\right)^d/4\right\}. \tag{6.8}$$

These two facts allow us to swiftly complete the proof of Theorem 6.1.2. Indeed, for each $(A,B) \in \mathcal{M}_\Lambda(H)$, let $C_\Lambda(A,B)$ be the subset of $C_\Lambda(A,B)'$ consisting of balanced homomorphisms. The $C_\Lambda(A,B)$'s are clearly disjoint. Letting

$C_\Lambda(0)$ be the complement of the union of the $C_\Lambda(A,B)$'s, we have that $w_\Lambda(C_\Lambda(0)) \leq 2^{-\Omega(d)} Z_\Lambda(\mathbb{Z}_m^d, H)$ as it consists of the unbalanced homomorphisms removed from the $C_\Lambda(A,B)'$'s (a collection with total weight at most $\exp\{-(m-\kappa/2)^d/4\} Z_\Lambda(\mathbb{Z}_m^d, H)$, by (6.8)) together with a subset of $D_\Lambda(0)$ (with total weight at most $2^{-\Omega(d)} Z_\Lambda(\mathbb{Z}_m^d, H)$). This establishes that our partition satisfies statement 1 of Theorem 6.1.2.

Statement 2 is immediate from the construction of the $C_\Lambda(A,B)$'s. Statements 3 and 4 follow from the corresponding statements for the $C_\Lambda(A,B)'$'s, since the sizes of $C_\Lambda(A,B)'$ and $C_\Lambda(A,B)$ differ by a multiplicative factor of no more than $1 \pm 2^{-\Omega(d)}$. Finally, statement 5 follows from (6.7) for the same reason.

We now begin the verification of (6.7) and (6.8), beginning with (6.7). Fix $(A,B) \in \mathcal{M}_\Lambda(H)$, $x \in \mathcal{E}$ and $k \in A$ (the case $y \in \mathcal{O}$ and $\ell \in B$ is analogous). If (F_1,F_2) is such that $x \notin F_1 \cup N(F_2)$, then since x is adjacent to vertices colored from B, and all vertices of A are adjacent to all vertices of B, we have the following: for f chosen from $C_\Lambda^{(F_1,F_2)}(A,B)'$ according to p_Λ, the probability that $f(x)=k$ is exactly λ_k/λ_A. Thus (6.7) will follow if we can show that the contribution to $w_\Lambda(C_\Lambda(A,B)')$ from those $C_\Lambda^{(F_1,F_2)}(A,B)'$'s with $x \in F_1 \cup N(F_2)$ is at most $2^{-\Omega(d)} w_\Lambda(C_\Lambda(A,B)')$. To establish this, note that

$$
\begin{aligned}
& \sum_{(F_1,F_2)} w_\Lambda(C_\Lambda^{(F_1,F_2)}(A,B)') \mathbf{1}_{\{x \in F_1 \cup N(F_2)\}} \\
= \ & \frac{1}{m^d} \sum_{y \in \mathcal{E}} \sum_{(F_1,F_2)} w_\Lambda(C_\Lambda^{(F_1,F_2)}(A,B)') \mathbf{1}_{\{y \in F_1 \cup N(F_2)\}} \\
\leq \ & \frac{1}{m^d} \sum_{(F_1,F_2)} |F_1 \cup N(F_2)| w_\Lambda(C_\Lambda^{(F_1,F_2)}(A,B)') \\
\leq \ & \frac{(2d+1)(m-\kappa)^d}{m^d} w_\Lambda(C_\Lambda(A,B)').
\end{aligned}
$$

The first equality follows from the symmetry of both $\mathbb{Z}_m^d$ and the construction of $C_\Lambda(A,B)'$. In the first inequality we reverse the order of summation, and in the second we bound $|F_1 \cup N(F_2)|$ by $(2d+1)(m-\kappa)^d$.

Now we consider (6.8), and so we again fix $(A,B) \in \mathcal{M}_\Lambda(H)$. A lower bound on $w_\Lambda(C_\Lambda^{(F_1,F_2)}(A,B)')$ (for $C_\Lambda^{(F_1,F_2)}(A,B)' \neq \emptyset$) is

$$\lambda_A^{m^d/2-|F_1 \cup N(F_2)|} \lambda_B^{m^d/2-|F_2 \cup N(F_1)|}. \tag{6.9}$$

As before, this is because every vertex in $\mathcal{E} \setminus F_1 \cup N(F_2)$ is adjacent only to vertices colored only from B and so may be given any color from A, with a similar argument for vertices from $\mathcal{O} \setminus F_2 \cup N(F_1)$ (note that in this lower bound we are using the assumption $\lambda_i \geq 1$ for all i).

For $\delta > 0$, an upper bound on the sum of the weights of those $f \in C_\Lambda^{(F_1,F_2)}(A,B)'$ in which a particular color k from A appears either on a proportion less than $(\lambda_k/\lambda_A - \delta)$ of $\mathcal{E}$, or on a proportion greater than $(\lambda_k/\lambda_A + \delta)$, is

$$\left(\sum_{\substack{i \leq (\lambda_k/\lambda_A - \delta)m^d/2 \\ i \geq (\lambda_k/\lambda_A + \delta)m^d/2}} \binom{m^d/2}{i} (\lambda_A - \lambda_k)^{m^d/2-i} \lambda_k^i \right) \lambda_B^{m^d/2} \lambda_H^{|F_1 \cup N(F_2)|+|F_2 \cup N(F_1)|}. \tag{6.10}$$

By standard Binomial concentration inequalities (see for example [34] or [3, Appendix A], we have

$$\sum_{\substack{i \leq (\lambda_k/\lambda_A - \delta)m^d/2 \\ i \geq (\lambda_k/\lambda_A + \delta)m^d/2}} \binom{m^d/2}{i} (\lambda_A - \lambda_k)^{m^d/2-i} \lambda_k^i \leq 2 \exp\left\{-\delta^2 m^d/2\right\} \lambda_A^{m^d/2}. \tag{6.11}$$

Now combining (6.9), (6.10) and (6.11) we find that for f chosen from non-empty $C_\Lambda^{(F_1,F_2)}(A,B)'$ according to p_Λ, the probability that a particular color appears either on a proportion less than $(\lambda_k/\lambda_A - \delta)$ of $\mathcal{E}$ or on a proportion greater than $(\lambda_k/\lambda_A + \delta)$ is at most

$$\frac{2\lambda_H^{2|F_1 \cup N(F_2)|+2|F_2 \cup N(F_1)|}}{\exp\{\delta^2 m^d/2\}} \leq \exp\left\{-\delta^2 m^d/2 + O(d(m-\kappa)^d)\right\}$$

(again using $\lambda_i \geq 1$ for all i as well as our upper bound on $|F_1| + |F_2|$). Repeating this argument for colors from B and applying the law of total probability and a union bound, we find that for f chosen from $C_\Lambda(A, B)'$ according to p_Λ, the probability that either there is some color k from A which fails to appear on a proportion between $(\lambda_k/\lambda_A - \delta)$ and $(\lambda_k/\lambda_A + \delta)$ of $\mathcal{E}$, or there is some color ℓ from B which fails to appear on a proportion between $(\lambda_\ell/\lambda_B - \delta)$ and $(\lambda_\ell/\lambda_B + \delta)$ of $\mathcal{O}$ is at most $\exp\{-\delta^2 m^d/2 + O(d(m-\kappa)^d)\}$. Taking $\delta = (1 - \kappa/(4m))^d$ gives the required result.

6.4 Proof of Theorem 6.1.4

Our strategy is to put an upper bound on the entropy of a uniformly chosen element of $\mathrm{Hom}(\mathbb{Z}_m^d, H)$ that is smaller than a trivial lower bound unless ε is suitably small. We build on ideas introduced by Kahn [40].

6.4.1 Entropy

In this section we very briefly review the entropy material that is relevant for the proof of Theorem 6.1.4. See [40] for an expanded treatment appropriate to the present application, or for example [50] for a very thorough discussion. In what follows, X, Y, etc. are discrete random variables, taking values in any finite set. Throughout, we take $\log = \log_2$.

The (binary) entropy function is $H(\alpha) = -\alpha \log \alpha - (1-\alpha)\log(1-\alpha)$. The *entropy* of the random variable X is $H(X) = \sum_x -p(x)\log p(x)$ where we write $p(x)$ for $\Pr(X = x)$ (and later $p(x|y)$ for $\Pr(X = x|Y = y)$). The inequality that makes entropy a useful tool for counting is

$$H(X) \leq \log|\mathrm{range}(X)|, \tag{6.12}$$

with equality if and only if X is uniform. For random variables X, Y and Z where

Y determines Z, we also have

$$H(X|Y) \leq H(X) \quad \text{and} \quad H(X|Y) \leq H(X|Z), \tag{6.13}$$

that is, dropping or lessening conditioning does not decrease entropy (here $H(X|Y) = \sum_y p(y) \sum_x -p(x|y) \log p(x|y)$ is a conditional entropy). We will also use the (conditional) chain rule: for $X = (X_1, \ldots, X_n)$ a random vector,

$$H(X|Y) = H(X_1|Y) + H(X_2|X_1, Y) + \cdots + H(X_n|X_1, \ldots, X_{n-1}, Y). \tag{6.14}$$

Finally, we will need the conditional version of Shearer's lemma from [40] (extending the original Shearer's lemma from [13]). For a random vector $X = (X_1, \ldots, X_m)$ and $A \subseteq [m] := \{1, \ldots, m\}$, set $X_A = (X_i : i \in A)$.

Lemma 6.4.1. *Let $X = (X_1, \ldots, X_m)$ be a random vector and $\mathcal{A}$ a collection of subsets (possibly with repeats) of $[m]$, with each element of $[m]$ contained in at least t members of $\mathcal{A}$. Then, for any partial order $\prec$ on $[m]$,*

$$H(X) \leq \frac{1}{t} \sum_{A \in \mathcal{A}} H(X_A|(X_i : i \prec A)),$$

where $i \prec A$ means $i \prec a$ for all $a \in A$.

6.4.2 Notation and definitions

It will be convenient to gather together all of our technical notation in a single place. We will also utilize the notation from Section 2.3. Recall that

$$\eta(H) = \max\{|A||B| \colon A, B \subseteq V(H),\ A \sim B\}$$

and

$$\mathcal{M}(H) = \{(A,B)\colon A,B \subseteq V(H),\ A \sim B,\ |A||B| = \eta(H)\}.$$

Define

$$\mathcal{S}(H) = \{A\colon (A,B) \in \mathcal{M}(H) \text{ for some } B\}.$$

For $A \subseteq V(H)$ let $n(A) = \{v \in V(H)\colon \{v\} \sim A\}$, and for $A, B \subseteq V(H)$ let $p(A,B)$ be the number of pairs $(a,b) \in A \times B$ with $a \nsim b$. Let

$$V^\star = \{x = (x_1, \ldots, x_d) \in V\colon x_d = 0,\ x \in \mathcal{E}\}$$

(a set of size $m^{d-1}/2$). For each $v \in V^\star$ set

$$\mathcal{C}(v) = \{v + (0, \ldots, 0, i)\colon 0 \leq i \leq m-1\}.$$

In other words, $\mathcal{C}(v)$ is the set of all vertices in V which agree with v on the first $d-1$ coordinates; note that unless $m = 2$, $\mathcal{C}(v)$ induces a cycle in $\mathbb{Z}_m^d$. (In the case $m = 2$, $\mathcal{C}(v)$ simply induces an edge; this slight difference between $m = 2$ and $m \geq 4$ is something that has to be accommodated throughout the proof.) Throughout the proof we think of $\mathcal{C}(v)$ as an *ordered* tuple of vectors $(v_0, v_1, \ldots, v_{m-1})$ with each $v_i = v + (0, \ldots, 0, i)$.

For $u \in \mathcal{C}(v)$ for some $v \in V^\star$, let $u'_+ = u + (0, \ldots, 0, 1)$ and $u'_- = u - (0, \ldots, 0, 1)$ (so $u'_+ = u'_-$ if and only if $m = 2$), and set

$$M_u = N(u) \setminus \{u'_+, u'_-\}$$

and

$$M_{\mathcal{C}(v)} = M_{v_0} \cup \cdots \cup M_{v_{m-1}}.$$

A key observation that drives our proof is that the subgraph of $\mathbb{Z}_m^d$ induced by $M_{\mathcal{C}(v)}$

is a disjoint union of $2d-2$ cycles of length m (when $m \geq 4$) or of $d-1$ disjoint edges (when $m=2$); this significantly restricts the appearance of an H-coloring on $M_{\mathcal{C}(v)}$ given its appearance on $\mathcal{C}(v)$.

To each $v \in V^\star$ with $|v| \geq 2m$ (where $|\cdot|$ indicates the sum of the coordinates) associate a $w(v) \in V^\star$ with $|w(v)| = |v| - 2m$ and with $w(v) < v$ in the usual component-wise partial order on $\mathbb{Z}^d$. For $|v| < 2m$ we do not define a $w(v)$, but it will prove convenient to adopt the convention in this case that $M_w = \emptyset$. From now on, whenever w appears, it will be $w(v)$ for whatever $v \in V^\star$ is under consideration.

We will use $(A_0, \ldots, A_{m-1})$ to indicate a tuple with each $A_i \subseteq V(H)$, and when $(A_0, \ldots, A_{m-1})$ appears as a range of summation it will vary over all possible such tuples. We will use $\mathrm{alt}(A, B)$ for the tuple $(A, B, \ldots, A, B)$, and $n(A_0, \ldots, A_{m-1})$ for the tuple $(n(A_0), \ldots, n(A_{m-1}))$. We denote by $g(A_0, \ldots, A_{m-1})$ the number of ways of choosing $(x_0, \ldots, x_{m-1})$ with $x_i \in A_i$ for each i and with $x_0 \sim \cdots \sim x_{m-1} \sim x_0$ (that is, with the x_i's, taken consecutively, forming a cycle).

6.4.3 Events and probabilities

Now let f be uniformly chosen from $\mathrm{Hom}(\mathbb{Z}_m^d, H)$. We define a number of events in the associated probability space. For $A \subseteq V(H)$ and $v \in V^\star$, let

$$Q_{v,A} = \{f(N(v)) = A\},$$

$$R_{v,A} = \{f(M_v) = A\},$$

$$Q_{\mathcal{C}(v),(A_0,\ldots,A_{m-1})} = \cap_{i=0}^{m-1} Q_{v_i,A_i}$$

and

$$R_{\mathcal{C}(v),(A_0,\ldots,A_{m-1})} = \cap_{i=0}^{m-1} R_{v_i,A_i}.$$

To denote the probability of each of these events, we will replace the leading upper case letter with the corresponding lower case letter; so, for example,

$$q_{v,A} = \Pr\left(Q_{v,A}\right).$$

For $u \in \mathcal{C}(v)$ for some $v \in V^\star$ let $R_u = \{f(y) : y \in M_u\}$ be the random variable indicating the palette of colors used on M_u, and let

$$T_{\mathcal{C}(v)} = (R_{v_0}, \ldots, R_{v_{m-1}}).$$

Finally, define ε (depending on d, m and H, but by the symmetry of $\mathbb{Z}_m^d$ independent of v) by

$$1 - \varepsilon = \sum_{(A,B)\in\mathcal{M}(H)} r_{\mathcal{C}(v),\mathrm{alt}(A,B)}.$$

6.4.4 A partial order on V

For $0 \leq k \leq (m-1)(d-1)$, let

$$L_k = \left\{ x \in V \colon \sum_{i=1}^{d-1} x_i = k \right\}.$$

We refer to the L_k's as the *levels* of V; note that they partition V. Following the approach of [40], we wish to put a partial order on V that satisfies (6.15) and (6.16) below. We will achieve this by putting an order $\prec$ on the indices of the levels, as follows. Begin by ordering the odd natural numbers in the usual order, up to $m-1$. Next put 0, then $m+1$, then 2, then $m+3$, etc., interleaving the standard order of the evens and the odds. This order for $m = 2$ is used in [40], and begins $1 \prec 0 \prec 3 \prec 2 \prec 5 \prec 4 \ldots$. For $m = 4$, it begins $1 \prec 3 \prec 0 \prec 5 \prec 2 \prec 7 \prec 4 \ldots$, and for $m = 6$ it begins $1 \prec 2 \prec 5 \prec 0 \prec 7 \prec 2 \prec 9 \prec 4 \prec \ldots$. These orders are constructed specifically to satisfy the follow property: for each even $i \in \mathbb{N}$ we have $x \prec i$ for all $x \in X_i$ and

$y \prec x$ for all $x \in X_i$ and $y \in Y_i$, where $X_i = \{i-m+1, i-1, i+1, i+m-1\} \cap \mathbb{N}$ (or $\{i-1, i+1\} \cap \mathbb{N}$ if $m = 2$) and $Y_i = \{i-3m+1, i-2m-1, i-2m+1, i-m-1\} \cap \mathbb{N}$ (or $\{i-5, i-3\} \cap \mathbb{N}$ if $m = 2$).

We use $\prec$ to obtain a partial order (which we shall also call $\prec$) on V by declaring $L_i \prec L_j$ if and only if $i \prec j$. This partial order has two properties that will be critically important for us. For the first of these, note that for $v \in V^\star$, if $v \in L_i$ for some i (necessarily even), then $\mathcal{C}(v) \subseteq L_i$ and $M_{\mathcal{C}(v)} \subseteq \cup_{x \in X_i} L_x$, and so

$$M_{\mathcal{C}(v)} \subseteq \{x \colon x \prec \mathcal{C}(v)\}. \tag{6.15}$$

For the second property, note that since $M_w \subseteq \cup_{y \in Y_i} L_y$ for $v \in L_i$ we have

$$M_w \subseteq \{x \colon x \prec M_{\mathcal{C}(v)}\}. \tag{6.16}$$

6.4.5 The proof of Theorem 6.1.4

We will show that $\varepsilon < 2^{-\Omega(d)}$ (with the implicit constant depending on m and H). From this, Theorem 6.1.4 follows. To see this, first observe that for $(A, B) \in \mathcal{M}(H)$ we have $Q_{\mathcal{C}(v),\mathrm{alt}(A,B)} \supseteq R_{\mathcal{C}(v),\mathrm{alt}(A,B)}$. Indeed, consider any $f \in R_{\mathcal{C}(v),\mathrm{alt}(A,B)}$. For each even i we must have $f(v_i) \sim a$ for all $a \in A$, and so since $(A, B) \in \mathcal{M}(H)$, we must have $f(v_i) \in B$; similarly, for odd i we must have $f(v_i) \in A$. It follows that

$$1 - \varepsilon \leq \sum_{(A,B) \in \mathcal{M}(H)} q_{\mathcal{C}(v),\mathrm{alt}(A,B)}.$$

Now let $e = xy$ be an edge of $\mathbb{Z}_m^d$; by symmetry we may assume that $e = v_0 v_1$ for some $v = v_0 \in V^\star$. The event that e is ideal contains the event $\cup_{(A,B) \in \mathcal{M}(H)} Q_{\mathcal{C}(v),\mathrm{alt}(A,B)}$ (a union of disjoint events), and so the probability that e is ideal is at least $1 - \varepsilon$.

To bound ε we consider the entropy $H(f)$ of an $f \in \mathrm{Hom}(\mathbb{Z}_m^d, H)$, chosen uniformly. We first put a trivial lower bound on $H(f)$:

$$H(f) = \log |\mathrm{Hom}(\mathbb{Z}_m^d, H)| \geq \frac{m^d}{2} \log \eta(H), \tag{6.17}$$

the equality from (6.12) and the inequality obtained by choosing any $(A, B) \in \mathcal{M}(H)$ and considering only pure-(A, B) colorings (as defined in Section 6.1). The bulk of the proof will be devoted to finding an upper bound on $H(f)$ which, for ε too large, is smaller than this trivial lower bound.

We will upper bound $H(f)$ by an application of Shearer's lemma (with conditioning), that is, Lemma 6.4.1. For $m \geq 4$, we take as our covering family $\{M_{\mathcal{C}(v)} \colon v \in V^\star\}$ together with $2d-2$ copies of $\mathcal{C}(v)$ for each $v \in V^\star$. For $m = 2$ we take $\{M_{\mathcal{C}(v)} \colon v \in V^\star\}$ together with $d-1$ copies of $\mathcal{C}(v)$ for each $v \in V^\star$. Each vertex of $\mathbb{Z}_m^d$ is covered $2d-2$ times by this family (in the case $m \geq 4$) or $d-1$ times (in the case $m = 2$) and so, bearing (6.13), (6.15) and (6.16) in mind we have

$$H(f) \leq \sum_{v \in V^\star} H(f\restriction_{\mathcal{C}(v)} | f\restriction_{M_{\mathcal{C}(v)}}) + \left(\frac{1 + \mathbf{1}_{\{m=2\}}}{2d-2}\right) \sum_{v \in V^\star} H(f\restriction_{M_{\mathcal{C}(v)}} | f\restriction_{M_w}), \tag{6.18}$$

where $f\restriction_S$ denotes the restriction of f to the set $S \subseteq V$ (note that this is our only use of the order $\prec$). For the first term on the right-hand side of (6.18) we expand out the conditional entropy and use (6.12) to get

$$\begin{aligned}
& H(f\restriction_{\mathcal{C}(v)} | f\restriction_{M_{\mathcal{C}(v)}}) \\
\leq & \sum_{(A_0,\ldots,A_{m-1})} r_{\mathcal{C}(v),(A_0,\ldots,A_{m-1})} H\left(f(\mathcal{C}(v)) | \left\{T_{\mathcal{C}(v)} = (A_0, \ldots, A_{m-1})\right\}\right) \\
\leq & \sum_{(A_0,\ldots,A_{m-1})} r_{\mathcal{C}(v),(A_0,\ldots,A_{m-1})} \log\left(g(n(A_0, \ldots, A_{m-1}))\right). \qquad (6.19)
\end{aligned}$$

We now turn to the second term on the right-hand side of (6.18). For $|v| \leq 2m-1$

we use (6.12) to naively bound

$$H(f\restriction M_{\mathcal{C}(v)}|f\restriction M_w) \leq \left(\frac{2d-2}{1+\mathbf{1}_{\{m=2\}}}\right) m \log |V(H)|; \tag{6.20}$$

this will ultimately not be too costly since there are not too many such v. Specifically, the number of such v is exactly the number of vectors $(a_1, \ldots, a_{d-1}) \in \{0, \ldots, m-1\}^{d-1}$ with $\sum_{i=0}^{d-1} a_i \leq 2m-2$ and even; this is at most the number of solutions to $\sum_{i=0}^{d} a_i = 2m-2$ in non-negative integers, which is at most $\binom{2m+d-3}{2m-2}$.

For $|v| \geq 2m$ we use (6.13) and (6.14) to obtain

$$\begin{aligned} H(f\restriction M_{\mathcal{C}(v)}|f\restriction M_w) &\leq H(f\restriction M_{\mathcal{C}(v)}|R_w) \\ &= H(f\restriction M_{\mathcal{C}(v)}, T_{\mathcal{C}(v)}|R_w) \\ &\leq H(T_{\mathcal{C}(v)}|R_w) + H(f\restriction M_{\mathcal{C}(v)}|T_{\mathcal{C}(v)}), \end{aligned} \tag{6.21}$$

the equality holding since $f\restriction M_{\mathcal{C}(v)}$ determines $T_{\mathcal{C}(v)}$. For the second term on the right hand side of (6.21) we expand out the conditional entropy and then use (6.12) to get

$$\begin{aligned} & H(f\restriction M_{\mathcal{C}(v)}|T_{\mathcal{C}(v)}) \\ = & \sum_{(A_0,\ldots,A_{m-1})} r_{\mathcal{C}(v),(A_0,\ldots,A_{m-1})} H(f\restriction M_{\mathcal{C}(v)}|\{T_{\mathcal{C}(v)} = (A_0, \ldots, A_{m-1})\}) \\ \leq & \sum_{(A_0,\ldots,A_{m-1})} r_{\mathcal{C}(v),(A_0,\ldots,A_{m-1})} \left(\frac{2d-2}{1+\mathbf{1}_{\{m=2\}}}\right) \log(g(A_0, \ldots, A_{m-1})). \end{aligned} \tag{6.22}$$

Here we use that $M_{\mathcal{C}(v)}$ consists of $2d-2$ disjoint cycles (in the case $m \geq 4$) and $d-1$ disjoint edges (in the case $m = 2$).

Inserting (6.19), (6.20), (6.21) and (6.22) into (6.18), combining with (6.17), summing over $v \in V^\star$ (noting that $|V^\star| = m^{d-1}/2$) and using the symmetry of $\mathbb{Z}_m^d$ we obtain

$$
\begin{aligned}
& m \log \eta(H) \\
\leq \quad & \frac{2\binom{2m+d-3}{2m-2} \log |V(H)|}{m^{d-2}} + \left(\frac{1+\mathbf{1}_{\{m=2\}}}{2d-2} \right) H(T_{\mathcal{C}(v)}|R_w) \\
& + \sum_{(A_0,\ldots,A_{m-1})} r_{\mathcal{C}(v),(A_0,\ldots,A_{m-1})} \log \left(g(A_0,\ldots,A_{m-1}) g(n(A_0,\ldots,A_{m-1})) \right).
\end{aligned} \tag{6.23}
$$

We now focus on the sum on the right-hand side of (6.23). Using the trivial bound

$$g(A_0,\ldots,A_{m-1}) \leq \prod_{i=0}^{m-1} |A_i| \tag{6.24}$$

together with the observation that for any $(A,B) \in \mathcal{M}(H)$ we have $n(A) = B$ and $n(B) = A$, we have

$$g(\mathrm{alt}(A,B)) g(n(\mathrm{alt}(A,B))) \leq \eta(H)^m \tag{6.25}$$

for any such (A,B) (actually we have equality in (6.25), but we will not need it). On the other hand, we claim that if $(A_0,\ldots,A_{m-1})$ is not of the form $\mathrm{alt}(A,B)$ for some $(A,B) \in \mathcal{M}(H)$ then there is a constant $\delta(H) \geq 1$ such that

$$g(A_0,\ldots,A_{m-1}) g(n(A_0,\ldots,A_{m-1})) \leq \eta(H)^m - \delta(H). \tag{6.26}$$

To see this, note first that if there is an $A \in (A_0,\ldots,A_{m-1})$ with $A \notin \mathcal{S}(H)$, A_0 say, then from (6.24) we have

$$g(A_0,\ldots,A_{m-1}) g(n(A_0,\ldots,A_{m-1})) \leq \prod_{i=0}^{m-1} |A_i||n(A_i)|,$$

and since each of the terms in the product above is at most $\eta(H)$ and one ($|A_0||n(A_0)|$) is strictly less than $\eta(H)$, we get (6.26). So we may assume that $(A_0,\ldots,A_{m-1}) \in$

$\mathcal{S}(H)^m$, but is not of the form $\mathrm{alt}(A, B)$. Since $(A, B) \in \mathcal{M}(H)$ is equivalent to $A, B \in \mathcal{S}(H)$ and $A = n(B)$, $B = n(A)$, we may assume without loss of generality that $A_1 \neq n(A_0)$. We have

$$g(A_0, \ldots, A_{m-1}) \leq (|A_0||A_1| - p(A_0, A_1)) \prod_{i=2}^{m-1} |A_i|$$

and

$$g(n(A_0, \ldots, A_{m-1})) \leq (|n(A_0)||n(A_1)| - p(n(A_0), n(A_1))) \prod_{i=2}^{m-1} |n(A_i)|.$$

If one of $p(A_0, A_1)$, $p(n(A_0), n(A_1))$ is non-zero, then as before the product of these two bounds is strictly less than $\eta(H)^m$, giving (6.26) in this case. If they are both 0 then we have $A_0 \sim A_1$ and $n(A_0) \sim n(A_1)$, so $A_1 \subseteq n(A_0)$ and $n(A_0) \subseteq A_1$, so $A_1 = n(A_0)$, a contradiction.

Recalling the definition of ε, together (6.25) and (6.26) yield

$$\begin{aligned}
& \sum_{(A_0,\ldots,A_{m-1})} r_{\mathcal{C}(v),(A_0,\ldots,A_{m-1})} \log\left(g(A_0, \ldots, A_{m-1}) g(n(A_0, \ldots, A_{m-1}))\right) \\
\leq \;& \varepsilon \log(\eta(H)^m - \delta(H)) + (1 - \varepsilon) \log \eta(H)^m \\
= \;& m \log \eta(H) + \varepsilon \log\left(1 - \frac{\delta(H)}{\eta(H)^m}\right) \\
\leq \;& m \log \eta(H) - \frac{\varepsilon \delta(H) \log e}{\eta(H)^m}
\end{aligned}$$

(recall $\log = \log_2$). Inserting into (6.23) we get

$$\frac{\varepsilon \delta(H) \log e}{\eta(H)^m} \leq \frac{2\binom{2m+d-3}{2m-2} \log |V(H)|}{m^{d-2}} + \left(\frac{1 + \mathbf{1}_{\{m=2\}}}{2d-2}\right) H(T_{\mathcal{C}(v)}|R_w). \tag{6.27}$$

The final entropy term we need to analyze is $H(T_{\mathcal{C}(v)}|R_w)$. A naive upper bound from (6.12) is

$$H(T_{\mathcal{C}(v)}|R_w) \leq |V(H)|m,$$

the right-hand side being the logarithm of the size of the range of possible values. Inserting this into (6.27) we have

$$\frac{\varepsilon\delta(H)\log e}{\eta(H)^m} \le \frac{2\binom{2m+d-3}{2m-2}\log|V(H)|}{m^{d-2}} + \left(\frac{1+\mathbf{1}_{\{m=2\}}}{2d-2}\right)|V(H)|m, \tag{6.28}$$

showing that $\varepsilon \le c/d$ for some constant c depending on H and m.

The information that $\varepsilon = o(1)$ as $d \to \infty$ allows us to strengthen our bound on $H(T_{\mathcal{C}(v)}|R_w)$, via the following key lemma.

Lemma 6.4.2. *For any* $(A,B) \in \mathcal{M}(H)$*,*

$$\Pr(R_{\mathcal{C}(v),\mathrm{alt}(A,B)}|R_{w,A}) \ge 1 - \frac{(3m-1)\varepsilon}{r_{w,A}},$$

and also

$$\sum_{A\notin\mathcal{S}(H)} r_{w,A} \le \varepsilon.$$

Proof. Choose $w_1,\ldots,w_{2m-1} \in V^\star$ with $w < w_1 < \cdots < w_{2m-1} < v$ in the usual partial ordering of $\mathbb{Z}^d$. Then

$$\begin{aligned}\left(R_{\mathcal{C}(v),\mathrm{alt}(A,B)}\right)^c \cap R_{w,A} \quad\subset\quad & (R_{w,A}\cap(R_{w_1,B})^c) \cup (R_{w_1,B}\cap(R_{w_2,A})^c) \cup \cdots \\ & \cup \left(R_{w_{2m-1},B}\cap(R_{v_0,A})^c\right) \cup (R_{v_0,A}\cap(R_{v_1,B})^c)\cdots \\ & \cup \left(R_{v_{m-2},A}\cap\left(R_{v_{m-1},B}\right)^c\right),\end{aligned}$$

and each of the $3m-1$ events on the right hand side occurs with probability less that ε, by symmetry of $\mathbb{Z}_m^d$. Therefore

$$\begin{aligned}\Pr(\left(R_{\mathcal{C}(v),(A,B)}\right)^c|R_{w,A}) \quad&=\quad \frac{\Pr\left(\left(R_{\mathcal{C}(v),(A,B)}\right)^c\cap R_{w,A}\right)}{r_{w,A}} \\ &\le\quad \frac{(3m-1)\varepsilon}{r_{w,A}}.\end{aligned}$$

Also, $r_{w,A} \geq r_{\mathcal{C}(w),\mathrm{alt}(A,B)}$ implies

$$\sum_{A \in \mathcal{S}(H)} r_{w,A} \geq \sum_{A \in \mathcal{S}(H)} r_{\mathcal{C}(w),\mathrm{alt}(A,B)} = \sum_{(A,B) \in \mathcal{M}(H)} r_{\mathcal{C}(w),\mathrm{alt}(A,B)} = 1 - \varepsilon.$$

$\square$

We now partition $\mathcal{S}(H)$ by $\mathcal{S}(H) = \mathcal{S}_1(H) \cup \mathcal{S}_2(H)$, where $A \in \mathcal{S}_1(H)$ if and only if $r_{w,A} \leq 2(3m-1)\varepsilon$ (note that this partition depends on d as well as on H, and for fixed m and H it may change for different values of d). For convenience we also write $\mathcal{S}_0(H)$ for the complement of $\mathcal{S}(H)$ (in the power set of $V(H)$). Expanding out the conditional entropy we have

$$H(T_{\mathcal{C}(v)}|R_w) = \sum_{i=0}^{2} \sum_{A \in \mathcal{S}_i(H)} r_{w,A} H(T_{\mathcal{C}(v)}|R_{w,A}).$$

Trivially (from (6.12) and the second statement of Lemma 6.4.2),

$$\sum_{A \in \mathcal{S}_0(H)} r_{w,A} H(T_{\mathcal{C}(v)}|R_{w,A}) \leq \varepsilon |V(H)| m. \tag{6.29}$$

For the remaining two terms of the sum, we need to do a little groundwork. For each A, $-H(T_{\mathcal{C}(v)}|R_{w,A})$ is the sum over all $(A_0, \ldots, A_{m-1})$ of

$$\Pr(\{T_{\mathcal{C}(v)} = (A_0, \ldots, A_{m-1})\}|R_{w,A}) \log \left(\Pr(\{T_{\mathcal{C}(v)} = (A_0, \ldots, A_{m-1})\}|R_{w,A}) \right)$$

(by definition of entropy) and so

$$H(T_{\mathcal{C}(v)}|R_{w,A}) \leq \sum_{(A_0,\ldots,A_{m-1})} H\left(\Pr(\{T_{\mathcal{C}(v)} = (A_0, \ldots, A_{m-1})\}|R_{w,A})\right). \tag{6.30}$$

For $A \in \mathcal{S}_1(H)$, we cannot do any better than bounding all $2^{|V(H)|m}$ entropy terms in (6.30) by 1, leading to

$$\begin{aligned}\sum_{A\in\mathcal{S}_1(H)} r_{w,A}H(T_{\mathcal{C}(v)}|R_{w,A}) &\leq 2^{|V(H)|m}\sum_{A\in\mathcal{S}_1(H)} r_{w,A} \\ &\leq 2(3m-1)2^{|V(H)|(m+1)}\varepsilon, \qquad (6.31)\end{aligned}$$

since there are at most $2^{|V(H)|}$ summands and each is at most $2(3m-1)\varepsilon$. For $A \in \mathcal{S}_2(H)$, on the other hand, we know by Lemma 6.4.2 and the definition of $\mathcal{S}_2(H)$ that

$$\Pr(\{T_{\mathcal{C}(v)} = (A_0,\ldots,A_{m-1})\}|R_{w,A}) \leq \frac{(3m-1)\varepsilon}{r_{w,A}} \leq \frac{1}{2}$$

if $(A_0,\ldots,A_{m-1}) \neq \mathrm{alt}(A,B)$, while

$$\Pr(\{T_{\mathcal{C}(v)} = (A_0,\ldots,A_{m-1})\}|R_{w,A}) \geq 1 - \frac{(3m-1)\varepsilon}{r_{w,A}} \geq \frac{1}{2}$$

if $(A_0,\ldots,A_{m-1}) = \mathrm{alt}(A,B)$. We may therefore replace each of the entropy terms in (6.30) by $H((3m-1)\varepsilon/r_{w,A})$, leading to

$$\begin{aligned}&\sum_{A\in\mathcal{S}_2(H)} r_{w,A}H(T_{\mathcal{C}(v)}|R_{w,A}) \\ \leq\ & 2^{|V(H)|m}\sum_{A\in\mathcal{S}_2(H)} r_{w,A}H\left(\frac{(3m-1)\varepsilon}{r_{w,A}}\right) \\ \leq\ & 2^{|V(H)|m}\left(\sum_{A\in\mathcal{S}_2(H)} r_{w,A}\right)H\left(\frac{|\mathcal{S}_2(H)|(3m-1)\varepsilon}{\sum_{A\in\mathcal{S}_2(H)} r_{w,A}}\right) \qquad (6.32)\end{aligned}$$

with (6.32) an application of Jensen's inequality. Now we use the fact that $\varepsilon \leq c/d$ to conclude that the argument of the entropy term in (6.32) is bounded above by $C\varepsilon$ for some constant depending on m and H (this utilizes Lemma 6.4.2 and the fact that $\sum_{A\in\mathcal{S}_1(H)} r_{w,A}$ is at most $c\varepsilon$) to get

$$\sum_{A\in\mathcal{S}_2(H)} r_{w,A} H(T_{\mathcal{C}(v)}|R_{w,A}) \leq CH(C\varepsilon). \tag{6.33}$$

We now combine (6.29), (6.31) and (6.33) with (6.27) to find that there are constants $c_i, i = 1, \ldots, 4$ (all depending on both m and H) such that

$$c_1\varepsilon \leq \frac{d^{c_2}}{m^d} + \frac{c_3 H(c_4\varepsilon)}{d}.$$

Using $H(x) \leq 2x\log(1/x)$ for $x \leq 1/2$ (a simple power series argument) this becomes

$$c_1\varepsilon \leq \frac{d^{c_2}}{m^d} + \frac{c_3\varepsilon}{d}\log\frac{1}{c_4\varepsilon}, \tag{6.34}$$

from which it follows that $\varepsilon \leq 2^{-\Omega(d)}$.

BIBLIOGRAPHY

1. J. Alexander, J. Cutler, and T. Mink, *Independent sets in graphs with given minimum degree,* Electron. J. Combin. **19** (2012), #P37.

2. N. Alon, *Independent sets in regular graphs and sum-free subsets of finite groups,* Israel J. Math. **73** (1991), 247-256.

3. N. Alon and J. Spencer, *The Probabilistic Method,* Wiley, New York, 2000.

4. I. Benjamini, O. Häggström and E. Mossel, *On random graph homomorphisms into $\mathbb{Z}$,* J. Comb. Theory Ser. B **78** (2000), 86-114.

5. B. Bollobás, *Modern Graph Theory,* Springer, New York, 1998.

6. B. Bollobás and I. Leader, *Edge-isoperimetric inequalities in the grid,* Combinatorica **11** (1991), 299-314.

7. C. Borgs, J. Chayes, M. Dyer and P. Tetali, *On the sampling problem for H-colorings on the hypercubic lattice,* Graphs, Morphisms, and Statistical Physics, DIMACS Ser. Disc. Math. Theoret. Comput. Sci. **63** (2004), 13-28.

8. C. Borgs, J. Chayes, A. Frieze, J. Kim, P. Tetali, E. Vigoda and V. Vu, *Torpid Mixing of Some Monte Carlo Chain Algorithms in Statistical Physics,* Proc. IEEE FOCS (1999), 218-229.

9. G. Brightwell and P. Winkler, *Graph homomorphisms and phase transitions,* J. Combin. Theory Ser. B **77** (1999), 221-262.

10. G. Brightwell and P. Winkler, *Hard constraints and the Bethe lattice: adventures at the interface of combinatorics and statistical physics,* Proc. Int'l. Congress of Mathematicians Vol. III (Li Tatsien, ed.), Higher Education Press, Beijing (2002), 605-624.

11. G. Brightwell and P. Winkler, *Random colorings of a Cayley tree,* Bolyai Soc. Math. Stud. **1** (2002) Contemporary Combinatorics, 247-276.

12. T. Carroll, D. Galvin, and P. Tetali, *Matchings and independent sets of a fixed size in regular graphs,* J. Combin. Theory Ser. A **116** (2009), 1219-1227.

13. F. Chung, P. Frankl, R. Graham, and J. Shearer, *Some intersection theorems for ordered sets and graphs,* J. Combin. Theory Ser. A **48** (1986), 23-37.

14. J. Cutler and A. J. Radcliffe, *Extremal problems for independent set enumeration,* Electron. J. Combin. **18(1)** (2011) #P169.

15. J. Cutler and A. J. Radcliffe, *Extremal graphs for homomorphisms,* J. Graph Theory **67** (2011), 261-284.

16. J. Cutler and A. J. Radcliffe, *Extremal graphs for homomorphisms II,* manuscript.

17. R. Diestel, *Graph Theory,* Springer, Heidelberg, 2005.

18. R. Dobrushin, *Prescribing a system of random variables by the help of conditional distributions,* Theory Probab. Appl. **15** (1970), 469-497.

19. J. Engbers, *Extremal H-colorings of graphs with fixed minimum degree,* manuscript.

20. J. Engbers and D. Galvin, *H-coloring bipartite graphs,* J. Combin. Theory Ser. B **102** (2012), 726-742.

21. J. Engbers and D. Galvin, *H-coloring tori,* J. Combin. Theory Ser. B **102** (2012), 1110-1133.

22. J. Engbers and D. Galvin, *Counting independent sets of a fixed size in graphs with given minimum degree,* submitted.

23. D. Galvin, *On homomorphisms from the Hamming cube to $\mathbb{Z}$,* Israel J. Math. **138** (2003), 189-213.

24. D. Galvin, *Bounding the partition function of spin systems,* Electron. J. Combin. **13** (2006), #R72.

25. D. Galvin, *Sampling* 3*-colourings of regular bipartite graphs,* Electron. J. Probab. **12** (2007), 481-497.

26. D. Galvin, *A threshold phenomenon for random independent sets in the discrete hypercube,* Combin. Probab. Comput. **20** (2011) 27-51.

27. D. Galvin, *Two problems on independent sets in graphs,* Discrete Math. **311** (2011), 2105-2112.

28. D. Galvin, *Maximizing H-colorings of a regular graph,* to appear in J. Graph Theory, doi: 10.1002/jgt.21658.

29. D. Galvin, *Counting colorings of a regular graph,* manuscript, arXiv:1205.2718.

30. D. Galvin and J. Kahn, *On phase transition in the hard-core model on $\mathbb{Z}^d$,* Combin. Probab. Comput. **13** (2004), 137-164.

31. D. Galvin, F. Martinelli, K. Ramanan, and P. Tetali, *The multi-state hard core model on a regular tree,* SIAM J. of Discrete Math. **25** (2011), 894-916.

32. D. Galvin and P. Tetali, *On weighted graph homomorphisms,* Graphs, Morphisms, and Statistical Physics, DIMACS Ser. in Discrete Math. Theoret. Comput. Sci. **63** (2005), 97-104.

33. D. Galvin and Y. Zhao, *The number of independent sets in a graph with small maximum degree,* Graphs Combin. **27** (2011), 177-186.

34. W. Hoeffding, *Probability inequalities for sums of bounded random variables,* J. Amer. Statist. Assoc. **58** (1963), 13-30.

35. G. Hopkins and W. Staton, *Some identites arising from the Fibonacci numbers of certain graphs,* Fib. Quart. **22** (1984), 255-258.

36. H. Hua, *A sharp upper bound for the number of stable sets in graphs with a given number of cut edges,* Appl. Math. Lett. **22** (2009), 1380-1385.

37. M. Jerrum, *A very simple algorithm for estimating the number of k-colourings of a low-degree graph,* Random Structures Algorithms **7** (1995), 157-165.

38. J. Jonasson, *Uniqueness of uniform random colorings of regular trees,* Statist. Probab. Lett. **57** (2002), 243-248.

39. J. Kahn, *An entropy approach to the hard-core model on bipartite graphs,* Combin. Probab. Comput. **10** (2001), 219-237.

40. J. Kahn, *Range of cube-indexed random walk,* Israel J. Math. **124** (2001), 189-201.

41. J. Kahn, *Entropy, independent sets and antichains: a new approach to Dedekind's problem,* Proc. Amer. Math. Soc. **130** (2002), 371-378.

42. G. Katona, *A theorem of finite sets,* Theory of graphs (Proc. Colloq., Tihany, 1966), Academic Press, New York (1968) 187-207.

43. A. Korshunov and A. Sapozhenko, *The number of binary codes with distance* 2, Probl. Kibern. **40** (1983), 111-130. (Russian)

44. J. Kruskal, *The number of simplices in a complex,* Mathematical optimization techniques, Univ. of Califoria Press, Berkeley, Calif. (1963), 251-278.

45. H.-F. Law and C. McDiarmid, *Independent sets in graphs with given minimum degree,* manuscript, arXiv:1210.1497.

46. F. Lazebnik, *On the greatest number of 2 and 3 colorings of a (v, e)-graph,* J. Graph Theory **13** (1989), 203-214.

47. C. Lin and S. Lin, *Trees and forests with large and small independent indices,* Chin. J. Math. **23** (1995), 199-210.

48. N. Linial, *Legal colorings of graphs,* Combinatorica **6** (1986), no. 1, 49-54.

49. P.-S. Loh, O. Pikhurko, and B. Sudakov, *Maximizing the number of q-colorings,* Proc. Lon. Math. Soc. **101** (2010), 655-696.

50. R.J. McEliece, *The Theory of Information and Coding,* Addison-Wesley, London, 1977.

51. A. Mazel and Y. Suhov, *Random surfaces with two-sided constraints: an application of the theory of dominant ground states,* J. Stat. Phys. **64** (1991), 111-134.

52. R. Merrifield and H. Simmons, *Topological Methods in Chemistry,* Wiley, New York, 1989.

53. P. O'Neil, *Asymptotics and random matrices with row-sum and column sum-restrictions,* Bull. Amer. Math. Soc. **75** (1969), 1276-1282.

54. A. Pedersen and P. Vestergaard, *Bounds on the number of vertex independent sets in a graph,* Taiwanese J. Math. **10** (2006), 1575-1587.

55. R. Peled, *High-Dimensional Lipschitz Functions are Typically Flat,* arXiv:1005.4636.

56. H. Prodinger and R. Tichy, *Fibonacci numbers of graphs,* Fib. Quart. **20** (1982), 16-21.

57. K. Ramanan, A. Sengupta, I. Ziedins, and P. Mitra, *Markov random fields of multicasting in tree networks,* Adv. in Appl. Probab. **34** (2002), 58-84.

58. J. Rowlinson and B. Widom, *New model for the study of liquid-vapor phase transitions,* J. Chem. Phys. **52** (1970), 1670-1684.

59. J. Salas and A. Sokal, *Absence of phase transition for antiferromagnetic Potts models via the Dobrushin uniqueness theorem,* J. Stat. Phys. **86** (1997), 551-579.

60. A.A. Sapozhenko, *On the number of independent sets in bipartite graphs with large minimum degree,* DIMACS Technical Report (2000), no. 2000-25.

61. E. Seneta, *Non-negative matrices and Markov chains,* Springer Series in Statistics no. 21, New York, 2006.

62. A. Sokal, *A personal list of unsolved problems concerning lattice gases and antiferromagnetic Potts models,* Markov Process. Related Fields **7** (2001), 21-38.

63. T. Tao, *An epsilon of room, I: real analysis: pages from year three of a mathematical blog,* Graduate Studies in Mathematics vol. 117, American Mathematical Society, Providence, RI, 2010.

64. H. Wilf, *Backtrack: an $O(1)$ expected time algorithm for the graph coloring problem,* Inform. Process. Lett. **18** (1984), no. 3, 119-121.

65. G. Wingard, *Properties and applications of the Fibonacci polynomial of a graph,* Ph.D. thesis, University of Mississippi, May 1995.

66. N. Wormold, *Models of random regular graphs,* Surveys in Combinatorics (J.D. Lamb and D.A. Preece, eds., London Math. Soc. Lecture Note Ser. **276**) (1999), 239-298.

67. Y. Zhao, *The number of independent sets in a regular graph,* Combin. Probab. Comput. **19** (2010), 315-320.

68. Y. Zhao, *The bipartite swapping trick on graph homomorphisms,* SIAM J. Discrete Math. **25** (2011), 660-680.

69. A.A. Zykov, *On some properties of linear complexes,* Mat. Sbornik N. S. **24(66)** (1949), 163-188.

This document was prepared & typeset with pdfLaTeX*, and formatted with* NDdiss2$_\varepsilon$ *classfile (v3.2013β[2013/01/31]) provided by Sameer Vijay and updated by Megan Patnott.*