

Variations on a Theme: Graph Homomorphisms

by

David E. Roberson

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Combinatorics & Optimization

Waterloo, Ontario, Canada, 2013

© David E. Roberson 2013

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

This thesis investigates three areas of the theory of graph homomorphisms: cores of graphs, the homomorphism order, and quantum homomorphisms.

A core of a graph X is a vertex minimal subgraph to which X admits a homomorphism. Hahn and Tardif have shown that, for vertex transitive graphs, the size of the core must divide the size of the graph. This motivates the following question: when can the vertex set of a vertex transitive graph be partitioned into sets which each induce a copy of its core? We show that normal Cayley graphs and vertex transitive graphs with cores half their size always admit such partitions. We also show that the vertex sets of vertex transitive graphs with cores less than half their size do not, in general, have such partitions.

Next we examine the restriction of the homomorphism order of graphs to line graphs. Our main focus is in comparing this restriction to the whole order. The primary tool we use in our investigation is that, as a consequence of Vizing's theorem, this partial order can be partitioned into intervals which can then be studied independently. We denote the line graph of X by $L(X)$. We show that for all $n \geq 2$, for any line graph Y strictly greater than the complete graph K_n , there exists a line graph X sitting strictly between K_n and Y . In contrast, we prove that there does not exist any connected line graph which sits strictly between $L(K_n)$ and K_n , for n odd. We refer to this property as being " n -maximal", and we show that any such line graph must be a core and the line graph of a regular graph of degree n .

Finally, we introduce quantum homomorphisms as a generalization of, and framework for, quantum colorings. Using quantum homomorphisms, we are able to define several other quantum parameters in addition to the previously defined quantum chromatic number. We also define two other parameters, projective rank and projective packing number, which satisfy a reciprocal relationship similar to that of fractional chromatic number and independence number, and are closely related to quantum homomorphisms. Using the projective packing number, we show that there exists a quantum homomorphism from X to Y if and only if the quantum independence number of a certain product graph achieves $|V(X)|$. This parallels a well known classical result, and allows us to construct examples of graphs whose independence and quantum independence numbers differ. Most importantly, we show that if there exists a quantum homomorphism from a graph X to a graph Y , then $\bar{\vartheta}(X) \leq \bar{\vartheta}(Y)$, where $\bar{\vartheta}$ denotes the Lovász theta function of the complement. We prove similar monotonicity results for projective rank and the projective packing number of the complement, as well as for two variants of $\bar{\vartheta}$. These immediately imply that all of these parameters lie between the quantum clique and quantum chromatic numbers, in particular yielding a quantum analog of the well known "sandwich theorem". We also briefly investigate the quantum homomorphism order of graphs.

Acknowledgements

First, I want to thank my advisor Chris Godsil. Without his guidance I would have never been able to make it through my graduate career. He has repeatedly pointed me towards many interesting questions and helpful resources. Without him this thesis would have not been possible.

Utmost thanks goes to my committee members Bruce Richter, Penny Haxell, John Watrous, and Claude Tardif. Their helpful comments and questions have significantly improved the quality of this thesis. I am gracious for the time and effort they contributed towards bettering this work.

Close friends have greatly enriched my time at the University of Waterloo. Without them my graduate experience would not have been nearly as rewarding as it was. I am deeply thankful for these friends as well as my family who has been very supportive of me through all these years.

Key to my success in the writing of this thesis was the support of Laura Mančinska. She kept me going through the difficult moments and times of stress. She also first introduced me to the problem of quantum colorings, from which a majority of this thesis sprang. I owe her a great debt of gratitude.

To my friends and family, and to Laura Mančinska especially.

Table of Contents

List of Figures	ix
1 Introduction	1
2 Background on Homomorphisms	5
2.1 Basic Definitions and Properties	6
2.2 Types of Homomorphisms	7
2.3 Graph Parameters Defined Via Homomorphisms	8
2.4 Homomorphic Equivalence and the Homomorphism Order	12
2.5 Cores	13
2.6 Graph Products and Other Constructions	17
2.7 Transitive Graphs	21
2.8 Cayley Graphs	25
2.9 Cores of Transitive Graphs	27
2.10 Lattice Properties of \mathcal{G}	31
2.11 Density and Universality	33
3 Cores of Vertex Transitive Graphs	37
3.1 Half-Sized Cores	38
3.2 Attempts at Generalization	41
3.3 Normal Cayley Graphs	42

3.4	Alternative Proof of Theorem 2.9.3	44
3.5	Remarks	45
4	The Homomorphism Order of Line Graphs	46
4.1	Basic Definitions and Properties	47
4.2	The Order	51
4.3	The Intervals	52
4.4	Density	57
4.5	Gaps	63
4.6	d -Maximal Graphs	70
4.7	Another Gap	76
4.8	Remarks and Open Questions	77
5	Quantum Information Background	80
5.1	Introduction to Quantum Theory	81
5.2	General States and Measurements	83
5.3	Shared Systems	88
6	Quantum Homomorphisms	94
6.1	Some Linear Algebra	97
6.2	Definition of Quantum Homomorphism	98
6.3	Measurement Graphs	103
6.4	Homomorphism Game	105
6.5	The Quantum Strategy	106
6.6	Basic Properties of Quantum Homomorphisms	111
6.7	Vector Colorings	114
6.8	The Homomorphic Product and Projective Packings	118
6.9	Defining Quantum Parameters	124

6.10 Quantum Chromatic Number	126
6.11 Quantum Independence Number	128
6.12 No Quantum No-Homomorphism Lemma	130
6.13 Projective Rank	132
6.14 Relation to $\bar{\vartheta}^+$ and Monotonicity of Projective Rank	136
6.15 Projective Rank of Some Special Graphs	139
6.16 Lattice Properties of Quantum Homomorphisms	142
6.17 The Quantum Homomorphism Order and Measurement Graphs	143
6.18 Discussion of Results	147
6.19 Open Questions	150
References	153
Index	159

List of Figures

2.1	Example of a homomorphism.	7
2.2	Homomorphism from C_5 to K_3	10
4.1	Examples of graphs and their line graphs.	48
4.2	The graph \hat{X} (left) and its line graph (right).	54
4.3	The Petersen graph (left) and its line graph (right).	54
4.4	The graph \hat{Y} (left) and its line graph (right).	57
4.5	A portion of the graphs Y (left) and $T(Y, y)$ (right).	59
4.6	Example of V_y and corresponding E_i for $i \in \{0, 1, \dots, d\}$ for $d = 5$	63
4.7	The graph M_3 (left) and its line graph (right).	66
4.8	The graph \hat{Z}	68
4.9	The graph M_4	69

Chapter 1

Introduction

A graph homomorphism is a map between the vertex sets of two graphs which preserves adjacency. From this simple definition springs forth a rich and varied field of study. In this thesis we will investigate the following three areas of the field of graph homomorphisms: cores of graphs, the homomorphism order, and quantum homomorphisms.

A core of a graph X is a vertex minimal subgraph of X to which X admits a homomorphism. In some sense, the core of a graph X is the smallest possible graph which retains all of the homomorphic information of X , such as chromatic number, clique number, etc. We will mainly focus on cores of graphs which have a high degree of symmetry, specifically those of vertex transitive graphs. For a vertex transitive graph X , a result of Hahn and Tardif [23] states that the number of vertices in the core of X must divide the number of vertices in X itself. This result raises the question: Can the vertex set of a vertex transitive graph be partitioned into sets which each induce copies of its core? It is precisely this question that we investigate in Chapter 3. We will see that, though in general the answer is no, there exist some interesting classes of graphs which do have such partitions. In particular, we show that any vertex transitive graph whose core is half its size must simply be two disjoint copies of its core along with some edges between the copies. We then see that this result cannot be extended to vertex transitive graphs whose core is less than half its size by exhibiting a family of vertex transitive graphs X_n for $n \geq 3$ such that the core of X_n is $1/n$ the size of X_n but none of the X_n can be partitioned into copies of their core. However, we are able to prove that an important class of Cayley graphs, known as normal Cayley graphs, always admit partitions into copies of their cores. The majority of the work in this chapter is from [43].

Using the composition of maps it is not hard to see that graph homomorphisms are

transitive, i.e. if X has a homomorphism to Y , and Y has a homomorphism to Z , then X has a homomorphism to Z . Moreover, the identity map is a homomorphism from X to itself for any graph X . This implies that the relation “has a homomorphism to” is a preorder. This relation is not a partial order since it is possible for graphs X and Y to have homomorphisms between them in both directions but not be isomorphic. However, if we define such graphs to be “homomorphically equivalent”, then the relation “has a homomorphism to” can be thought of as a partial order on classes of homomorphically equivalent graphs. This partial order is known as the “homomorphism order of graphs” and there has been a great deal of work on this order, including a book by Hell and Nešetřil [26] which gives an excellent introduction to the topic. In Chapter 4, we investigate the restriction of this order to line graphs. One interesting and useful feature of this restriction is that it can be partitioned into the intervals $[K_n, K_{n+1})$ consisting of the line graphs above K_n and strictly below K_{n+1} . This partition is a consequence of the fact that the clique and chromatic numbers of a line graph differ by at most one, which follows from Vizing’s theorem.

One of the main focuses of Chapter 4 is to compare the restriction to line graphs to the full homomorphism order of graphs. Our primary basis for comparison of these orders is a property known as density. A partial order is dense if for any two elements X and Z such that X is strictly less than Z in the partial order, there exists an element Y sitting strictly above X and strictly below Z in the order. The homomorphism order of graphs is known to be dense everywhere above the graph K_2 . In contrast to this, we show that when restricted to line graphs there are an infinite number of “gaps” in the order. A gap is essentially a counterexample to density. In other words, a gap is a pair of elements X and Z such that X is strictly less than Z but there does not exist any element Y strictly between them. The gaps we exhibit all lie between the graphs K_2 and K_3 , and so it is possible that the homomorphism order of line graphs is still dense everywhere above K_3 . However, we also show that for odd $n \geq 5$, there does not exist any *connected* line graph strictly above the line graph of K_n and strictly below K_n (which is a line graph). On the other hand, we prove that if Y is any line graph strictly above the graph K_n , then there exists a line graph X strictly between K_n and Y . Lastly, we show that if X is a connected graph whose line graph is strictly less than K_{n+1} , but no connected line graph lies strictly between the line graph of X and K_{n+1} , then X must be regular of degree n and its line graph is a core.

In Chapter 6 we introduce quantum homomorphisms, which were originally motivated by a game played between two players, Alice and Bob, and a referee. The purpose of the game is for Alice and Bob to convince the referee that they have a homomorphism from a graph X to a graph Y . To play, Alice and Bob each receive a vertex of X from the referee and then respond with a vertex of Y . Alice and Bob are not allowed to communicate

during the game, though they can agree on a strategy beforehand. To win, they must respond with the same vertex of Y if they were given the same vertex of X , and they must respond with adjacent vertices of Y if they were given adjacent vertices of X . It is easy to see that this game can be won with certainty if there exists a homomorphism from X to Y . It is not much more difficult to see that, for classical strategies, this condition is also necessary. Thus the existence of a perfect classical strategy of this game encodes the existence of a homomorphism.

It is known that allowing two separated parties to share, and perform measurements on, a quantum state can sometimes enable them to perform certain tasks that are impossible to perform classically. Considering this, it is natural to ask if the above described “homomorphism game” can sometimes be won by a quantum strategy even if no appropriate homomorphism exists. The answer turns out to be yes, as was shown in [6, 5]. Inspired by the equivalence between classical strategies and homomorphisms, we say that a graph X has a *quantum homomorphism* to Y if there exists a winning quantum strategy for the corresponding homomorphism game. It has been shown [8] that the existence of such a strategy is equivalent to the existence of an assignment of projectors to the elements of $V(X) \times V(Y)$ satisfying certain orthogonality and completeness conditions (in fact, we define quantum homomorphisms in terms of such assignments and then show that this is equivalent to the existence of a winning quantum strategy).

Fascinatingly, quantum homomorphisms seem to behave very similarly to homomorphisms. For instance, quantum homomorphisms are transitive, cannot map larger complete graphs to smaller ones, and a connected graph X has a quantum homomorphism to a graph Y if and only if it has a quantum homomorphism to one of its components. On the other hand, many graph parameters which are homomorphism monotone, i.e. $f(X) \leq f(Y)$ if X has a homomorphism to Y , are not quantum homomorphism monotone. Some examples are the chromatic and clique numbers. This is not surprising as these parameters are defined in terms of homomorphisms and not quantum homomorphisms. Somewhat surprisingly however, we are able to show that the Lovász theta function of the complement [34], which can also be defined via homomorphisms, is quantum homomorphism monotone. We also prove analogous theorems for two variants of the Lovász theta function, introduced respectively by Schrijver [47] and Szegedy [48].

Motivated by the definition of quantum homomorphisms, we define and investigate two graph parameters, $\tilde{\alpha}$ and ξ_f , which are based on assigning projectors to vertices of a graph such that adjacent vertices receive projectors which are orthogonal to each other. These parameters turn out to be somewhat reciprocal to each other similarly to how independence number and fractional chromatic number are related to one another. More specifically, we show that $|V(X)|/\tilde{\alpha}(X) \leq \xi_f(X)$ with equality if X is vertex transitive. We also prove

that ξ_f and $\tilde{\alpha}$ of the complement are quantum homomorphism monotone.

Using quantum homomorphisms, we define quantum analogs of several well known graph parameters. We primarily study the quantum independence number, denoted α_q , and the quantum chromatic number, denoted χ_q , which was previously defined in [3]. Using the parameter $\tilde{\alpha}$, we show that there exists a quantum homomorphism from X to Y if and only if the quantum independence number of a certain product graph achieves value $|V(X)|$. This parallels a well known classical result, allowing us to construct examples of graphs X such that $\alpha(X) \ll \alpha_q(X)$. We are also able to obtain effective bounds on α_q and χ_q using the monotonicity results mentioned above. These bounds allow us to show that the quantum odd girth of a Kneser graph is equal to its odd girth.

Lastly, we briefly discuss the quantum homomorphism order of graphs. We show that this partial order is a homomorphic image of the homomorphism order of graphs, and that it is isomorphic to an induced suborder of the homomorphism order of infinite graphs. We further show that the quantum homomorphism order of graphs is a lattice with the same meet and join operations as the homomorphism order. We also discuss what quantum homomorphism equivalence classes “look like”, noting that they must be unions of homomorphism equivalence classes. Several results of Chapter 6 are from a joint work with Laura Mančinska [42].

To prepare the reader for the main results of this thesis mentioned above, we have given the relevant background material on homomorphisms in Chapter 2, as well as an introduction to the basics of quantum information in Chapter 5.

Chapter 2

Background on Homomorphisms

The central notion of this thesis is the graph homomorphism. As graph homomorphisms have been extensively studied, the field contains a broad array of results and ideas. We have no hope of covering the whole of them here, but rather we will provide a subset designed to prepare the reader for the results presented in the remaining chapters. As this is background material, the results presented are not new, nor are they due to the author of this thesis. For a more thorough introduction to the theory of homomorphisms, we refer the reader to [19, 23, 26].

Though we have endeavored to be efficient, not all material presented here is necessary for our results. However, everything is intended to benefit the reader in their understanding of our work. The rest of this chapter is outlined as follows:

In Section 2.1, we introduce the basic definitions and a few simple properties of homomorphisms which will be ubiquitous throughout this thesis. Section 2.2 introduces the various special cases of homomorphisms one obtains by adding additional constraints, such as isomorphisms and retractions. We then give examples of some well known graph parameters which can be formulated naturally in terms of homomorphisms in Section 2.3. We also introduce the concept of homomorphism monotone parameters in this section.

Section 2.4 defines the equivalence relation known as “homomorphic equivalence”, and uses this to define the homomorphism order of graphs which is central to the ideas of Chapter 4. In Section 2.5 we introduce the notion of cores, which are the focus of Chapter 3, and prove several basic properties of these objects. We, in particular, focus on their relation to homomorphic equivalence. Cores will also play a role in some results of Chapter 4.

Section 2.6 introduces several well known graph products and other constructions that will be used throughout the thesis. We also give some details about each of these construc-

tions and discuss some of their basic properties. In Section 2.7, we define various types of transitivity such as vertex and arc transitivity. We also prove the no-homomorphism lemma and use it to prove the clique-coclique bound for vertex transitive graphs. In Section 2.8, we introduce Cayley graphs as natural examples of vertex transitive graphs, and prove a theorem of Sabidussi which states that any vertex transitive graph is homomorphically equivalent to some Cayley graph. We then use this result to give another proof of the clique-coclique bound.

We present several results on the cores of various types of transitive graphs in Section 2.9. In particular, we give a theorem of Hahn and Tardif which states that the size of the core of a vertex transitive graph must divide the size of the graph, as well as a result of Cameron showing that non-edge transitive graphs are either cores or have complete cores. We end the section by giving a proof that Kneser graphs are cores. Finally, Section 2.10 focuses on the homomorphism order of graphs. Specifically, we show that it is a lattice with meet and join operations corresponding to categorical product and disjoint union of graphs, respectively.

2.1 Basic Definitions and Properties

In this thesis, we will always use “graph” to mean “simple finite graph”, unless explicitly stated otherwise. Given graphs X and Y , a function $\varphi : V(X) \rightarrow V(Y)$ is a *homomorphism* if $\varphi(x) \sim \varphi(y)$ whenever $x \sim y$, where \sim denotes “is adjacent to”. Note that this definition implies that if $\varphi(x) = \varphi(y)$, then x and y are not adjacent in X . Thus the preimage $\varphi^{-1}(y)$ of a vertex in Y is necessarily an independent set in X . We will refer to the preimages $\varphi^{-1}(y)$ for $y \in V(Y)$ as the *fibres* of φ . We will use $X \rightarrow Y$ to denote the existence of a homomorphism from X to Y , and we write $\varphi : X \rightarrow Y$ to specify that φ is such a homomorphism. An example of a homomorphism has been given in Figure 2.1.

Note that if φ is a homomorphism from a graph X to a graph Y , it need not be surjective, i.e. there may be vertices in Y to which φ maps no vertices of X . We define the *image* of φ , denoted $\text{Im}(\varphi)$, to be the subgraph of Y induced by the vertices

$$\{y \in V(Y) : \exists x \in V(X) \text{ s.t. } \varphi(x) = y\}.$$

An important property of homomorphisms is that they are transitive: if $X \rightarrow Y$ and $Y \rightarrow Z$, then $X \rightarrow Z$. To see this, consider homomorphisms $\varphi_1 : X \rightarrow Y$ and $\varphi_2 : Y \rightarrow Z$, and adjacent vertices $x, y \in V(X)$. By definition, $\varphi_1(x) \sim \varphi_1(y)$, and similarly $\varphi_2(\varphi_1(x)) \sim \varphi_2(\varphi_1(y))$. Thus, $\varphi_2 \circ \varphi_1$ is a homomorphism from X to Z .




Figure 2.1: Example of a homomorphism.

2.2 Types of Homomorphisms

By imposing extra constraints on homomorphisms, we can obtain many other types of graph morphisms, some of which have been studied thoroughly in their own right. We take some time here to introduce some of these that will be used throughout this thesis.

An *isomorphism* is a bijective homomorphism which also preserves nonadjacency. Equivalently, a homomorphism is an isomorphism if and only if it is surjective and its inverse is also a homomorphism. If there exists an isomorphism from graph X to graph Y , then we say that X and Y are *isomorphic* and write $X \cong Y$. Note that the inverse of an isomorphism is also an isomorphism and thus X has an isomorphism to Y if and only if Y has an isomorphism to X . Isomorphisms have received attention mostly due to the graph isomorphism problem: deciding if two graphs are isomorphic to each other. This problem is one of only a few problems in NP which is neither known to be either polynomial time solvable nor NP-complete.

An *automorphism* is an isomorphism from a graph to itself. The set of automorphisms of a graph X , denoted $\text{Aut}(X)$, forms a group under composition. We will discuss the automorphism group of a graph in more detail when we introduce vertex transitive graphs, which will be the focus of Chapter 3.

An *endomorphism* of a graph X is a homomorphism from X to itself. These differ from automorphisms in that their image need not be all of X . Since the inverse of a homomorphism need not be a homomorphism (or even a function), the set of endomorphisms of a graph X , denoted $\text{End}(X)$, is not necessarily a group. However, $\text{End}(X)$ is a monoid (a set with an associative binary operation and an identity element) under composition. Note that $\text{Aut}(X) \subseteq \text{End}(X)$ as an automorphism is exactly an endomorphism

which is also an isomorphism. We will refer to endomorphisms which are not bijective as *proper endomorphisms*. Note that an endomorphism is an automorphism if and only if it is bijective.

A *retraction* is a homomorphism from X to a subgraph Y that acts as identity on its image, i.e. $\varphi : X \rightarrow Y$ is a retraction if $\varphi(\varphi(x)) = \varphi(x)$ for all $x \in V(X)$. For a homomorphism $\varphi : X \rightarrow Y$, we denote by $\varphi|_{X'}$ the restriction of φ to the subgraph X' of X . Using this notation, an endomorphism, φ , of X is a retraction if and only if $\varphi|_Y = \text{id}$ for Y the image of φ . We say that a subgraph Y of X is a *retract* if there exists a retraction whose image is Y . Note that, since a retraction acts as identity on its image, any retract of a graph X must be an *induced* subgraph of X . Endomorphisms and retractions are essential in the study of cores, which we will introduce in Section 2.5, and are the focus of Chapter 3 as well as an important part of Chapter 4.

2.3 Graph Parameters Defined Via Homomorphisms

One of the nice things about the theory of homomorphisms is that many graph parameters which are normally defined in a combinatorial way, also admit definitions purely in terms of homomorphisms. This in some sense unifies the study of these parameters as a special case of the study of homomorphisms. In this section we will introduce several such graph parameters, giving both the usual definition and the equivalent formulation in terms of homomorphisms. Many of the parameters we see will play an important part in the following chapters, in particular Chapter 6.

Probably one of the most well known graph parameters is the *chromatic number*. To define this parameter we must first define what a *coloring* of a graph is. For a graph X , a coloring of X is an assignment of “colors” from a given set to the vertices of X such that adjacent vertices receive different colors. An *n -coloring* is when the set of colors from which to choose has size n . More formally, a function $f : V(X) \rightarrow S$ is a coloring of X if $f(x) \neq f(y)$ whenever $x \sim y$, and f is an n -coloring if $|S| = n$. Typically, the set S is taken to be $[n] := \{1, 2, \dots, n\}$. We say that a graph X is *n -colorable* if there exists some n -coloring of X . The chromatic number of X , denoted $\chi(X)$, is then defined to be the minimum n such that X is n -colorable. For a given coloring $f : V(X) \rightarrow S$, we will refer to the sets $f^{-1}(s)$ for $s \in S$ as the *color classes* of f .

It turns out that an n -coloring is equivalent to a homomorphism to the complete graph K_n . To see this, note that inequality and adjacency are the same relation on $V(K_n)$, and thus a function $f : V(X) \rightarrow V(K_n)$ is an n -coloring if and only if it is a homomor-

phism. With this equivalence, we can define the chromatic number of a graph in terms of homomorphisms as follows:

$$\chi(X) = \min\{n \in \mathbb{N} : X \rightarrow K_n\}.$$

A *clique* of a graph X is a subset of $V(X)$ which induces a complete subgraph of X , i.e. a subset of vertices which are all pairwise adjacent. An *n-clique* is simply a clique of size n . The size of the largest clique in a graph X is known as the *clique number* of X , and is denoted $\omega(X)$. An *independent set* of a graph X is a subset of $V(X)$ such that no two vertices in the subset are adjacent. It is easy to see that a subset $S \subseteq V(X)$ is an independent set in X if and only if S is a clique in \overline{X} . One important property of independent sets, as they relate to homomorphisms, is that the inverse image of any independent set must be an independent set. If this were not the case, then adjacency would not be preserved. The *independence number* of a graph X , denoted $\alpha(X)$, is the size of the largest independent set in X . Both independent sets and cliques naturally arise in many different problems, and they are closely related to graph colorings as well.

To see how one can define $\omega(X)$ in terms of homomorphisms, consider a homomorphism $\varphi : K_n \rightarrow X$ for some graph X . Since every vertex of K_n is adjacent to every other vertex, no two of its vertices can be identified by φ . Furthermore, every pair of vertices in the image of φ must be adjacent for the same reason. Therefore the image of φ must be an n -clique. Using this, we can define both clique and independence numbers in terms of homomorphisms:

$$\begin{aligned}\omega(X) &= \max\{n \in \mathbb{N} : K_n \rightarrow X\} \\ \alpha(X) &= \omega(\overline{X}).\end{aligned}$$

The *odd girth* of a graph X , denoted $\text{og}(X)$, is the length of the shortest odd cycle in X . One can define odd girth in terms of homomorphisms from odd cycles, similarly to clique number. However, unlike complete graphs, the image of an odd cycle need not be an odd cycle of the same length. For instance, Figure 2.2 shows a homomorphism from C_5 to $C_3 = K_3$. On the other hand, the image of an odd cycle under a homomorphism must contain an odd cycle of equal or lesser length. To see this, note that if C is an odd cycle and φ is a homomorphism from C such that $\text{Im}(\varphi)$ does not contain any odd cycle, then $\text{Im}(\varphi)$ is bipartite and thus $\text{Im}(\varphi) \rightarrow K_2$. By the transitivity of homomorphisms, this implies that $C \rightarrow K_2$ which is a contradiction since no odd cycle can be 2-colored. Therefore the image of an odd cycle C under a homomorphism must contain an odd cycle, and this odd cycle must be of equal or lesser length since its image cannot contain more

vertices than C . Due to this property of odd cycles, we can define odd girth as follows:

$$\text{og}(X) = \min\{n \in \mathbb{N}, n \text{ odd} : C_n \rightarrow X\}.$$




Figure 2.2: Homomorphism from C_5 to K_3 .

There are many other graph parameters that can be defined via homomorphisms like the above. However, unlike the above it is sometimes nontrivial to show that the homomorphism definition is equivalent to the usual definition. An example of a parameter for which this is the case is the *fractional chromatic number*, denoted χ_f .

The fractional chromatic number of a graph X is defined by the linear program (LP) below, where $\mathcal{I}(X)$ denotes the set of all independent sets in X , and $\mathcal{I}(X, x)$ is the set of all independent sets in X containing the vertex x .

$$\begin{aligned} \min \quad & \sum_{I \in \mathcal{I}(X)} w_I \\ \text{s.t.} \quad & \sum_{I \in \mathcal{I}(X, x)} w_I \geq 1 \text{ for all } x \in V(X) \end{aligned}$$

If we further restrict the variables w_I in the above to be integer valued, then we obtain an integer program whose value is equal to the chromatic number of X . Solutions to the above LP are referred to as “fractional colorings”, and solutions to its dual given below are known as “fractional cliques”. By LP duality, the values of these two LPs are equal.

$$\begin{aligned} \max \quad & \sum_{x \in V(X)} w_x \\ \text{s.t.} \quad & \sum_{x \in I} w_x \leq 1 \text{ for all } I \in \mathcal{I}(X) \end{aligned}$$

To define the fractional chromatic number using homomorphisms, we must first introduce what are known as the *Kneser graphs*. The Kneser graph, $K_{n:r}$, has the r -subsets of $[n]$ as its vertices such that two are adjacent if they are disjoint. These graphs are related to the Erdős-Ko-Rado theorem which states that, for $n > 2r$, a family of pairwise intersecting r -subsets of $[n]$ has size at most $\binom{n-1}{r-1}$, and the only families meeting this bound are those which consist of all r -subsets containing some element fixed element i for $i \in [n]$. The Kneser graphs are of interest to us due to the following formulation of fractional chromatic number:

$$\chi_f(X) = \min \left\{ \frac{n}{r} \in \mathbb{Q} : X \rightarrow K_{n:r} \right\}.$$

The proof that this minimum is obtained and that it is equal to the value of the LP above is nontrivial and we omit it, but it can be found in [19].

The utility of defining these parameters in this manner lies in the following observation: if Y is n -colorable, then $Y \rightarrow K_n$, and if X is a graph such that $X \rightarrow Y$, then $X \rightarrow K_n$ and thus X is n -colorable. Therefore $X \rightarrow Y \Rightarrow \chi(X) \leq \chi(Y)$. Similarly, $X \rightarrow Y \Rightarrow \omega(X) \leq \omega(Y)$ and $\chi_f(X) \leq \chi_f(Y)$. We say that a graph parameter f for which $X \rightarrow Y$ implies $f(X) \leq f(Y)$ is *homomorphism monotone*. Note that under this definition, odd girth is not homomorphism monotone since $X \rightarrow Y$ implies $\text{og}(X) \geq \text{og}(Y)$ as opposed to the reverse inequality. The above implications are often of more use in their contrapositive forms, e.g. if $\chi(X) > \chi(Y)$ then $X \not\rightarrow Y$. Statements of this form (sufficient conditions for the nonexistence of a homomorphism) are often referred to as “no-homomorphism lemmas”. We will investigate this idea further in Section 2.7.

We take a moment here to note that the above discussion of homomorphisms from odd cycles implies that $C_{2n+1} \rightarrow C_{2m+1} \Rightarrow m \leq n$. To prove the converse, consider the cycles C_{2n+1} with vertices $\{u_1, u_2, \dots, u_{2n+1}\}$ and C_{2n-1} with vertices $\{v_1, v_2, \dots, v_{2n-1}\}$ where the indices indicate the order the vertices appear on the cycle. The map $\varphi : V(C_{2n+1}) \rightarrow V(C_{2n-1})$ given by

$$\varphi(u_i) = \begin{cases} v_i & \text{if } i = 1, 2, \dots, 2n-1 \\ v_1 & \text{if } i = 2n \\ v_2 & \text{if } i = 2n+1 \end{cases}$$

can easily be seen to be a homomorphism. Therefore $C_{2n+1} \rightarrow C_{2n-1}$ and by transitivity of homomorphisms, $C_{2n+1} \rightarrow C_{2m+1}$ whenever $m \leq n$.

In Chapter 6, we will see some more graph parameters which have formulations in terms of homomorphisms. As with the above, these parameters will easily be seen to be homomorphism monotone. But, unlike the parameters we have defined here, they will also be quantum homomorphism monotone, a concept we will define in Chapter 6.

2.4 Homomorphic Equivalence and the Homomorphism Order

The material in this section requires the following definition.

Definition. A *partially ordered set*, or simply *poset*, is a set \mathcal{P} equipped with a binary relation “ \leq ” satisfying the following properties:

- Reflexivity: $x \leq x$ for all $x \in \mathcal{P}$;
- Transitivity: $x \leq y$ and $y \leq z$ implies $x \leq z$ for all $x, y, z \in \mathcal{P}$;
- Antisymmetry: $x \leq y$ and $y \leq x$ implies $x = y$ for all $x, y \in \mathcal{P}$.

We refer to the relation “ \leq ” as a *partial order* on \mathcal{P} . For a partial order “ \leq ”, we will sometimes use the notation $x < y$ as shorthand for $x \leq y$ and $x \neq y$.

It is useful to consider the properties that “ \rightarrow ” has as a relation on graphs, and furthermore to consider whether it is a partial order. We have already seen above that “ \rightarrow ” is transitive. It is easy to see that it is reflexive as well, since the identity function is a homomorphism and thus $X \rightarrow X$ for all graphs X . We have also already seen examples that show that “ \rightarrow ” is not symmetric; two odd cycles of different lengths only have a homomorphism between them in one direction. One could also consider whether “ \rightarrow ” is antisymmetric, i.e. that $X \rightarrow Y$ and $Y \rightarrow X$ together imply that $X \cong Y$. This is, in fact, not the case. It is easy to see that if X is a bipartite graph, then $X \rightarrow K_2$ since X is 2-colorable. Furthermore, if X contains at least one edge, then $K_2 \rightarrow X$. Therefore $X \rightarrow K_2$ and $K_2 \rightarrow X$ for any nonempty bipartite graph X , and there are clearly many such graphs which are not isomorphic to K_2 .

Since “ \rightarrow ” is not antisymmetric, it is not a partial order. However, one can remedy this with the notion of homomorphic equivalence. We say that two graphs X and Y are *homomorphically equivalent* if $X \rightarrow Y$ and $Y \rightarrow X$. We will denote this by writing $X \equiv Y$. It is not hard to see that “ \equiv ” is in fact an equivalence relation on the class of all graphs. It is clearly reflexive since $X \rightarrow X$ for all graphs X as noted above. It is also clearly symmetric by design, and it is transitive by the transitivity of “ \rightarrow ”. Since “ \equiv ” is an equivalence relation, it partitions the class of graphs into equivalence classes. We refer to these classes as *homomorphic equivalence classes*, and we denote the class to which a graph X belongs by $\mathcal{H}(X)$.

Using homomorphic equivalence classes, we can construct a partial order induced by “ \rightarrow ”. We say that $\mathcal{H}(X) \leq \mathcal{H}(Y)$ if $X \rightarrow Y$. Though “ \rightarrow ” was not antisymmetric, this relation is since $\mathcal{H}(X) = \mathcal{H}(Y)$ if and only if $X \equiv Y$ if and only if $\mathcal{H}(X) \leq \mathcal{H}(Y)$ and $\mathcal{H}(Y) \leq \mathcal{H}(X)$. Note that “ \leq ” is well defined since $X \rightarrow Y$ if and only if $X' \rightarrow Y'$ for all graphs $X' \in \mathcal{H}(X)$, $Y' \in \mathcal{H}(Y)$. Also note that it is easy to see that “ \leq ” inherits the reflexivity and transitivity of “ \rightarrow ”. Since “ \leq ” is antisymmetric as well, it is a partial order on the set of homomorphic equivalence classes of graphs. We refer to this partially ordered set as the *homomorphism order of graphs* and denote it by \mathcal{G} . We will explore the poset \mathcal{G} in more depth in Section 2.10 and it will be central to the focus of Chapter 4.

Homomorphic equivalence also relates to certain graph parameters like those discussed in Section 2.3. For example, if $X \equiv Y$, then by definition we have that $\chi(X) \leq \chi(Y)$ and $\chi(Y) \leq \chi(X)$. Of course this implies that $\chi(X) = \chi(Y)$, and similarly we can see that $\omega(X) = \omega(Y)$, $\chi_f(X) = \chi_f(Y)$, etc. More generally, if f is any homomorphism-monotone graph parameter, then $X \equiv Y \Rightarrow f(X) = f(Y)$. In fact, any parameter g such that $X \rightarrow Y \Rightarrow g(X) \geq g(Y)$ (for example odd girth), will also have this property.

2.5 Cores

An important notion in the theory of graph homomorphisms, and in particular regarding homomorphic equivalence, is that of cores. A graph X is said to be a *core* if all of its endomorphisms are automorphisms. Equivalently, X is a core if it has no proper endomorphisms. Some examples of graphs which are cores are listed below.

- The complete graphs K_n
- The odd cycle C_{2n+1}
- The odd wheel W_{2k+1}
- The Kneser graph $K_{n:r}$ for $n > 2r$.

It is not too difficult to see that the first two examples above are cores, however the third requires much more work. We will see a proof that the Kneser graphs are cores which uses the Erdős-Ko-Rado theorem in Section 2.9. In general, it is NP-hard to determine whether a graph is a core [25].

Recall from Section 2.4 that there exist graphs X and Y such that $X \equiv Y$ but X and Y are not isomorphic. For cores however, this cannot happen.

Lemma 2.5.1. *Suppose graphs X and Y are cores. Then $X \equiv Y$ if and only if $X \cong Y$.*

Proof. One direction is straightforward. If $X \cong Y$, then there exist isomorphisms from X to Y and from Y to X . Since an isomorphism is a homomorphism, we have that $X \rightarrow Y$ and $Y \rightarrow X$ and thus $X \equiv Y$. This direction does not even require that X and Y are cores.

Now suppose that X and Y are cores such that $X \equiv Y$. This implies that there exist homomorphisms $\varphi_1 : X \rightarrow Y$ and $\varphi_2 : Y \rightarrow X$. We will show that these must in fact be isomorphisms. Suppose that φ_1 is not injective. Then $\varphi_2 \circ \varphi_1$ is an endomorphism of X which is not injective, and thus not a bijection. This contradicts the fact that X is a core. Similarly, φ_2 must be injective. Together these imply that $|V(X)| \leq |V(Y)|$ and $|V(X)| \geq |V(Y)|$ and thus $|V(X)| = |V(Y)|$. Since φ_1 is injective and $|V(X)| = |V(Y)|$, it is also surjective and thus a bijection. Similarly, φ_2 is a bijection. All that is left to show is that φ_1 preserves nonadjacency.

Since φ_i is a bijection for $i = 1, 2$, the endomorphism $\varphi_2 \circ \varphi_1$ is also a bijection from $V(X)$ to itself. Since $\varphi_2 \circ \varphi_1$ also preserves adjacency, it induces a bijection from $E(X)$ to itself. This implies that no nonedge of X is mapped to an edge of X . This in turn implies that φ_1 must preserve nonadjacency, and is thus an isomorphism. Therefore $X \cong Y$. \square

Note that we have actually proven a stronger statement:

Lemma 2.5.2. *Suppose graphs X and Y are cores such that $X \equiv Y$. Then any homomorphism from X to Y must be an isomorphism.* \square

We typically will only need Lemma 2.5.1, but Lemma 2.5.2 will be used below to help prove Lemma 2.5.10.

The above lemma implies that each homomorphic equivalence class contains at most one core (up to isomorphism). We will see later that each homomorphic equivalence class in fact contains exactly one core.

We say that a graph Y is a *core of X* if Y is a core, Y is a subgraph of X , and $X \rightarrow Y$. The first thing to note about this definition is that if X is a core, then X is a core of X . Another important property of this definition is given in the following lemma:

Lemma 2.5.3. *If Y is a core of a graph X , then $Y \equiv X$.*

Proof. By definition, $X \rightarrow Y$, so we only need to show that $Y \rightarrow X$. However, any subgraph of X has a homomorphism to X , in particular the inclusion map is such a homomorphism. Therefore $Y \equiv X$. \square

Combining the above lemma with Lemma 2.5.1 we obtain the following corollary:

Corollary 2.5.4. *If Y and Y' are cores of a graph X , then $Y \cong Y'$.*

Proof. By Lemma 2.5.3, $Y \equiv X$ and $Y' \equiv X$ and thus $Y \equiv Y'$. However, both Y and Y' are cores and thus by Lemma 2.5.1 we have $Y \cong Y'$. \square

Note that this does not mean that a graph has only one core; it may contain many distinct cores, but they all must be isomorphic. Also note that we have not shown that every graph has a core. To show this, we will prove the following stronger statement:

Lemma 2.5.5. *Let $M = \min\{|V(Z)| : Z \text{ is a subgraph of } X \text{ such that } X \rightarrow Z\}$. Then a subgraph Y of X is a core of X if and only if $X \rightarrow Y$ and $|V(Y)| = M$.*

Proof. Suppose that Y is a subgraph of X such that $X \rightarrow Y$ and $|V(Y)| = M$. To show that Y is a core of X we only need to show that Y is a core. Suppose not, and let φ be a homomorphism from X to Y . Since Y is not a core, there exists an endomorphism ψ of Y such that ψ is not a bijection. This implies that the image of ψ has strictly fewer vertices than Y . However, that would imply that $\psi \circ \varphi$ is a homomorphism from X to a subgraph of X with strictly fewer vertices than Y , which is a contradiction.

Now suppose that Y is a core of X such that $|V(Y)| > M$. Then there exists a subgraph Z of X such that $X \rightarrow Z$ and $|V(Z)| < |V(Y)|$. Let φ_Y and φ_Z be homomorphisms from X to Y and Z respectively. Then the restriction of $\varphi_Y \circ \varphi_Z$ to Y is an endomorphism of Y whose image necessarily has fewer vertices than Y . This contradicts the assumption that Y was a core of X . \square

The above lemma can be phrased as follows: a core of a graph X is a vertex minimal endomorphic image of X . As a corollary to this lemma and Corollary 2.5.4 we obtain the following:

Corollary 2.5.6. *Every graph has a unique core up to isomorphism.*

Proof. Since $X \rightarrow X$ and $V(X)$ is finite for all graphs X , Lemma 2.5.5 implies that every graph has at least one core. On the other hand, Corollary 2.5.4 states that any two cores of a graph X are isomorphic. \square

Due to the above corollary, we can refer to *the core of X* , which we will denote by X^\bullet . Note that when we say that X^\bullet is the core of X , we do not mean that X^\bullet is a subgraph of X , but rather a representative of the isomorphism class of a core of X . In contrast, when we say that Y is a core of X , we always mean that Y is a subgraph of X which is isomorphic to X^\bullet . As mentioned above, if X is a core, then X is a core of X , and now we can even say that X is *the* core of X .

As we have seen through some of the lemmas above, cores relate nicely to the notion of homomorphic equivalence. This is perhaps best emphasized by the following result:

Lemma 2.5.7. *If X and Y are graphs, then $X \equiv Y$ if and only if $X^\bullet \cong Y^\bullet$.*

Proof. Suppose that $X \equiv Y$. By Lemma 2.5.3 we have that $X^\bullet \equiv X$ and $Y^\bullet \equiv Y$ and thus $X^\bullet \equiv Y^\bullet$ since “ \equiv ” is transitive. Therefore, since X^\bullet and Y^\bullet are cores, we have that $X^\bullet \cong Y^\bullet$ by Lemma 2.5.1.

Conversely, suppose that $X^\bullet \cong Y^\bullet$. Then, using Lemma 2.5.3 again, we have that $X \equiv X^\bullet \equiv Y^\bullet \equiv Y$ and thus $X \equiv Y$. \square

A special case of the above lemma gives us a useful corollary.

Corollary 2.5.8. *Suppose that X and Y are graphs such that Y is a core. Then $X \equiv Y$ if and only if $X^\bullet \cong Y$.*

Proof. Suppose that X and Y are graphs and Y is a core. Then by Lemma 2.5.7, $X \equiv Y$ if and only if $X^\bullet \cong Y^\bullet$, but of course $Y^\bullet \cong Y$. \square

Earlier we gave some examples of graphs which are cores. Here we give some examples of graphs which are not cores, along with the cores of those graphs.

- The empty graph (on more than one vertex) whose core is K_1 .
- The even cycles C_{2k} , $k \in \mathbb{N}$, whose core is K_2 .
- The even wheels W_{2k} , $k \in \mathbb{N}$, whose core is K_3 .
- Any nonempty bipartite graph other than K_2 , whose core is K_2 .

Note that these examples are in fact all special cases of graphs whose core is a complete graph. This case arises quite a lot in the study of cores and we can characterize the graphs having this property with the following lemma.

Lemma 2.5.9. *If X is a graph, the core of X is K_n if and only if $\omega(X) = n = \chi(X)$.*

Proof. Suppose $X^\bullet = K_n$. Then $X \equiv K_n$ and thus

$$\omega(X) = \omega(K_n) = n = \chi(K_n) = \chi(X).$$

Conversely, if $\omega(X) = n = \chi(X)$, then $K_n \rightarrow X \rightarrow K_n$ and therefore $X \equiv K_n$. Since K_n is a core, applying Corollary 2.5.8 gives that K_n is the core of X . \square

Above we mentioned that we would see that every homomorphic equivalence class contains exactly one core (up to isomorphism). In fact, we have seen more than that, and the following statement is an apt summary of the results of this section: For any graph X , the unique (up to isomorphism) core contained in $\mathcal{H}(X)$ is X^\bullet . Furthermore, X^\bullet is the unique vertex minimal graph contained in $\mathcal{H}(X)$. Due to this, the partial order “ \leq ” defined on homomorphic equivalence classes in the previous section could equivalently be defined on cores.

Before we move on, we present one more important property of cores of graphs which will be useful in later chapters:

Lemma 2.5.10. *If Y is a core of a graph X , then Y is a retract of X .*

Proof. It suffices to exhibit a retraction of X whose image is Y . Let φ be a homomorphism from X to Y , and let σ be the restriction of φ to Y . Then σ is a homomorphism from Y to itself. Since Y is a core, by Lemma 2.5.2, the map σ must in fact be an isomorphism from Y to itself and thus $\sigma \in \text{Aut}(Y)$. Therefore, there exists $\sigma^{-1} \in \text{Aut}(Y)$ such that $\sigma^{-1} \circ \sigma$ is the identity map on Y . Now consider the map $\sigma^{-1} \circ \varphi$. Clearly this is a homomorphism since both φ and σ^{-1} are homomorphisms. Furthermore, it is a homomorphism to Y since φ is a homomorphism to Y and $\sigma^{-1} \in \text{Aut}(Y)$. Since $\sigma = \varphi|_Y$, we have that

$$(\sigma^{-1} \circ \varphi)|_Y = \sigma^{-1} \circ \varphi|_Y = \sigma^{-1} \circ \sigma = \text{id}.$$

Therefore, $\sigma^{-1} \circ \varphi$ is a retraction from X to Y and therefore Y is a retract of X . \square

This lemma makes life a bit easier when working with cores, since working with an arbitrary endomorphism is usually more cumbersome than working with retractions. Recall that a retract of a graph X is necessarily an induced subgraph of X , and therefore Lemma 2.5.10 gives the following corollary.

Corollary 2.5.11. *If Y is a core of X , then Y is an induced subgraph of X .* \square

Though this corollary is simple, it will be important in Chapter 4. We will return to the notion of cores and homomorphic equivalence in Sections 2.8, 2.9, and 2.10. Cores in particular will be a central feature of Chapters 3 and 4.

2.6 Graph Products and Other Constructions

In this section we will define several graph products and other graph constructions that will show up throughout this thesis. We will also discuss some of the basic properties of these constructions which will be important in later sections.

Given graphs X and Y , we define the following five graph products which all have vertex set $V(X) \times V(Y)$:

- *Cartesian product* ($X \square Y$): $(x, y) \sim (x', y')$ if either $x = x'$ and $y \sim y'$, or $x \sim x'$ and $y = y'$.
- *Categorical product* ($X \times Y$): $(x, y) \sim (x', y')$ if $x \sim x'$ and $y \sim y'$.
- *Strong product* ($X \boxtimes Y$): $(x, y) \sim (x', y')$ if $x \sim x'$ and $y \sim y'$, or $x = x'$ and $y \sim y'$, or $x \sim x'$ and $y = y'$.
- *Disjunctive product* ($X * Y$): $(x, y) \sim (x', y')$ if $x \sim x'$ or $y \sim y'$.
- *Lexicographic product* ($X[Y]$): $(x, y) \sim (x', y')$ if $x \sim x'$, or $x = x'$ and $y \sim y'$.

Note that if we want to take the, say, Cartesian product of r graphs X_1, X_2, \dots, X_r , then we can denote this by

$$\square_{i=1}^r X_i,$$

and similarly for any of the other first four products given above. In the special case where $X_i = X$ for all i , we write $X^{\square r}$, or the analogous notation for the other products.

We first consider the Cartesian product. An important property of this product is that for any fixed $x \in V(X)$, the set $\{(x, y) : y \in V(Y)\}$ induces a subgraph of $X \square Y$ which is isomorphic to Y . Similarly, for fixed $y \in V(Y)$, the set $\{(x, y) : x \in V(X)\}$ induces a subgraph which is isomorphic to X . In fact, it is easy to see that the same is true for the strong, disjunctive, and lexicographic products. This, in particular, implies that $X \rightarrow X \square Y$ and $Y \rightarrow X \square Y$. Given this fact, a natural question to ask is for which graphs X does $X \square X \rightarrow X$? A result of Larose, Laviolette, and Tardif [33] states that a graph X satisfies $X \square X \rightarrow X$ if and only if X is homomorphically equivalent to a normal Cayley graph, which will be defined in Chapter 3.

The categorical product appears in various places in graph theory, and is of particular importance to the homomorphism order of graphs as we will see in Section 2.10 below. One of the most important properties of this product is that $X \times Y \rightarrow X$ and $X \times Y \rightarrow Y$ for all graphs X and Y . To see this, consider the map $(x, y) \mapsto x$ for all $(x, y) \in V(X \times Y)$. If $(x, y) \sim (x', y')$, then $x \sim x'$ and therefore this map is a homomorphism from $X \times Y$ to X . The map $(x, y) \mapsto y$ analogously is a homomorphism from $X \times Y$ to Y . As a consequence of this, we see that

$$\chi(X \times Y) \leq \min\{\chi(X), \chi(Y)\}.$$

Perhaps the most well known unsolved problem in the theory of graph homomorphisms is the resolution of a conjecture of Hedetniemi [24] stating that equality holds above. As the above inequality also holds for any homomorphism monotone graph parameter, many different versions of it have been conjectured since the original. Perhaps most significantly, Zhu [53] recently proved the fractional version of Hedetniemi's conjecture, i.e. that

$$\chi_f(X \times Y) = \min\{\chi_f(X), \chi_f(Y)\}.$$

The strong graph product, also sometimes referred to as the *normal product*, arises in Shannon theory when considering asymptotic capacities of noisy channels. Note that vertices (x, y) and (x', y') are adjacent in $X \boxtimes Y$ if and only if they are adjacent in either $X \square Y$ or $X \times Y$.

The disjunctive product, also sometimes referred to as the *co-normal product*, can be used to define fractional versions of various graphs parameters. For example, the fractional chromatic number can be defined as follows [46]:

$$\chi_f(X) = \lim_{n \rightarrow \infty} \sqrt[n]{\chi(X^{*n})} = \inf_n \sqrt[n]{\chi(X^{*n})}.$$

The disjunctive product is actually related to the strong product through the following identity: $\overline{X * Y} = \overline{X} \boxtimes \overline{Y}$. To see this, suppose that (x, y) and (x', y') are distinct and adjacent in $\overline{X * Y}$. This implies that $x \not\sim x'$ in X and $y \not\sim y'$ in Y and either $x \neq x'$ or $y \neq y'$. Without loss of generality suppose that $x \neq x'$. Then $x \sim x'$ in \overline{X} and either $y = y'$ or $y \sim y'$ in \overline{Y} . In both cases, $(x, y) \sim (x', y')$ in $\overline{X} \boxtimes \overline{Y}$. The reverse implication is similar. The disjunctive product also has a relationship to the edge union of two graphs, which we will define below.

We will not be using the general lexicographic product, but rather the special case in which Y is the empty graph $\overline{K_n}$ for some $n \in \mathbb{N}$. In this case, we refer to $X \overline{[K_n]}$ as a *multiple* of X , as this graph can be obtained from X by creating $n - 1$ clones of each of the vertices in X . To be more explicit,

$$V(X \overline{[K_n]}) = \{(x, i) : x \in V(X), i \in [n]\},$$

and $(x, i) \sim (x', j)$ if and only if $x \sim x'$ in X . It is not hard to see that the map $(x, i) \mapsto x$ for all $x \in V(X)$ and $i \in [n]$ is a homomorphism from $X \overline{[K_n]}$ to X . Furthermore, if $f : V(X) \rightarrow [n]$ is any function, then the vertices $\{(x, f(x)) : x \in V(X)\}$ induce a copy of X in $X \overline{[K_n]}$. Therefore $X \rightarrow X \overline{[K_n]}$ and thus $X \equiv X \overline{[K_n]}$. This property will be of importance in Section 2.7 when we see that any vertex transitive graph is homomorphically equivalent to some Cayley graph.

Note that the notation for the first three products above is designed to resemble (a typical drawing of) the product of K_2 with itself. For example, $K_2 \square K_2 \cong C_4$ which can be drawn with its vertices and edges forming the corners and sides of a square, respectively. This allows one to easily remember and distinguish the notation for these three products.

The products above all interact with homomorphisms in different ways, but one property they have in common is the following:

Lemma 2.6.1. *Let \cdot be any of the five products defined above. Suppose that W, X, Y, Z are graphs such that $X \rightarrow W$ and $Y \rightarrow Z$. Then*

$$X \cdot Y \rightarrow W \cdot Z.$$

Proof. We give the proof for the lexicographic product; the proofs for the other products are similar. Suppose W, X, Y, Z are graphs such that $\varphi_1 : X \rightarrow W$ and $\varphi_2 : Y \rightarrow Z$. Define $\varphi : V(X[Y]) \rightarrow V(W[Z])$ as follows:

$$\varphi(x, y) = (\varphi_1(x), \varphi_2(y)).$$

We claim that this is a homomorphism. Indeed, if $(x, y) \sim (x', y')$ in $X[Y]$, then either $x \sim x'$, or $x = x'$ and $y \sim y'$. In the former case, $\varphi_1(x) \sim \varphi_1(x')$ in W and therefore $\varphi(x, y) \sim \varphi(x', y')$ in $W[Z]$. In the latter case, $\varphi_1(x) = \varphi_1(x')$ and $\varphi_2(y) \sim \varphi_2(y')$ in Z , and thus $\varphi(x, y) \sim \varphi(x', y')$ in $W[Z]$. \square

Though this lemma is simple, it can be useful for various purposes. For example, one can easily show that $\chi(X * Y) \leq \chi(X)\chi(Y)$ by using the above lemma to note that

$$X * Y \rightarrow K_{\chi(X)} * K_{\chi(Y)} \cong K_{\chi(X)\chi(Y)}.$$

There are two more graph constructions that will be of use to us throughout this thesis. The first is the *disjoint union* (or sometimes simply *union*) of graphs. For graphs X_1 and X_2 , the disjoint union of X_1 and X_2 , denoted $X_1 \cup X_2$, is the graph with vertex set $(V(X_1) \times \{1\}) \cup (V(X_2) \times \{2\})$ such that vertices (x, i) and (x', j) are adjacent if $i = j$ and $x \sim x'$ in X_i . Equivalently, we can assume that X_1 and X_2 have disjoint vertex sets (and thus edge sets) and take $X_1 \cup X_2$ to be the graph with vertex set $V(X_1) \cup V(X_2)$ and edge set $E(X_1) \cup E(X_2)$. It is quite obvious that both X and Y have homomorphisms to $X \cup Y$, but we will see in Section 2.10 that this is also the minimal (in terms of the homomorphism order) graph with this property. In this way, the disjoint union is related to the categorical product, which, as we will also see in Section 2.10, is the maximal graph that has homomorphisms to both X and Y .

The last construction we will present here is only defined for graphs with the same vertex set. For graphs X and Y with the same vertex set V , the *edge union* of X and Y , denoted $X + Y$, is the graph with vertex set V , such that vertices $z, z' \in V$ are adjacent if $z \sim z'$ in X or $z \sim z'$ in Y . As an example, we noted above that for graphs X and Y , two elements of $V(X) \times V(Y)$ are adjacent in $X \boxtimes Y$ if and only if they are adjacent in either $X \square Y$ or $X \times Y$. Therefore, $X \boxtimes Y = (X \square Y) \cup (X \times Y)$.

In the special case where $Y = \overline{X}$, we see that $X + Y \cong K_n$ where $n = |V(X)|$. As mentioned above, the edge union and disjunctive product are related: Suppose that X and Y have the same vertex set V . The vertices $\{(z, z) : z \in V\}$ induce a subgraph in $X * Y$ which is isomorphic to $X + Y$.

Note that if we desire we can define the edge union of two arbitrary graphs, by simply relabeling the vertices of one of the graphs and adding any possibly missing vertices. However, this obviously depends on how the relabelling is done and is therefore not unique given two graphs X and Y .

It is worth pointing out that both the disjoint union and edge union can be viewed as special cases of the same construction. Given two graphs X and Y , consider the graph with vertex set $V(X) \cup V(Y)$ and edge set $E(X) \cup E(Y)$. In this case, $V(X)$ and $V(Y)$ may intersect nontrivially, and so may $E(X)$ and $E(Y)$. Then the disjoint union is the special case where $V(X)$ and $V(Y)$ are disjoint, and the edge union is the special case where $V(X) = V(Y)$. We will not use this construction in its full generality, but we mention it for completeness.

2.7 Transitive Graphs

In this section and the next two we investigate graphs with a high degree of symmetry. We will mainly focus on vertex transitive graphs, though we will also spend some time dealing with arc transitive graphs and non-edge transitive graphs in Section 2.9. It is not obvious how the symmetry of a graph should relate to homomorphisms to or from that graph, but we will see several examples below of how these two aspects interact.

We begin by defining vertex transitive graphs. Recall that for a graph X , the group of automorphisms of X is denoted $\text{Aut}(X)$. This group acts on the vertices of X , and if it acts transitively, we say that X is *vertex transitive*. To be more explicit, a graph X is vertex transitive if for all $x, x' \in V(X)$, there exists $\sigma \in \text{Aut}(X)$ such that $\sigma(x) = x'$. Some examples of vertex transitive graphs are listed below.

- The complete graph K_n .
- The empty graph $\overline{K_n}$.
- The Kneser Graph $K_{n,r}$.
- The cycle C_n .
- Cayley graphs (defined in Section 2.8).

Note that the first two graphs listed above are complements. More generally, a graph X is vertex transitive if and only if \overline{X} is vertex transitive.

An *arc* of a graph X is an *ordered pair* of adjacent vertices in X . This is distinguished from an edge of X which is an *unordered* pair of adjacent vertices in X . Since automorphisms map edges to edges, they also map arcs to arcs. Therefore, $\text{Aut}(X)$ can also be viewed as acting on arcs, or edges, of X . As with vertices, we say that X is *arc transitive* if $\text{Aut}(X)$ acts transitively on the set of arcs of X . Again, to be explicit, a graph X is arc transitive if for any two ordered pairs (x, y) and (x', y') of adjacent vertices of X , there exists $\sigma \in \text{Aut}(X)$ such that $\sigma(x) = x'$ and $\sigma(y) = y'$.

Similarly, X is *edge transitive* if for any two unordered pairs $\{x, y\}$ and $\{x', y'\}$ of adjacent vertices, there exists $\sigma \in \text{Aut}(X)$ such that either $\sigma(x) = x'$ and $\sigma(y) = y'$, or $\sigma(x) = y'$ and $\sigma(y) = x'$. A graph can also be *non-edge transitive*, which is defined analogously to edge transitive except that we replace unordered pairs of adjacent vertices with unordered pairs of distinct nonadjacent vertices. Obviously, a graph X is non-edge transitive if and only if \overline{X} is edge transitive. There are many more types of transitivity that arise in graph theory, but those mentioned above are the most relevant for the work here.

Though the definition above does not imply this, when we say a graph is arc transitive, we will implicitly mean that it has no isolated vertices (or is an empty graph). Since removing isolated vertices does not change the homomorphic equivalence class, this assumption does not make much difference, but it ensures that all arc transitive graphs are also vertex transitive. Note that the converse is not true: there exist vertex transitive graphs which are not arc transitive. For example, the vertices and edges of the truncated tetrahedron form a vertex but not edge transitive graph.

The first result we present deals with homomorphisms to vertex transitive graphs. Recall that in Section 2.3, we mentioned a type of result known as a “no-homomorphism lemma”. These are essentially sufficient conditions for showing that no homomorphism

exists from a graph X to another graph Y . A simple example of such a condition is $\chi(X) > \chi(Y)$, since if $X \rightarrow Y$, then any n -coloring of Y can be translated into an n -coloring of X via the composition of homomorphisms. The more involved example we will see below considers a parameter known as the *independence ratio* of a graph. The independence ratio of a graph X , denoted $i(X)$, is simply equal to $\alpha(X)/|V(X)|$. This parameter is somewhat related to the chromatic number by the well known lower bound $1/i(X) = |V(X)|/\alpha(X) \leq \chi(X)$. The following lemma is known as “*the no-homomorphism lemma*”, and was originally proved (in a stronger form) in [2], but the proof we present here comes from [23]. Note that the lemma is given in the contrapositive form, i.e. as a necessary condition for the existence of a homomorphism.

Lemma 2.7.1. *Suppose that X and Y are graphs and Y is vertex transitive. If $X \rightarrow Y$, then $i(X) \geq i(Y)$.*

Proof. Let $\mathcal{S}(Y)$ denote the family of independent sets of size $\alpha(Y)$ in Y . Since Y is vertex transitive, every vertex of Y is contained in the same number, say m , of elements of $\mathcal{S}(Y)$. Counting ordered pairs of the form (S, y) where $y \in S \in \mathcal{S}(Y)$ in two different ways yields

$$\alpha(Y)|\mathcal{S}(Y)| = m|V(Y)|. \quad (2.1)$$

Now suppose that $\varphi : X \rightarrow Y$. Recall that for any homomorphism, the inverse image of an independent set must be an independent set. Therefore, for all $S \in \mathcal{S}(Y)$ we have that $|\varphi^{-1}(S)| \leq \alpha(X)$. Summing up this inequality for all elements of $\mathcal{S}(Y)$, we obtain

$$\sum_{S \in \mathcal{S}(Y)} |\varphi^{-1}(S)| \leq \alpha(X)|\mathcal{S}(Y)|. \quad (2.2)$$

However, each $x \in V(X)$ contributes exactly m to the sum above, since $\varphi(x)$ belongs to exactly m members of $\mathcal{S}(Y)$. Therefore,

$$\sum_{S \in \mathcal{S}(Y)} |\varphi^{-1}(S)| = m|V(X)|. \quad (2.3)$$

Combining (2.1), (2.2), and (2.3), we obtain

$$i(X) = \frac{\alpha(X)}{|V(X)|} \geq \frac{m}{|\mathcal{S}(Y)|} = \frac{\alpha(Y)}{|V(Y)|} = i(Y). \quad \square \quad (2.4)$$

One application of the no-homomorphism lemma is a quick proof that $C_{2\ell+1} \not\rightarrow C_{2k+1}$ if $\ell < k$. Since $\alpha(C_{2\ell+1}) = \ell$ and $\alpha(C_{2k+1}) = k$, we have

$$i(C_{2\ell+1}) = \frac{\ell}{2\ell+1} < \frac{k}{2k+1} = i(C_{2k+1}),$$

and the result follows.

If we consider the special case of when the graph X above is a complete graph, we can obtain what is known as the “clique-coclique bound” for vertex transitive graphs. Originally proven in [11], this result has many different proofs including one which uses Cayley graphs that we will see later.

Lemma 2.7.2. *If Y is a vertex transitive graph, then*

$$\alpha(Y)\omega(Y) \leq |V(Y)|.$$

Proof. Let $n = \omega(Y)$. Then $K_n \rightarrow Y$, and by the no-homomorphism lemma we have that $i(K_n) \geq i(Y)$. Since $i(K_n) = 1/n = 1/\omega(Y)$, we obtain

$$\frac{1}{\omega(Y)} \geq \frac{\alpha(Y)}{|V(Y)|} \Rightarrow \alpha(Y)\omega(Y) \leq |V(Y)|. \quad \square$$

If we consider the proof of the no-homomorphism lemma in the case where $i(X) = i(Y)$, we see that something quite strong can be said about the homomorphisms from X to Y :

Lemma 2.7.3. *Suppose X and Y are graphs such that Y is vertex transitive and $i(X) = i(Y)$. If $\varphi : X \rightarrow Y$, then*

$$|\varphi^{-1}(S)| = \alpha(X)$$

for all $S \in \mathcal{S}(Y)$.

Proof. Since $i(X) = i(Y)$, the inequality in (2.4) must be satisfied with equality, and therefore (2.2) must be satisfied with equality. Since (2.2) was the sum of the inequalities $|\varphi^{-1}(S)| \leq \alpha(X)$ for all $S \in \mathcal{S}(Y)$, we must have that all of those inequalities are satisfied with equality, thus yielding the result. \square

Applying this lemma to the clique-coclique bound scenario we obtain the following:

Lemma 2.7.4. *Let Y be a vertex transitive graph such that $\alpha(Y)\omega(Y) = |V(Y)|$. If S is a maximum independent set and T is a maximum clique of Y , then*

$$|S \cap T| = 1.$$

Proof. Let $n = \omega(Y)$. Since T is a maximum clique of Y , there exists a homomorphism $\varphi : K_n \rightarrow Y$ whose image is the subgraph induced by T . Since S is a maximum independent set of Y , by Lemma 2.7.3 we have that $|\varphi^{-1}(S)| = \alpha(K_n) = 1$. Therefore, exactly one

vertex of K_n is mapped to a vertex of S . However, since the image of φ is T , this implies that exactly one of T is contained in S . \square

In Section 2.9 we will see that Lemma 2.7.3 can be used to prove that the Kneser graph $K_{n:r}$ for $n > 2r$ is a core.

2.8 Cayley Graphs

In this section we will introduce a natural subclass of vertex transitive graphs known as Cayley graphs. Though there are vertex transitive graphs which are not Cayley graphs, we will see that every vertex transitive graph is homomorphically equivalent to some Cayley graph. This emphasizes the importance of considering Cayley graphs when investigating homomorphisms of vertex transitive graphs in general.

For a group G and subset $C \subseteq G \setminus \{1\}$ such that $g^{-1} \in C$ for all $g \in C$, we define the *Cayley graph*, $X(G, C)$, to be the graph with vertex set G such that $g \sim h$ if $gh^{-1} \in C$. We typically refer to C as the *connection set* of $X(G, C)$, and the property $g \in C \Rightarrow g^{-1} \in C$ is referred to as being “inverse closed”. Note that this is equivalent to saying that $C^{-1} = C$, where $C^{-1} = \{g \in G : g^{-1} \in C\}$.

For $a \in G$, define $f_a : G \rightarrow G$ to be the map given by $f_a(g) = ga$ for all $g \in G$. This map is clearly bijective, and $f_a(g)f_a(h)^{-1} = (ga)(ha)^{-1} = gaa^{-1}h^{-1} = gh^{-1}$. Therefore $g \sim h$ if and only if $f_a(g) \sim f_a(h)$ and thus f_a is an automorphism of $X(G, C)$ for all $a \in G$. Furthermore, if $g, h \in G$, then $f_{g^{-1}h}(g) = gg^{-1}h = h$, and thus $X(G, C)$ is vertex transitive. The maps f_a for $a \in G$ are known as *right translations*. It is important to note that the *left translations* are not necessarily automorphisms. Indeed, if $gh^{-1} \in C$, then $(ag)(ah)^{-1} = agh^{-1}a^{-1}$ need not be in C . In Section 3.3, we will investigate the subclass of Cayley graphs for which the left translations are in fact automorphisms.

In order to present the main result of this section, we first need to define the stabilizer of a vertex. For a graph X and subgroup G of $\text{Aut}(X)$, the *stabilizer* of $x \in V(X)$ with respect to G is denoted by G_x , and is defined as follows:

$$G_x = \{\sigma \in G : \sigma(x) = x\}.$$

In other words, it is the set of automorphisms of X which fix the vertex x . Note that the stabilizer of a vertex is not just a subset of $\text{Aut}(X)$, but is also a subgroup. Furthermore, if $\sigma \in G$ is such that $\sigma(x) = x'$, then the left coset σG_x consists of exactly those automorphisms of X in G which map x to x' .

The theorem below was originally proven by Sabidussi in [44].

Theorem 2.8.1 (Sabidussi). *Let X be a vertex transitive graph. Then some multiple of X is a Cayley graph.*

Proof. We will construct a Cayley graph which is isomorphic to a multiple of X . Let $x^* \in V(X)$, $G = \text{Aut}(X)$, and define

$$C = \{\sigma \in G : x^* \sim \sigma(x^*)\}.$$

Note that C is inverse closed since if $x^* \sim \sigma(x^*)$, then $\sigma^{-1}(x^*) \sim x^*$, since σ^{-1} is an automorphism. Also, since $x^* \not\sim x^*$, the identity automorphism is not contained in C . Therefore, the graph $Y = X(\text{Aut}(X), C)$ is a Cayley graph.

Define $G(x) = \{\sigma \in G : \sigma^{-1}(x^*) = x\}$ for all $x \in V(X)$. Then $\sigma_1 \in G(x_1)$ is adjacent to $\sigma_2 \in G(x_2)$ in Y if and only if $\sigma_1\sigma_2^{-1} \in C$ if and only if $x^* \sim (\sigma_1\sigma_2^{-1})(x^*)$ in X if and only if $\sigma_1^{-1}(x^*) \sim \sigma_2^{-1}(x^*)$ in X if and only if $x_1 \sim x_2$ in X . Since X is vertex transitive, the $G(x)$ are nonempty for all $x \in V(X)$ and they form a partition of $\text{Aut}(X)$. Furthermore, the $G(x)$ all have the same size since $G(x)^{-1}$ is a left coset of G_{x^*} for all $x \in V(X)$. Therefore, the vertex set of Y can be partitioned into the equal sized sets $G(x)$ for $x \in V(X)$, such that a vertex of $G(x_1)$ is adjacent in Y to a vertex of $G(x_2)$ if and only if $x_1 \sim x_2$ in X . Thus Y is isomorphic to a multiple of X , specifically $X \left[\overline{K_n} \right]$ where $n = |G_{x^*}|$. \square

Recall from Section 2.6 that all multiples of a graph X are homomorphically equivalent to X . Therefore we have the following corollary.

Corollary 2.8.2. *If X is a vertex transitive graph, then X is homomorphically equivalent to some Cayley graph.* \square

The advantage of Theorem 2.8.1 and Corollary 2.8.2 is that it allows one to generalize certain results on Cayley graphs to all vertex transitive graphs. Since the concrete algebraic structure of Cayley graphs often makes them easier to work with than vertex transitive graphs in general, this can be a significant help.

As an example of this type of application of these results, we give a second proof of the clique-coclique bound we proved above.

Lemma 2.7.2. *If Y is a vertex transitive graph, then*

$$\alpha(Y)\omega(Y) \leq |V(Y)|.$$

Proof. We first prove the result for Cayley graphs, then apply Theorem 2.8.1 to prove it for vertex transitive graphs in general.

Let Y be a Cayley graph. Thus $Y = X(G, C)$ for some group G and some inverse closed $C \subseteq G \setminus \{1\}$. Let $S \subseteq G$ be an independent set of Y of size $\alpha(Y)$, and let $T \subseteq G$ be a clique of size $\omega(Y)$ in Y . We claim that the sets $S^{-1}g = \{h^{-1}g : h \in S\}$ for $g \in T$ are pairwise disjoint. If not, then there exists distinct $g_1, g_2 \in T$ and $h_1, h_2 \in S$ such that $h_1^{-1}g_1 = h_2^{-1}g_2 \Rightarrow g_1g_2^{-1} = h_1h_2^{-1}$. However, $g_1 \sim g_2$ for distinct $g_1, g_2 \in T$ and thus $g_1g_2^{-1} \in C$, whereas $h_1h_2^{-1} \notin C$ since $h_1 \not\sim h_2$. Since the $S^{-1}g$ are pairwise disjoint, we must have that

$$|V(Y)| \geq |S||T| = \alpha(Y)\omega(Y).$$

Now suppose that Y is a vertex transitive graph. By Theorem 2.8.1, there exists an $m \in \mathbb{N}$ such that $X = Y \overline{[K_m]}$ is a Cayley graph. Therefore, by the above,

$$\alpha(X)\omega(X) \leq |V(X)|.$$

We now simply rewrite the above parameters for X in terms of the corresponding parameters for Y . First, and most obviously, we have that $|V(X)| = m|V(Y)|$. Second, since X and Y are homomorphically equivalent, we have that $\omega(X) = \omega(Y)$. Lastly, we can see either directly from the definition of the lexicographic product, or by applying the no-homomorphism lemma for $X \rightarrow Y$ and $Y \rightarrow X$, that $\alpha(X) = m\alpha(Y)$. Therefore the above equation becomes

$$m\alpha(Y)\omega(Y) \leq m|V(Y)|,$$

and canceling the m produces the desired result. \square

In Section 6.12 we will prove a “quantum” version of the clique-coclique bound through a proof similar to the above.

2.9 Cores of Transitive Graphs

In this section we will introduce some results concerning the cores of various types of transitive graphs. In particular, these results will allow us to reduce the possibilities for the cores of specific graphs.

The first theorem we present is a result due to Welzl [50], and states that the vertex transitivity of a graph is inherited by its core.

Theorem 2.9.1. *Let X be a graph. If X is vertex transitive then so is X^\bullet .*

Proof. Let Y be a core of X and let ρ be a retraction from X to Y . Let $x, y \in V(Y)$. Since X is vertex transitive, there exists $\sigma \in \text{Aut}(X)$ such that $\sigma(x) = y$. Consider the homomorphism $\varphi = (\rho \circ \sigma)|_Y$. Since ρ has Y as its image, the map φ is a homomorphism from Y to itself, and is therefore an automorphism since Y is a core. Furthermore, since ρ is identity on Y , we have that

$$\varphi(x) = \rho(\sigma(x)) = \rho(y) = y.$$

Therefore Y is vertex transitive. □

It is not hard to see that the above proof can be generalized to many other types of transitivity such as arc and edge transitivity. We will use this theorem and its arc transitive cousin in Chapter 3 to help determine the structure of the cores of a certain class of graphs.

A nice application of the above theorem is the following sufficient condition for a vertex transitive graph to be a core.

Corollary 2.9.2. *If X is a vertex transitive graph such that $|V(X)|$ and $\alpha(X)$ are relatively prime, then X is a core.*

Proof. By the no-homomorphism lemma, the independence ratios of X and X^\bullet are equal. However, as $|V(X)|$ and $\alpha(X)$ are relatively prime, no ratio of integers with a denominator less than $|V(X)|$ can be equal to $i(X)$. Since $|V(X^\bullet)| \leq |V(X)|$ trivially, we must have equality here and thus $X \cong X^\bullet$. □

Though determining the independence number of a graph is NP-hard (even given that the graph is vertex transitive [10]), if a nonempty vertex transitive graph has a prime number of vertices, then we can immediately deduce that it is a core.

The next theorem is the basis of our work in Chapter 3. It was originally proven in [23] by Hahn and Tardif, and we give (a modified version of) their proof here.

Theorem 2.9.3. *If X is a vertex transitive graph and φ is an endomorphism of X whose image is a core Y of X , then all of the fibres $\varphi^{-1}(y)$, $y \in V(Y)$, have the same size and thus $|V(Y)|$ divides $|V(X)|$.*

Proof. Let $y^* \in V(Y)$, $G = \text{Aut}(X)$, and set

$$S = \{(y, \sigma) : y \in V(Y), \sigma \in G, \varphi \circ \sigma(y) = y^*\}.$$

We proceed by counting the elements of S in two different ways. First, for any $\sigma \in \text{Aut}(X)$, the map $(\varphi \circ \sigma)|_Y$ is an automorphism of Y , so it must be bijective. Therefore, there exists

a unique $y \in V(Y)$ such that $\varphi \circ \sigma(y) = y^*$, i.e. $(y, \sigma) \in S$. Thus $|S| = |\text{Aut}(X)|$. Second, for any $y \in V(Y)$ and $x \in \varphi^{-1}(y^*)$, the set of automorphisms of X mapping y to x is a left coset of G_y , and has cardinality $|\text{Aut}(X)|/|V(X)|$. Thus,

$$|S| = |V(Y)| \cdot |\varphi^{-1}(y^*)| \cdot \frac{|\text{Aut}(X)|}{|V(X)|}.$$

Combining these two expressions for $|S|$, we obtain $|\varphi^{-1}(y^*)| = |V(X)|/|V(Y)|$, which clearly does not depend on y^* . \square

In light of this theorem, a natural question to ask is what can be said about vertex transitive graphs X , such that $|V(X)| = 2|V(X^\bullet)|$. This is in some sense the simplest nontrivial case, and so it seems possible that graphs satisfying this property may have some special structure. In Section 3.1, we address precisely this question and find that very much can be said about the structure of such graphs, and even more can be said if we further assume arc transitivity.

Note that the above theorem can be modified to show that the number of arcs in the core of an arc transitive graph must divide the number of arcs in the graph, but this does not seem to be a particularly useful result. Instead, the following theorem shows that arc transitivity actually gives us something stronger, as well as more interesting. In the proof below we will use $N_X(x)$ to denote the set of vertices adjacent to vertex x in graph X . This theorem was originally proved by Godsil in [18].

Theorem 2.9.4. *If X is an arc transitive graph, then the valency of X^\bullet divides the valency of X .*

Proof. First note that if X is arc transitive and $G = \text{Aut}(X)$, then G_x acts transitively on $N_X(x)$, for all $x \in V(X)$. Now suppose that ρ is a retraction from X to a core Y of X . Furthermore, let $y_1, y_2 \in V(Y)$ such that $y_1 \sim y_2$, and let

$$S = \{(y, \sigma) : y \in N_Y(y_1), \sigma \in G_{y_1}, \rho \circ \sigma(y) = y_2\}.$$

As in the proof of Theorem 2.9.3, we proceed by counting the elements of S in two ways. First note that since σ fixes y_1 and ρ is a retraction, the map $\rho \circ \sigma$ fixes y_1 . Therefore, $\rho \circ \sigma$ maps $N_Y(y_1)$ to itself. Furthermore, it must do this bijectively since otherwise $\rho \circ \sigma \circ \rho$ would be an endomorphism of X whose image contains fewer vertices than Y , a contradiction. This implies that for each $\sigma \in G_{y_1}$, there exists a unique $y \in N_Y(y_1)$ such that $\rho \circ \sigma(y) = y_2$, i.e. $(y, \sigma) \in S$. Therefore we have that $|S| = |G_{y_1}|$.

Alternatively, for any $y \in N_Y(y_1)$ and $x \in N_X(y_1) \cap \rho^{-1}(y_2)$, the set of automorphisms in G_{y_1} mapping y to x is a left coset of the subgroup of automorphisms in G_{y_1} which fix y and thus has size $|G_{y_1}|/|N_X(y_1)|$. Therefore,

$$|S| = |N_Y(y_1)| \cdot |\rho^{-1}(y_2) \cap N_X(y_1)| \cdot \frac{|G_{y_1}|}{|N_X(y_1)|}.$$

Combining our two expressions for $|S|$ we obtain

$$|\rho^{-1}(y_2) \cap N_X(y_1)| = \frac{|N_X(y_1)|}{|N_Y(y_1)|}.$$

Since $N_X(y_1)$ and $N_Y(y_1)$ are the valencies of X and Y respectively, and $|\rho^{-1}(y_2) \cap N_X(y_1)|$ must be an integer, we have proved the desired result. \square

As an immediate consequence of this result, we can see that any nonbipartite connected arc transitive graph with prime valency must be a core, since K_2 is the only core with valency 1.

We mentioned earlier that the case of graphs having complete cores comes up relatively frequently. One example of a result in which this is the case is the following which was originally proved in [7]. Their proof used a notion somewhat dual to that of a core known as the *hull* of a graph. The proof we present here is similar in spirit, but differs in presentation.

Lemma 2.9.5. *If X is non-edge transitive, then X is either a core or has a complete core.*

Proof. Suppose that X is non-edge transitive and not a core. Let ρ be a retraction from X to a core Y of X . Since X is not a core, there exists distinct $x, y \in V(X)$ such that $\rho(x) = \rho(y)$. Note that this implies that x and y are necessarily not adjacent, and thus $\{x, y\}$ is a non-edge. If Y is complete, we are done, otherwise there exists distinct $x', y' \in V(Y)$ such that $x' \not\sim y'$. Since X is non-edge transitive, without loss of generality, there exists $\sigma \in \text{Aut}(X)$ such that $\sigma(x') = x$ and $\sigma(y') = y$. However, this implies that $\rho \circ \sigma \circ \rho$ is an endomorphism of X whose image has strictly fewer vertices than Y , a contradiction. \square

To end this section we give an example of a proof that a specific class of graphs are all cores. The graphs in question are the Kneser graphs $K_{n,r}$ for $n > 2r$. Recall that $K_{n,r}$ is the graph whose vertices are the r -subsets of $[n]$ such that disjoint subsets are adjacent. The proof uses the following result known as the Erdős-Ko-Rado theorem.

Theorem 2.9.6 (Erdős, Ko, Rado). *Let $n, r \in \mathbb{N}$ satisfy $n > 2r$. Then*

$$\alpha(K_{n:r}) = \binom{n-1}{r-1},$$

and the only independent sets which meet this bound are the sets

$$F_i = \{S \subseteq [n] : |S| = r, i \in S\}$$

for $i \in [n]$. □

The following theorem was originally proven by Hahn and Tardif in [23], but the proof we present is essentially that of [19].

Theorem 2.9.7. *Let $n, r \in \mathbb{N}$ satisfy $n > 2r$. Then $K_{n:r}$ is a core.*

Proof. First note that for any $S \in V(K_{n:r})$, we have

$$\{S\} = \bigcap_{i \in S} F_i.$$

Now suppose that φ is an endomorphism of $K_{n:r}$. Since $K_{n:r}$ is vertex transitive, we can apply the no-homomorphism lemma to obtain $|\varphi^{-1}(F_i)| = \alpha(K_{n:r})$, and thus by the Erdős-Ko-Rado theorem $\varphi^{-1}(F_i) = F_{g(i)}$ for some function $g : [n] \rightarrow [n]$. This implies that for $S' \in V(K_{n:r})$,

$$\varphi^{-1}(S') = \varphi^{-1} \left(\bigcap_{i \in S'} F_i \right) = \bigcap_{i \in S'} \varphi^{-1}(F_i) = \bigcap_{i \in S'} F_{g(i)} = \{S \subseteq [n] : |S| = r, g(S') \subseteq S\},$$

where $g(S') = \{g(i) : i \in S'\}$. However, the last set above is clearly nonempty since $|g(S')| \leq r$ and thus every vertex of $K_{n:r}$ has a nonempty preimage. This implies that φ is surjective and thus an automorphism. Therefore, $K_{n:r}$ has no proper endomorphisms and must be a core. □

For similar results to the one above, the reader is encouraged to look at [7, 20] or Chapter 6 of [19].

2.10 Lattice Properties of \mathcal{G}

Let \mathcal{P} be a partially ordered set with partial order \leq , and let $x, y \in \mathcal{P}$. An element $w \in \mathcal{P}$ is the *meet* of x and y if the following conditions hold:

- $w \leq x$ and $w \leq y$;
- for any $u \in \mathcal{P}$ such that $u \leq x$ and $u \leq y$, we have that $u \leq w$.

Note that this definition guarantees that if there exists a meet of two elements, then it is unique. We denote the meet of elements x and y by $x \wedge y$. The meet is also sometimes referred to as the greatest lower bound for obvious reasons.

Analogously, we can consider the least upper bound of x and y . We say that z is the *join* of x and y if the following conditions are met:

- $x \leq z$ and $y \leq z$;
- for any $v \in \mathcal{P}$ such that $x \leq v$ and $y \leq v$, we have that $z \leq v$.

Again, this definition guarantees uniqueness of the join. We denote the join of elements x and y by $x \vee y$.

A partially ordered set for which every two elements have both a meet and a join is known as a *lattice*. The purpose of this section is to show that \mathcal{G} is in fact a lattice.

Note that to show an element w is the meet of elements x and y in a partially ordered set it is both necessary and sufficient to show that $u \leq w$ if and only if $u \leq x$ and $u \leq y$. Similarly, to show that z is the join of x and y , it is necessary and sufficient to show that $z \leq v$ if and only if $x \leq v$ and $y \leq v$.

The following two lemmas show that the meet and join operations on \mathcal{G} correspond to the categorical product and disjoint union of graphs respectively.

Lemma 2.10.1. *Let X and Y be graphs. The meet of the homomorphic equivalence classes $\mathcal{H}(X)$ and $\mathcal{H}(Y)$ is the homomorphic equivalence class $\mathcal{H}(X \times Y)$.*

Proof. We must show that $\mathcal{H}(Z) \leq \mathcal{H}(X)$ and $\mathcal{H}(Z) \leq \mathcal{H}(Y)$ if and only if $\mathcal{H}(Z) \leq \mathcal{H}(X \times Y)$ for any graph Z . However, by the results of Section 2.4 we can see that it suffices to show that $Z \rightarrow X$ and $Z \rightarrow Y$ if and only if $Z \rightarrow X \times Y$. Suppose that $Z \rightarrow X \times Y$. We saw in Section 2.6 that $X \times Y \rightarrow X$ and $X \times Y \rightarrow Y$; therefore $Z \rightarrow X$ and $Z \rightarrow Y$.

Conversely, suppose that $\varphi_X : Z \rightarrow X$ and $\varphi_Y : Z \rightarrow Y$ are homomorphisms. Define $\varphi : V(Z) \rightarrow V(X \times Y)$ to be the map given by

$$\varphi(z) = (\varphi_X(z), \varphi_Y(z))$$

for all $z \in V(Z)$. We claim that this is a homomorphism. Suppose that z and z' are adjacent vertices of Z . Since φ_X and φ_Y are homomorphisms, we have that $\varphi_X(z) \sim \varphi_X(z')$ and $\varphi_Y(z) \sim \varphi_Y(z')$. Therefore, $\varphi(z) \sim \varphi(z')$. \square

Lemma 2.10.2. *Let X and Y be graphs. The join of the homomorphic equivalence classes $\mathcal{H}(X)$ and $\mathcal{H}(Y)$ is the homomorphic equivalence class $\mathcal{H}(X \cup Y)$.*

Proof. It suffices to show that $X \rightarrow Z$ and $Y \rightarrow Z$ if and only if $X \cup Y \rightarrow Z$. Suppose that $\varphi_X : X \rightarrow Z$ and $\varphi_Y : Y \rightarrow Z$ are homomorphisms. Define $\varphi : V(X \cup Y) \rightarrow V(Z)$ as follows:

$$\varphi(w) = \begin{cases} \varphi_X(w) & \text{if } w \in V(X) \\ \varphi_Y(w) & \text{if } w \in V(Y). \end{cases}$$

If $w \sim w'$ in $X \cup Y$, then either $w, w' \in V(X)$ or $w, w' \in V(Y)$. In the former case, $\varphi(w) = \varphi_X(w) \sim \varphi_X(w') = \varphi(w')$ since φ_X is a homomorphism. The latter case is similar and therefore φ is a homomorphism from $X \cup Y$ to Z .

Conversely, suppose that $\varphi : X \cup Y \rightarrow Z$ is a homomorphism. Clearly, the restrictions $\varphi|_X$ and $\varphi|_Y$ are homomorphisms from X to Z and Y to Z respectively. Therefore $X \rightarrow Z$ and $Y \rightarrow Z$. \square

Since we have shown that any two homomorphic equivalence classes have both a meet and a join, we have shown our desired result:

Theorem 2.10.3. *The partially ordered set \mathcal{G} is a lattice.* \square

2.11 Density and Universality

In this section we discuss two important properties of the order \mathcal{G} : density and universality. We will not need the material of this section for any of our results, but we will use these two properties of \mathcal{G} as a basis of comparison to the homomorphism order of line graphs in Chapter 4.

A partially ordered set \mathcal{P} is said to be *dense* if for any two elements $x, z \in \mathcal{P}$ such that $x < z$, there exists an element y such that $x < y < z$. Note that the order \mathcal{G} is not dense, since there does not exist any graph X satisfying $K_1 < X < K_2$. However, we will see that this is the only exception to density for \mathcal{G} .

To prove this, we will need the following result:

Theorem 2.11.1. *For any integers g and c , there exists a graph X such that*

$$\text{og}(X) \geq g \text{ and } \chi(X) \geq c.$$

□

We will not give a proof of this, but it follows from the well known result of Erdős [13] which states the same thing except with odd girth replaced by girth. There are also constructive proofs of this result, including one which uses Kneser graphs [23]. Along with this result, we will also need the following construction which is important to the study of the homomorphism order.

Definition. For graphs X and Y , define the *exponential graph* Y^X to be the graph whose vertices are the functions from $V(X)$ to $V(Y)$ such that $f \sim g$ if

$$f(x) \sim g(x') \text{ whenever } x \sim x'.$$

The exponential graph possesses many interesting properties, but the most important property for our purposes is the following.

Theorem 2.11.2. *Let X , Y , and Z be graphs. Then $X \times Y \rightarrow Z$ if and only if $Y \rightarrow Z^X$.*

Proof. Suppose that φ is a homomorphism from $X \times Y$ to Z . For each $y \in V(Y)$, define $f_y : V(X) \rightarrow V(Z)$ as

$$f_y(x) = \varphi(x, y)$$

for all $x \in V(X)$. For $y \sim y'$, we see that if $x \sim x'$, then

$$f_y(x) = \varphi(x, y) \sim \varphi(x', y') = f_{y'}(x').$$

Therefore, $y \mapsto f_y$ is a homomorphism from Y to Z^X .

Conversely, suppose that ϕ is a homomorphism from $Y \rightarrow Z^X$. Let $g_y = \phi(y)$ for all $y \in V(Y)$. Define a function $\psi : V(X \times Y) \rightarrow V(Z)$ as follows:

$$\psi(x, y) = g_y(x)$$

for all $x \in V(X)$, $y \in V(Y)$. If $y \sim y'$ in Y , then $g_y \sim g_{y'}$ in Z^X and thus

$$\psi(x, y) = g_y(x) \sim g_{y'}(x') = \psi(x', y')$$

whenever $x \sim x'$. Therefore ψ is a homomorphism. □

With these two theorems we can give a quick proof that \mathcal{G} is dense everywhere above K_2 . The proof we present is originally due to Nešetřil [38].

Theorem 2.11.3. *Let X and Z be graphs such that $X < Z$ and $Z \not\cong K_2$. Then there exists a graph Y such that $X < Y < Z$.*

Proof. If $X \cong K_1$, then we can simply take Y to be K_2 . Otherwise, we have that $\chi(X) > 1$ and $\chi(Z) > 2$. Note that this means we can assume that no component of Z is bipartite.

By Theorem 2.11.1, there exists a graph W whose odd girth is greater than that of any component of Z and whose chromatic number is greater than $\chi(X^Z)$. Note that X^Z does not have loops since $Z \not\rightarrow X$. Let $Y = X \cup (Z \times W)$. Clearly, $X \rightarrow Y$ since Y contains X as a component. Also, $Y \rightarrow Z$ since $X \rightarrow Z$ by assumption and $Z \times W \rightarrow Z$ for any graph W . So we only need to show that $Z \not\rightarrow Y \not\rightarrow X$.

Since $Z \not\rightarrow X$, there exists some component Z_0 of Z such that $Z_0 \not\rightarrow X$. If $Z \rightarrow Y$, then we must have that $Z_0 \rightarrow Z \times W \rightarrow W$. However, W has greater odd girth than Z_0 by assumption, a contradiction. Therefore $Z \not\rightarrow Y$.

If $Y \rightarrow X$, then we have that $Z \times W \rightarrow X$ and thus $W \rightarrow X^Z$ by Theorem 2.11.2. However, $\chi(W) > \chi(X^Z)$ by assumption and so this is impossible. Therefore, $Y \not\rightarrow X$ and thus $X < Y < Z$. \square

Given two posets, \mathcal{P} and \mathcal{P}' , with corresponding partial orders \leq and \leq' , we say that \mathcal{P} *embeds order-preservingly into \mathcal{P}'* if there exists a function $f : \mathcal{P} \rightarrow \mathcal{P}'$ such that $f(x) \leq' f(y)$ if and only if $x \leq x'$. The poset \mathcal{P}' is said to be *universal* for a class \mathbb{P} of posets if every poset in \mathbb{P} embeds order-preservingly into \mathcal{P}' .

As an example, consider a partially ordered set \mathcal{P} with partial order \leq . The map

$$y \mapsto \{x \in \mathcal{P} : x \leq y\}$$

embeds \mathcal{P} order-preservingly into its power set ordered by inclusion. Therefore, the power set of $[n]$ ordered by inclusion is universal for the class of posets of size at most n .

We will only be concerned with countable posets which are universal for the class of all countable posets. We will refer to such a poset as simply being universal. Note that it is not clear that such an object even exists. However, the existence of a universal partially ordered set has been proven many times [31, 30, 14]. More pertinent to our work is the following theorem [41] which we state without proof.

Theorem 2.11.4. *The order \mathcal{G} is universal.* \square

In some sense, this theorem is evidence of the richness of the order \mathcal{G} . Interestingly, even suborders of \mathcal{G} which seem very restrictive retain this richness. For instance, the suborder of \mathcal{G} consisting of planar cubic graphs is universal [29]. Another important example for

us is that the graphs with a given chromatic number greater than two induce a universal suborder of \mathcal{G} . We will revisit the notions of density and universality in [Chapter 4](#).

Chapter 3

Cores of Vertex Transitive Graphs

My algebraic graph theory instructor once commented that there do not seem to be any “interesting” examples of vertex transitive graphs which are neither cores, nor have complete cores. What he meant by this is that all other examples of vertex transitive graphs seem to simply be constructed by taking several copies of a vertex transitive core and then adding some edges between these copies. Of course, this may simply be due to practical reasons. Indeed, many results on cores are simply sufficient conditions for a graph to be either a core, or have a complete graph as a core. This, along with the fact that determining the core of a graph is a hard problem in general [25], could be the reason for the lack of “interesting” examples.

These caveats aside, the above mentioned comment motivated the following question: “Can the vertex set of a vertex transitive graph always be partitioned into subsets, each of which induce a copy of the graph’s core?” Recalling Theorem 2.9.3 lends some plausibility to answering this question in the affirmative. However, in general the answer to this question is “no”, and we will see an infinite class of counterexamples later. It is worth noting, however, that these counterexamples all have complete graphs as cores. In this chapter we will see two classes of graphs for which we can answer the above question with a “yes”. The first class is all vertex transitive graphs that have cores half their size, and the second is the class of normal Cayley graphs, whose definition we will see in Section 3.3. Using the tools we develop to prove these results, we are also able to give an alternative proof of Theorem 2.9.3.

3.1 Half-Sized Cores

Considering the result of Theorem 2.9.3, it is natural to consider what can be said about vertex transitive graphs whose cores are half their size, since this is in some sense the simplest nontrivial case. As we will see below, such graphs can always be partitioned into two copies of their core. The proof uses the full power of Theorem 2.9.3, i.e. that the fibres of a retraction onto a core of a vertex transitive graph must all have the same size, not simply that the order of the core must divide the order of the graph.

Theorem 3.1.1. *Suppose X is a vertex transitive graph with a core X_1 such that $|V(X_1)| = \frac{1}{2}|V(X)|$. Furthermore, let φ be a retraction of X whose image is X_1 , and let X_2 be the subgraph of X induced by the vertices $V(X) \setminus V(X_1)$. If Y is the bipartite graph consisting of the edges of X which have exactly one end in each of $V(X_1)$ and $V(X_2)$, then we have the following:*

1. $X_1 \cong X_2$;
2. $\varphi|_{X_2}$ is an isomorphism from X_2 to X_1 ;
3. Y is regular and all of its edges are of the form $\{x, \varphi(y)\}$ where $x \sim_{X_2} y$.

Proof. Since X is vertex transitive, its core X_1 is also vertex transitive and therefore they are both regular. Let d and d_1 be the degree of vertices in X and X_1 respectively. This means that the $V(X_1)$ side of Y is regular with degree $d - d_1$.

First note that by Theorem 2.9.3 the fibres of φ all have the same size, in this case two. We claim that the restriction of φ to X_2 , denoted $\varphi|_{X_2}$, is a bijection from the vertices of X_2 to those of X_1 . Indeed, for any vertex $x \in V(X_1)$, the fibre $\varphi^{-1}(x)$ must contain two vertices. Since φ is a retraction, the vertex contained in $\varphi^{-1}(x)$ that is not x must be from $V(X) \setminus V(X_1) = V(X_2)$. Thus for all $x \in V(X_1)$, there exists a vertex $y \in V(X_2)$ such that $\varphi(y) = x$. This implies that $\varphi|_{X_2}$ is onto $V(X_1)$, and since $|V(X_1)| = |V(X_2)|$, it must be injective as well. Therefore, $\varphi|_{X_2}$ is a bijection between the vertices of X_2 and X_1 which preserves adjacency and thus X_2 is isomorphic to a spanning subgraph of X_1 . So the degree of any vertex in X_2 is at most d_1 and thus the degree in Y of a vertex in $V(X_2)$ is at least $d - d_1$. But of course this means that the degree in Y of every vertex in $V(X_2)$ is exactly $d - d_1$ since the sum of the degrees on one side of a bipartite graph is equal to the sum of the degrees on the other side. Therefore X_2 is regular with degree d_1 and thus must be isomorphic to X_1 , and furthermore, the restriction of φ to X_2 is an isomorphism from X_2 to X_1 .

Note that we have already shown that Y is regular with degree $d - d_1$. Hence, all that is left to show is that the edges of Y have the appropriate form. Consider a vertex $x \in V(X_2)$ which is adjacent in X to a vertex $y' \in V(X_1)$. Since φ is a retraction, $\varphi(x) \sim \varphi(y') = y'$. However, since the restriction of φ to X_2 is an isomorphism, if $y \in V(X_2)$ is such that $\varphi(y) = y'$, then $x \sim y$ in X_2 . \square

The partition from the theorem above actually satisfies another interesting property: it is equitable. A partition $\{C_1, C_2, \dots, C_k\}$ is said to be *equitable* if for all $i, j \in \{1, \dots, k\}$, the number of neighbors a given vertex from C_i has in C_j depends only on i and j , and not on the particular vertex of C_i . Note that this definition is equivalent to saying that each cell of the partition induces a regular graph, and the edges between any two cells form a semiregular bipartite graph. From this perspective, it is easy to see that the partition given in Theorem 3.1.1 above is equitable, since X_1 , X_2 , and Y are all regular. Given an equitable partition $\{C_1, \dots, C_k\}$, its *quotient matrix* is the matrix whose ij -entry is the number of neighbors a vertex of C_i has in C_j . It is known (see Lemma 9.3.3 of [19]) that the eigenvalues of the quotient matrix of any equitable partition of a graph X , are also eigenvalues of X . Using this result along with Theorem 3.1.1, we can give a necessary condition for a vertex transitive graph to have a core of half its size.

Corollary 3.1.2. *Let X be a vertex transitive graph with a core X_1 such that $|V(X_1)| = \frac{1}{2}|V(X)|$, and let X_2 be the induced subgraph of X with vertex set $V(X) \setminus V(X_1)$. Then $\{V(X_1), V(X_2)\}$ is an equitable partition of X . Furthermore, if d is the valency of X and d_1 is the valency of X_1 , then $2d_1 - d \geq 0$ is an eigenvalue of X .*

Proof. The fact that $\{V(X_1), V(X_2)\}$ is an equitable partition is obvious from the above proof. Since it is an equitable partition, the eigenvalues of its quotient matrix are eigenvalues of X . Since X_1 and X_2 are both regular with valency d_1 , and Y is regular with valency $d - d_1$, the quotient matrix of this partition is

$$\begin{pmatrix} d_1 & d - d_1 \\ d - d_1 & d_1 \end{pmatrix}.$$

Orthogonal eigenvectors for this matrix are $e_1 + e_2$ and $e_1 - e_2$ with eigenvalues d and $2d_1 - d$ respectively, where e_i is the i^{th} standard basis vector. Note that part (3) of Theorem 3.1.1 implies that the degree of Y is at most the degree of X_2 , and thus $d_1 \geq d - d_1$. This of course implies that $2d_1 - d \geq 0$ and this completes the proof. \square

Consider a vertex transitive graph X on $2p$ vertices for some prime p . By Theorem 2.9.3, its core must have 1, 2, p , or $2p$ vertices. If we further know that X is not bipartite, then we can rule out the first two of these four cases. If d is the valency of X , then by the

corollary above, if X has no nonnegative integer eigenvalue of the same parity as d , then it is its own core.

It would be nice to be able to more fully describe the structure of the subgraph Y in Theorem 3.1.1 above. For instance, is Y vertex transitive? A more precise description of Y could lead to a complete characterization of when a graph X is vertex transitive with a core of half its size.

For arc transitive graphs, we can fully characterize those whose core is half their size. Recall from Section 2.9 that the core of an arc transitive graph is arc transitive. Furthermore, by Theorem 2.9.4 the degree of the core of an arc transitive graph must divide the degree of the graph. Combining this with the result of Theorem 3.1.1 we are able to prove the following:

Theorem 3.1.3. *A graph X is arc transitive with a core half its size if and only if it is isomorphic to either $X_1 \square \overline{K_2}$ or $X_1 \boxed{K_2}$ for some arc transitive core X_1 .*

Proof. First note that if X_1 is arc transitive, then both $X_1 \square \overline{K_2}$ and $X_1 \boxed{K_2}$ are arc transitive. This is trivial for $X_1 \square \overline{K_2}$ since it is simply two disjoint copies of X_1 . To see that $X_1 \boxed{K_2}$ is arc transitive, note that for any $\sigma \in \text{Aut}(X_1)$ and set of permutations $\{\pi_x \in S_2 : x \in V(X)\}$, the map $(x, i) \mapsto (\sigma(x), \pi_x(i))$ is an automorphism of $X_1 \boxed{K_2}$. If $X \cong X_1 \square \overline{K_2}$, then X is simply two disjoint copies of X_1 . In this case it is trivial to see that $X \equiv X_1$, and since X_1 is a core this implies that it is the core of X . If $X \cong X_1 \boxed{K_2}$, then the vertices of X can be viewed as having the form (x, i) such that $x \in V(X_1)$ and $i \in \{1, 2\}$, and two vertices (x, i) and (y, j) are adjacent if $x \sim y$ in X_1 . The map $(x, i) \mapsto x$ can be easily seen to be a homomorphism and thus $X \rightarrow X_1$. On the other hand, the vertices $\{(x, 1) : x \in V(X_1)\}$ induce a subgraph of X isomorphic to X_1 and thus $X_1 \rightarrow X$. Therefore $X \equiv X_1$ and X_1 is the core of X .

Conversely, suppose that X is an arc transitive graph with core X_1 such that

$$|V(X_1)| = \frac{1}{2}|V(X)|.$$

Let X_2, Y , and φ be as in the proof of Theorem 3.1.1. As discussed above, X_1 must be arc transitive, and it is a core since it is the core of X . All that is left to show is that either $X \cong X_1 \square \overline{K_2}$ or $X \cong X_1 \boxed{K_2}$. If we let d and d_1 be the valency of X and X_1 respectively, then by the proof of Theorem 3.1.1 it is easy to see that $d \leq 2d_1$. However, as mentioned above, d_1 must divide d , and thus either $d = d_1$ or $d = 2d_1$. In the former case, the subgraph Y consisting of the edges with one end in each of X_1 and X_2 is empty and thus $X \cong X_1 \square \overline{K_2}$. In the latter case, Y consists of all edges of the form given in (3)

of Theorem 3.1.1. To see that X must be isomorphic to $X_1 \overline{[K_2]}$ in this case, consider the map

$$\psi(x) = \begin{cases} (\varphi(x), 1) = (x, 1) & \text{if } x \in V(X_1) \\ (\varphi(x), 2) & \text{if } x \in V(X_2). \end{cases}$$

Since the image of φ is $V(X_1)$, this map is a bijection from $V(X)$ to $V(X_1 \overline{[K_2]})$. To show that ψ is an isomorphism, we must show that it preserves adjacency and non-adjacency. If $x \sim y$ in X , then $\psi(x) = (\varphi(x), i) \sim (\varphi(y), j) = \psi(y)$ for some $i, j \in [2]$ and thus ψ preserves adjacency. Now suppose that $x \not\sim y$ in X . If $x, y \in V(X_1)$ or $x, y \in V(X_2)$, then $\psi(x) \not\sim \psi(y)$ since φ is a isomorphism when restricted to either $V(X_1)$ or $V(X_2)$. Otherwise, without loss of generality, $x \in V(X_1)$ and $y \in V(X_2)$, and there exists a vertex $z \in v(X_2)$ such that $\varphi(z) = x$. Suppose for contradiction that $\psi(x) \sim \psi(y)$, this implies that $\varphi(x) \sim \varphi(y)$ and thus $x \sim_{X_2} z$. But then $\{x, y\}$ is of the form given in (3) from Theorem 3.1.1, and thus $x \sim y$ which contradicts our assumption. Therefore ψ preserves adjacency and non-adjacency and thus is an isomorphism. \square

3.2 Attempts at Generalization

In light of the result of Theorem 3.1.1, the obvious next step would be to attempt to generalize the argument to show that all vertex transitive graphs can be partitioned into copies of their cores. The first obstacle one runs into when attempting this is that it is not obvious how to select the desired partition. In the above case, each fibre of the retraction onto the core contains exactly two vertices, and thus one can choose the vertices fixed by the retraction as one cell, and the remaining vertices as the other cell of the partition. If the fibres contain more than two vertices, then there is no “natural” choice for the partition, indeed there may be many possible choices.

The second obstacle one runs into while trying to generalize the above result is that it is simply not true. In fact, for any integer $n \geq 3$, there exists a vertex transitive graph X with core Y such that $|V(X)|/|V(Y)| = n$ and there does not exist any partition of $V(X)$ into subsets which each induce a copy of Y . To see this, consider the line graphs of the graphs K_{2n} , denoted $L(K_{2n})$, for $n \geq 3$. The maximum degree of K_{2n} is $2n - 1$, and therefore $L(K_{2n})$ contains a clique of size $2n - 1$, i.e. $K_{2n-1} \rightarrow L(K_{2n})$. On the other hand, it is known that the graphs K_{2n} can be $(2n - 1)$ -edge-colored and thus $L(K_{2n})$ can be $(2n - 1)$ -colored, i.e. $L(K_{2n}) \rightarrow K_{2n-1}$. From these two facts it follows that $L(K_{2n})$ is homomorphically equivalent to K_{2n-1} , and since K_{2n-1} is a core, it is the core of $L(K_{2n})$.

The cliques of $L(K_{2n})$ correspond to sets of edges of K_{2n} such that any two edges in the set have a common endpoint. Since $n \geq 3$ and thus $2n - 1 \geq 5$, it is easy to see that

any $(2n - 1)$ -clique in $L(K_{2n})$ must correspond to a set of edges incident to a particular vertex of K_{2n} . If C_x and C_y are the two $(2n - 1)$ -cliques of $L(K_{2n})$ corresponding to the edges of K_{2n} incident to vertices x and y of K_{2n} , then the vertex of $L(K_{2n})$ corresponding to the edge xy is contained in both C_x and C_y . Therefore no two $(2n - 1)$ -cliques of $L(K_{2n})$ are disjoint, and thus it cannot be partitioned into copies of its core. Noting that $|E(K_{2n})| = \binom{2n}{2}$, we see that $|V(K_{2n-1})| = \frac{1}{n}|V(L(K_{2n}))|$, and thus these graphs provide counterexamples to generalizing Theorem 3.1.1 to any relative size smaller than $1/2$.

3.3 Normal Cayley Graphs

Recall that for a group G and inverse closed subset $C \subseteq G \setminus \{1\}$, the Cayley graph $X(G, C)$ has vertex set G and $g \sim h$ if $gh^{-1} \in C$. We say that a Cayley graph is *normal* if its connection set is closed under conjugation by any group element, i.e. $gCg^{-1} = C$ for all $g \in G$. When studying vertex transitive graphs, Cayley graphs are a natural subclass to consider since, as we saw in Corollary 2.8.2, every vertex transitive graph is a retract of some Cayley graph.

In Section 2.8, we saw that the right translations, $f_a(g) = ga$, are automorphisms of $X(G, C)$. If X is normal, then the left translations are automorphisms as well. To see this note that $(ag)(ah)^{-1} = a(gh^{-1})a^{-1} \in aCa^{-1} = C$ if and only if $gh^{-1} \in C$. In fact, an equivalent definition of normal Cayley graph is that the right translations are automorphisms.

The proofs of the results in this section borrow ideas from the proof of Lemma 2.7.2 given in Section 2.8 and the following theorem from [18], whose proof we reproduce here for the convenience of the reader.

Theorem 3.3.1 (Godsil). *Let X be a normal Cayley graph. If $\alpha(X)\omega(X) = |V(X)|$, then $\chi(X) = \omega(X)$.*

Proof. Let C be a clique of X of size $\omega(X)$, and S an independent set of size $\alpha(X)$. As in the proof of Lemma 2.7.2 given in Section 2.8, the sets

$$S^{-1}c, \quad c \in C$$

are pairwise disjoint. Since X is normal, the inverse map is an automorphism and thus these sets are independent. If $\alpha(X)\omega(X) = |V(X)|$, then these independent sets form a coloring with $\omega(X)$ colors. \square

we will show that any normal Cayley graph can be partitioned into copies of its core. To prove this, we will need the following two lemmas which will also allow us to give a new proof of Theorem 2.9.3. The first is a simple yet useful lemma which applies to all graphs, not just those which are vertex transitive.

Lemma 3.3.2. *Let φ be an endomorphism of X such that $\varphi(x) = \varphi(y)$ for distinct $x, y \in V(X)$, and let w and z be two vertices which appear in some core of X . Then there is no automorphism of X which maps the pair $\{w, z\}$ to the pair $\{x, y\}$.*

Proof. Suppose σ is such an automorphism of X , and ρ is a retraction onto a core of X containing w and z . Then the endomorphism $\varphi \circ \sigma \circ \rho$ has at least one fewer vertex in its image than is in the core of X , a contradiction. \square

We can apply this lemma to Cayley graphs to obtain the following:

Lemma 3.3.3. *Let $X = X(G, C)$ be a Cayley graph. If φ is an endomorphism of X whose image is a core Y of X , and $y \in V(Y)$, then the sets $a^{-1}V(Y)$, for $a \in \varphi^{-1}(y)$, are mutually disjoint.*

Proof. Suppose not. Then there exist distinct $a, b \in \varphi^{-1}(y)$ and distinct $c, d \in V(Y)$ such that $a^{-1}c = b^{-1}d$. Consider the map $\sigma : V(X) \rightarrow V(X)$ given by $\sigma(x) = xc^{-1}a$. Note that this is an automorphism of X since it is a right translation. However, $\sigma(c) = a$ and $\sigma(d) = dc^{-1}a = b$ by the above. But this contradicts Lemma 3.3.2, since $\varphi(a) = y = \varphi(b)$ and both c and d appear in Y . \square

With these two lemmas in hand, the proof of our main theorem is simple and straightforward.

Theorem 3.3.4. *Let X be a normal Cayley graph and Y be a core of X . Then there exists a partition $\{V_1, \dots, V_k\}$ of $V(X)$ such that each V_i induces a copy of Y .*

Proof. Let φ be a retraction from X to Y . Further, let $A = \varphi^{-1}(y)$ for some $y \in V(Y)$. By Theorem 2.9.3, A has size $|V(X)|/|V(Y)|$, and thus $|A||V(Y)| = |V(X)|$. Combining this with the fact that the sets $a^{-1}V(Y)$, for $a \in A$, are mutually disjoint by Lemma 3.3.3, we see that these sets must in fact partition $V(X)$. Furthermore, since X is normal, left translations are automorphisms, and therefore each set $a^{-1}V(Y)$, for $a \in A$, induces a copy of Y . \square

3.4 Alternative Proof of Theorem 2.9.3

Here we present a new proof of Theorem 2.9.3 which uses the techniques of section 3.3. Recall from Section 2.6 that a multiple of a graph X is simply the lexicographic product of X with an empty graph, i.e. $X \overline{[K_m]}$. Also recall that any multiple of X is homomorphically equivalent to X and thus they have the same core. The proof we give will use the result of Theorem 2.8.1 which states that for any vertex transitive graph X , there exists some multiple of X which is a Cayley graph.

Theorem 2.9.3. *If X is a vertex transitive graph and φ is an endomorphism of X whose image is a core Y of X , then all of the fibres $\varphi^{-1}(y)$, $y \in V(Y)$, have the same size and thus $|V(Y)|$ divides $|V(X)|$.*

Proof. We first prove it for Cayley graphs. Suppose that $X = X(G, C)$ is a Cayley graph, and φ is an endomorphism onto a core Y of X . Then for any $y \in V(Y)$, the sets $a^{-1}V(Y)$, for $a \in \varphi^{-1}(y)$, are mutually disjoint by Lemma 3.3.3. This implies that $|\varphi^{-1}(y)| \leq |V(X)|/|V(Y)|$ is true for all $y \in V(Y)$. However, the average size of a fibre of φ is clearly $|V(X)|/|V(Y)|$, and thus we must have equality in the above inequality for all $y \in V(Y)$.

Now suppose that X is a vertex transitive graph and φ is an endomorphism whose image is a core Y of X . By the result of Sabidussi mentioned above, there exists $m \in \mathbb{N}$ such that $Z = X \overline{[K_m]}$ is a Cayley graph. The vertices of Z are of the form (x, i) for $x \in V(X)$ and $i \in [m]$. The map ρ given by $\rho(x, i) = (x, 1)$ is easily seen to be a retraction onto a subgraph X' of Z isomorphic to X . If we define a map $\hat{\varphi} : V(X') \rightarrow V(X')$ by $\hat{\varphi}(x, 1) = (\varphi(x), 1)$, then $\hat{\varphi} \circ \rho$ is an endomorphism of Z onto a subgraph Y' isomorphic to Y , and is thus an endomorphism onto a core of Z . Therefore, by the first part of the proof, the fibres of $\hat{\varphi} \circ \rho$ all have the same size. However, each fibre of $\hat{\varphi} \circ \rho$ clearly has size m times the size of the corresponding fibre of φ , and therefore all fibres of φ must have the same size. \square

This proof is reminiscent of the proof of the clique-coclique bound (Lemma 2.7.2) that we gave in Section 2.7. Both proofs begin by proving the result for Cayley graphs and then use Theorem 2.8.1 to generalize the result to all vertex transitive graphs. Furthermore, they both use the disjointness of certain subsets of vertices to obtain upper bounds on the size of other subsets of vertices. This similarity is not coincidental, in fact one can use the clique-coclique bound along with the notion of the *hull* of a graph introduced in [7] to give a short proof of the above theorem.

3.5 Remarks

The motivation behind this work is to take steps toward fully describing the structure of vertex transitive graphs in terms of their cores. In the above special cases of normal Cayley graphs and vertex transitive graphs with cores half their size, our results hint towards a product structure for these classes of graphs. However, the examples of $L(K_{2n})$ show us that not all vertex transitive graphs fit this description, and thus some more general result would be required to describe all vertex transitive graphs in terms of their cores. Currently, we do not know of any vertex transitive graphs which neither can be partitioned into copies of their core, nor have a complete graph as their core, and so it is an interesting question as to whether such graphs exist.

Another question of interest is when the *edge* set of a vertex transitive graph can be partitioned into copies of its core. A necessary condition for this is that the valency of the core must divide the valency of the graph, and thus arc transitive graphs are a natural class of graphs to consider with respect to this question. Note that the graphs $L(K_{2n})$ for $n \geq 3$, *can* have their edge sets partitioned into copies of their cores. The partition is given by the cliques of $L(K_{2n})$ corresponding to the vertices of K_{2n} .

Chapter 4

The Homomorphism Order of Line Graphs

There has been a great deal of research into the homomorphism order of graphs, and we have seen a small sampling of the results on this topic in Sections 2.10 and 2.11. For a much more detailed introduction to this area of graph homomorphisms, we recommend the book by Hell and Nešetřil [26]. Here, we consider the homomorphism order of graphs restricted to line graphs. The study of the homomorphism order of graphs has shed light on other areas of graph theory, especially with respect to chromatic number and other similarly defined parameters. We believe that systematic study of the homomorphism order of line graphs will analogously bear fruit in terms of results on parameters such as edge chromatic number or circular edge chromatic number which has been studied in [1] and [35]. Our focus here will be on comparing the properties of this suborder to those of the full order. In particular, we consider the question of density of this suborder. We give several specific results, but overall impression is that this order appears to be “richer” just above the complete graphs and “sparser” just below them. The rest of the chapter is outlined as follows.

We begin in Section 4.1 by defining line graphs and discussing some of their properties which will be important for the rest of the chapter. In particular we discuss the clique and chromatic numbers of line graphs. In Section 4.2, we formally define the homomorphism order of line graphs and translate several results of the previous section into the language of partial orders. We also show that the homomorphism order of line graphs can be partitioned into intervals which can be studied individually. Section 4.3 sets out to investigate the “lowest” few intervals discussed in the prior section. We show that the first two intervals

have a simple structure but the “higher” ones appear to be more complex. In particular, we give two line graphs that are incomparable in the homomorphism order.

In Section 4.4 we show that for any line graph Y such that $K_d < Y$, there exists a line graph X such that $K_d < X < Y$. In contrast to this result, we show in Section 4.5 that there does not exist any connected graph X such that $L(K_d) < L(X) < K_d$ for odd $d \geq 5$. We also give an example of a connected line graph less than K_4 for which there exists no connected line graph strictly between it and K_4 . In Section 4.6, we aim to determine properties of connected graphs X such that $L(X) < K_{d+1}$ and there exist no connected graph Y such that $L(X) < L(Y) < K_{d+1}$. We show that any such graph must be d -regular and $L(X)$ must be a core. Section 4.7 presents two connected graphs X and Z which satisfy $L(X) < L(Z)$, but for which there exists no connected line graph Y such that $L(X) < L(Y) < L(Z)$, and neither $L(X)$ nor $L(Z)$ are homomorphically equivalent to a complete graph. We end the chapter with some remarks on our results and a discussion of a few open problems in Section 4.8.

4.1 Basic Definitions and Properties

In this section we will introduce the basic properties of line graphs which will be important for the remaining sections in this chapter. The material in this section is well known and is not the work of the author. For more information on line graphs, see [19].

We will use the notation $e = xy$ to denote that e is the edge between vertices x and y in some graph. Technically we should write $e = \{x, y\}$, since edges are unordered pairs of vertices, but the more succinct notation is also more readable. Note that $e = xy$ and $e = yx$ mean the same thing since we ignore the order.

For a graph X , we say that two edges e and f of X are *incident* (to one another) if they share an endpoint, i.e. $e = xy$ and $f = xy'$ for some distinct $x, y, y' \in V(X)$. We do not say that an edge is incident to itself. We will also say, for $e \in E(X)$ and $x \in V(X)$, that e is incident to x (and similarly x is incident to e) if $e = xy$ for some $y \in V(X)$.

Given a graph X , we define the *line graph* of X , denoted $L(X)$ to be the graph whose vertices are the edges of X , such that two edges of X are adjacent in $L(X)$ if they are incident in X . We have given some examples of some graphs and their corresponding line graphs in Figure 4.1, and we have listed some special graphs along with their line graphs below.

- $L(P_n) \cong P_{n-1}$ for $n \geq 2$;

- $L(\overline{K_n})$ is the null graph (graph containing no vertices);
- $L(C_n) \cong C_n$
- $L(K_{1,n}) \cong K_n$




Figure 4.1: Examples of graphs and their line graphs.

An important fact to note is that if X is connected, then $L(X)$ must be connected. This is not difficult to see, since paths in X can be translated into paths in $L(X)$ in a straightforward manner. The converse is not always true: it is possible for $L(X)$ to be connected while X is not. However, this is only possible if X contains isolated vertices. Since removing isolated vertices from a graph X clearly does not change $L(X)$, we typically ignore isolated vertices by assuming that any graphs we discuss do not have them. Under this assumption, a graph X is connected if and only if $L(X)$ is.

When first encountering line graphs, it is natural to ask if every graph is the line graph of some graph. This is in fact not the case, and it is not too difficult to see that there exists no graph X such that $L(X) \cong K_{1,3}$. This also implies that the property of being a line graph is not closed under taking subgraphs, since $L(K_{1,4}) \cong K_4$ and $K_{1,3}$ is a subgraph of K_4 . However, if we restrict to induced subgraphs, this is not the case.

Lemma 4.1.1. *Let X be a graph. Any induced subgraph of $L(X)$ is a line graph of a subgraph of X .*

Proof. Let X be a graph and $Y = L(X)$. Furthermore, let Z be the subgraph of Y induced by $S \subseteq V(Y) = E(X)$, and let W be the subgraph of X consisting of the edges in S . We claim that $Z = L(W)$. Obviously, both Z and $L(W)$ have S as their vertex set. For distinct $e, f \in S$, we have that $e \sim f$ in Z if and only if $e \sim f$ in Y if and only if e and f

are incident in X . On the other hand, we have that $e \sim f$ in $L(W)$ if and only if they are incident in W if and only if they are incident in X . Therefore $Z = L(W)$. \square

Though the above lemma is simple and straightforward, we will need it for an important result in the next section.

The last item in the above list states that $L(K_{1,n}) \cong K_n$. Considering this, it is easy to see that for any graph X and vertex $x \in V(X)$, the edges incident to x form a clique in $L(X)$ of size equal to the degree of x in X . If we let $\Delta(X)$ denote the maximum degree of X , then the above observation implies that $\omega(L(X)) \geq \Delta(X)$. Below we will see that in fact equality holds here in all but one case (up to homomorphic equivalence).

Having considered cliques of line graphs, it is natural to consider colorings as well. Traditionally, colorings of line graphs are not studied directly, but rather through edge colorings of graphs. An n -edge-coloring of a graph X is an assignment of the elements of $[n]$ to the edges of X such that no two incident edges receive the same color. Note that in any edge coloring of a graph, each color class forms a set of pairwise non-incident edges. Such an object is typically known as a *matching*. Thus matchings are to edge colorings as independent sets are to vertex colorings. Also, the matchings in a graph X are exactly the independent sets in $L(X)$.

The minimum number of colors required for an n -edge-coloring of X is known as the edge chromatic number of X and is denoted $\chi'(X)$. Obviously, an n -edge-coloring of a graph X is equivalent to an n -coloring of the *vertices* of $L(X)$, and thus $\chi'(X) = \chi(L(X))$. From the discussion of cliques of line graphs above, it is clear that $\chi'(X) \geq \Delta(X)$ for all graphs X . On the other hand, the following theorem of Vizing's shows that the edge chromatic number of a graph X can never be greater than $\Delta(X) + 1$.

Theorem 4.1.2 (Vizing). *If X is a graph with maximum degree d , then X has a $(d + 1)$ -edge-coloring.*

Since $\chi'(X) \geq \Delta(X)$ for any graph X , there are only two possibilities for $\chi'(X)$: either $\Delta(X)$ or $\Delta(X) + 1$. Because of this, graphs which have edge chromatic number equal to their maximum degree are referred to as *class I* graphs, and all other graphs are referred to as *class II*. In the next section we will see that the class of a graph corresponds well with where it sits in the homomorphism order of line graphs. It is worth pointing out, that though determining the maximum degree of graph is obviously trivial, determining whether it is class I or II, and thus its edge chromatic number, is NP-hard [28].

We mentioned above that, except for a special class of graphs, the maximum degree of a graph is equal to the clique number of its line graph. One exceptional example is that of K_3 . Since $L(K_3) \cong K_3$, we have that $\omega(L(K_3)) = 3 = \Delta(K_3) + 1$. More generally, any

graph that has maximum degree two and contains a K_3 will clearly have this property. This does not really cause any problems, but to avoid having to constantly make exceptions, we address this class of graphs here, and implicitly exclude it from our discussion for the rest of this chapter. However, we will explicitly exclude it in our theorem statements. Let \mathcal{T} denote the set of all (isomorphism classes of) graphs having maximum degree two and containing K_3 as a subgraph. We have the following:

Lemma 4.1.3. *If $X \in \mathcal{T}$, then $L(X) \cong K_3$.*

Proof. It suffices to show that $L(X)$ both contains a clique of size 3 and is 3-colorable. Since X contains K_3 as a subgraph, and $L(K_3) \cong K_3$, we have that $L(X)$ contains a 3-clique. Since X has maximum degree 2, it is 3-colorable by Vizing's theorem. \square

Since $L(K_{1,3}) \cong K_3$ and $K_{1,3} \notin \mathcal{T}$, if we only consider line graphs of graphs not in \mathcal{T} , then we don't actually lose any line graphs. Because of this, in the remaining sections of this chapter, when we consider the line graph of an arbitrary graph X , we will implicitly assume that $X \notin \mathcal{T}$, unless specifically stated otherwise. Note that this does not mean that if we consider an arbitrary *line* graph Y , then $Y \notin \mathcal{T}$, as we will certainly need to consider this case for our work.

Note that above we only discussed why the graphs contained in \mathcal{T} are exceptional to the equation $\omega(L(X)) = \Delta(X)$, and we did not address why these are the only exceptions. The following lemma deals with this case.

Lemma 4.1.4. *If $X \notin \mathcal{T}$, then $\omega(L(X)) = \Delta(X)$.*

Proof. We have seen that $\omega(L(X)) \geq \Delta(X)$ for all graphs X , so we only need to show that $\omega(X) \leq \Delta(X)$ for $X \notin \mathcal{T}$. To accomplish this we will show that if $X \notin \mathcal{T}$, then there exists a vertex in X such that the edges incident to it form a maximum clique in $L(X)$.

Suppose that $S \subseteq E(X)$ is a maximum clique in $L(X)$ and there does not exist any vertex in X to which every element of S is incident. If no such S exists then we are done. The condition on S implies that S necessarily contains at least three elements. Suppose $e_1, e_2 \in S$. Then there exists $x \in V(X)$ such that $e_1 = xy_1$ and $e_2 = xy_2$ where $y_1 \neq y_2$. By our assumption on S , there exists $e_3 \in S$ that is not incident to x . Since e_3 is incident to e_1 but not x , it must be incident to y_1 . Similarly, e_3 must be incident to y_2 . Therefore $e_3 = y_2y_3$. It is easy to see that S cannot contain any other elements, since they would need to be incident to two of x, y_1, y_2 in order to be incident to all of e_1, e_2, e_3 , but then they would in fact be equal to one of e_1, e_2, e_3 , a contradiction. This implies that $\omega(X) = 3$. Since x, y_1, y_2 induce a K_3 in X , and $X \notin \mathcal{T}$, we have that X must have a vertex y of degree 3. The edges incident to y in X induce a K_3 in $L(X)$ and so we are done. \square

Note that the existence of the set S in the proof above implied that $\omega(L(X)) = 3$, and therefore if $\Delta(X) \geq 4$, then no such S exists and thus any maximum clique of X corresponds to the edges incident to some maximum degree vertex of X .

4.2 The Order

In this section we will formally introduce the homomorphism order of line graphs, and translate some of the results of the preceding chapter into results about this order.

Recall that the homomorphism order of graphs, which we denote by \mathcal{G} , consists of the set of all homomorphic equivalence classes of graphs along with the relation “ \leq ” defined as $\mathcal{H}(X) \leq \mathcal{H}(Y)$ if $X \rightarrow Y$. The *homomorphism order of line graphs*, which we denote by \mathcal{L} , consists of the homomorphic equivalence classes containing at least one line graph, equipped with the same relation as that of \mathcal{G} . It is easy to see that any subset of a partially ordered set is a partially ordered set under the same relation, and such an object is known as a *suborder*. Therefore, \mathcal{L} is a *suborder* of \mathcal{G} .

Recall that the order \mathcal{G} could equivalently be defined on cores, since each homomorphic equivalence class contains exactly one core. We can do the same for \mathcal{L} , but it would be nice to know that the cores we define it on are line graphs themselves, and not just homomorphically equivalent to some line graph. The following lemma addresses this concern.

Lemma 4.2.1. *The core of a line graph is a line graph.*

Proof. By Corollary 2.5.11, the core of a graph is an induced subgraph, and by Lemma 4.1.1, induced subgraphs of line graphs are line graphs. \square

Recall from Chapter 2 that we typically abuse notation and write $X \leq Y$ when we should write $\mathcal{H}(X) \leq \mathcal{H}(Y)$. We will take advantage of this notation and write $X < Y$ whenever $X \rightarrow Y$ and $Y \not\rightarrow X$. Furthermore, if $X \not\rightarrow Y$ and $Y \not\rightarrow X$, then we will write $X \parallel Y$.

The following lemma shows how the results of the previous section allow us to essentially break the order \mathcal{L} up into pieces which we can then study individually.

Lemma 4.2.2. *If $X \notin \mathcal{T}$, and the maximum degree of X is d , then*

$$K_d \leq L(X) < K_{d+1}.$$

Proof. It is easy to see that for any graph Y , we have $K_{\omega(Y)} \leq Y \leq K_{\chi(Y)}$. Since $\omega(L(X)) \geq d$ and $\chi(L(X)) \leq d + 1$ by Theorem 4.1.2, we have that $K_d \leq L(X) \leq K_{d+1}$. So all that is left to show is that $K_{d+1} \not\prec L(X)$. However, Lemma 4.1.4 states that $\omega(L(X)) = d$ and therefore the result holds. \square

What this lemma says is that the order \mathcal{L} can be partitioned into the “intervals” $[K_d, K_{d+1}]_{\mathcal{L}} = \{Y \in \mathcal{L} : K_d \leq Y < K_{d+1}\}$. Furthermore, the line graphs in the interval $[K_d, K_{d+1}]_{\mathcal{L}}$ are exactly the line graphs of graphs with maximum degree d . This implies that if $\Delta(X) < \Delta(Y)$, then $L(X) < L(Y)$. Therefore, to understand the structure of \mathcal{L} , it is sufficient to investigate these intervals individually.

We mentioned in the previous section that the class of a graph X gives us information about where in \mathcal{L} the graph $L(X)$ is located. The following result formalizes this statement.

Lemma 4.2.3. *If $X \notin \mathcal{T}$, then $L(X) \equiv K_d$ if and only if $\Delta(X) = d$ and X is class I.*

Proof. First note that if $\Delta(X) = d$ and X is class I, then $\omega(L(X)) = d = \chi(L(X))$ and thus $L(X) \equiv K_d$. Conversely, if $L(X) \equiv K_d$ for some $d \in \mathbb{N}$, then $\omega(L(X)) = d = \chi(L(X))$ and thus $\Delta(X) = d$ by Lemma 4.1.4 and X is class I since $\chi'(X) = d$. \square

This means that class I graphs correspond exactly to the line graphs which are homomorphically equivalent to complete graphs, i.e., the line graphs on the endpoints of the intervals given above. We noted above that given two graphs with different maximum degrees, it is trivial to deduce how their line graphs compare in the homomorphism order. If X and Y are graphs such that $\Delta(X) = \Delta(Y)$, then the class of the graphs can sometimes give us information about how they compare. In particular, if X is class I, then we immediately have that $L(X) \leq L(Y)$, and equality holds if and only if Y is class I as well. From this we see that if X and Y are such that $L(X) \parallel L(Y)$, then $\Delta(X) = \Delta(Y)$ and both X and Y are class II graphs.

4.3 The Intervals

In the previous section we saw that the order \mathcal{L} can be partitioned into the intervals $[K_n, K_{n+1}]_{\mathcal{L}}$ for $n \in \mathbb{N}$. Here, we take a look at the first three such intervals and discuss their structure. We will see that the interval $[K_2, K_3]_{\mathcal{L}}$ differs greatly from the same interval in \mathcal{G} , which we denote by $[K_2, K_3]_{\mathcal{G}}$, and similarly for other intervals.

The first interval we consider is $[K_1, K_2]_{\mathcal{L}}$. This case is not particularly interesting as it is easy to see that if $K_2 \not\prec X$, then X is an empty graph and therefore homomorphically

equivalent to K_1 . Thus we see that $[K_1, K_2]_{\mathcal{G}} = \{K_1\}$. Since $K_n \in \mathcal{L}$ for all n , we have that $[K_1, K_2]_{\mathcal{L}} = [K_1, K_2]_{\mathcal{G}} = \{K_1\}$. In fact, we have that $[K_1, K_2]_{\mathcal{L}} = [K_1, K_2]_{\mathcal{G}} = \{K_1, K_2\}$. This interval is actually a special case in \mathcal{G} since it represents the only *gap* in \mathcal{G} . A gap in a partially ordered set is an ordered pair, (X, Y) , of elements of the poset such that $X < Y$ and there does not exist any element Z such that $X < Z < Y$. We will see that \mathcal{L} differs from \mathcal{G} in this respect and has infinitely many gaps. Note that if (X, Y) is a gap we will say that it is a *gap below* Y and a *gap above* X .

The interval $[K_2, K_3]_{\mathcal{L}}$ is more interesting than the above case, but its structure is still quite simple. To see this, consider a graph X such that $L(X) \in [K_2, K_3]_{\mathcal{L}}$. As noted in the previous section, X must have maximum degree two, and is therefore the disjoint union of paths and cycles. This implies that $L(X)$ is also the disjoint union of paths and cycles. There are then two cases: $L(X)$ does not contain any odd cycles, or it does. In the former case, $L(X) \rightarrow K_2$ and thus $L(X) \equiv K_2$ since $L(X) \geq K_2$ by assumption. In the latter case, $L(X)$ has some shortest odd cycle C . Note that C must have length at least five since $K_3 \not\rightarrow L(X)$. Obviously, any component of $L(X)$ that is either a path or even cycle has a homomorphism to C . Furthermore, any odd cycle in $L(X)$ has length as least as great as that of C and therefore has a homomorphism to C as discussed in Section 2.3. This implies that $L(X) \equiv C$ and thus every line graph in $[K_2, K_3]_{\mathcal{L}}$ is homomorphically equivalent to either K_2 or some odd cycle of length at least five. Conversely, since $K_2 \in \mathcal{L}$ and $L(C_n) \cong C_n$ for all n , we have that

$$[K_2, K_3]_{\mathcal{L}} = \{K_2\} \cup \{C_n : n \geq 5, n \text{ odd}\}.$$

Furthermore, since $C_{2n+1} \rightarrow C_{2m+1}$ for $m \leq n$, the interval $[K_2, K_3]_{\mathcal{L}}$ is a *total order*, i.e. either $X \leq Y$ or $Y \leq X$ for all $X, Y \in [K_2, K_3]_{\mathcal{L}}$. This differs wildly from \mathcal{G} since the interval $[K_2, K_3]_{\mathcal{G}}$ is universal for countable posets.

Another significant difference between the two orders is that \mathcal{L} contains many gaps in this interval. Indeed, we can see from the above that there exists no graphs X such that $C_{2n+1} < L(X) < C_{2n-1}$ for $n \geq 2$, and therefore (C_{2n+1}, C_{2n-1}) is a gap in \mathcal{L} for all $n \geq 2$.

The interval $[K_3, K_4]_{\mathcal{L}}$ is not so easy to analyze as the two above. In fact, we do not know much at all about the structure of $[K_3, K_4]_{\mathcal{L}}$. The difficulty is simply that graphs with maximum degree three are not so easily characterized as those with maximum degree two. Though we are not able to say too much about its general structure, we have identified a pair of incomparable elements of $[K_3, K_4]_{\mathcal{L}}$. These two elements are the line graphs of the Petersen graph and a K_4 with one edge subdivided. Both are pictured along with their line graphs below. We will use the shorthand ‘‘Pete’’ to refer to the Petersen graph.

To show that these two line graphs are incomparable is somewhat tedious, but the following lemmas allow us to simplify things quite a bit. To succinctly state the first




Figure 4.2: The graph \hat{X} (left) and its line graph (right).




Figure 4.3: The Petersen graph (left) and its line graph (right).

lemma, we need some additional terminology. For graphs X and Y , and independent set $S \subseteq V(X)$, we say that a map $\varphi : V(X) \rightarrow V(Y)$ *identifies* the vertices in S if $\varphi(x) = \varphi(x')$ for all $x, x' \in S$. Furthermore, for any graph X and independent set $S \subseteq V(X)$, let X/S denote the graph with vertex set $(V(X) \setminus S) \cup \{x_S\}$ such that $x, y \in V(X) \setminus S$ are adjacent in X/S if $x \sim y$ in X , and x_S is adjacent to $z \in V(X) \setminus S$ in X/S if there exists $w \in S$ such that $w \sim z$ in X . In other words, X/S is the graph obtained from X by identifying all of the vertices in S .

Lemma 4.3.1. *Let X and Y be graphs and $S \subseteq V(X)$ be an independent set. If $\varphi : X \rightarrow Y$ is a homomorphism which identifies the vertices of S , then $X/S \rightarrow Y$.*

Proof. Let $y^* \in V(Y)$ be the vertex to which φ maps the elements of S . Consider the map

$\varphi' : V(X/S) \rightarrow V(Y)$ defined as follows

$$\varphi'(x) = \begin{cases} \varphi(x) & \text{if } x \in V(X) \setminus S \\ y^* & \text{if } x = x_S \end{cases}$$

Clearly, if $x_1, x_2 \in V(X) \setminus S$ are adjacent in X/S , then $\varphi'(x_1) = \varphi(x_1)$ is adjacent to $\varphi'(x_2) = \varphi(x_2)$ in Y . Furthermore, if $x \in V(X) \setminus S$ is adjacent to x_S in X/S , then there exists $x' \in S$ such that $x \sim x'$ in X and therefore $\varphi'(x) = \varphi(x)$ is adjacent to $\varphi'(x_S) = y^* = \varphi(x')$ in Y . \square

The above lemma is actually a special case of Proposition 2.11 in [23]. The lemma is somewhat obvious, but combined with the structure of \mathcal{L} it proves to be quite useful. One example of its usefulness is the following lemma.

Lemma 4.3.2. *Let Y be a line graph. If $Y < K_4$, then $W_{2n+1} \not\rightarrow Y$ for all $n \in \mathbb{N}$. Conversely, if $K_4 \leq Y$, then $W_{2n+1} \rightarrow Y$ for all $n \in \mathbb{N}$.*

Proof. Since $W_{2n+1} \rightarrow K_4$ for all $n \in \mathbb{N}$, then second statement holds trivially.

To prove the first statement, we use induction on n . If $n = 1$, then $W_{2n+1} = W_3 \cong K_4$ which does not have a homomorphism to Y by assumption. Now consider W_{2n+1} for $n \geq 2$. Let us refer to the vertex of W_{2n+1} of degree $2n + 1$ as w . Since $Y < K_4$, it is the line graph of a graph with maximum degree at most 3. Therefore Y has maximum degree at most 4. However, w has degree $2n + 1 \geq 5$ for $n \geq 2$. Therefore Y does not contain W_{2n+1} as a subgraph, and thus no homomorphism from W_{2n+1} to Y can be injective. Therefore, any homomorphism from W_{2n+1} to Y must identify two nonadjacent vertices $x_1, x_2 \in V(W_{2n+1})$. Since they are nonadjacent, neither of them can be w . It is easy to see that $W_{2n+1}/\{x_1, x_2\}$ must contain a strictly smaller odd wheel which by induction does not have a homomorphism to Y . Therefore, since $W_{2n+1}/\{x_1, x_2\}$ has no homomorphism to Y , by Lemma 4.3.1 neither does W_{2n+1} . \square

Using the above two lemmas makes proving the incomparability of $L(\text{Pete})$ and $L(\widehat{X})$ much less messy, but it still involves some tedious case analysis. We only give a partial proof below, but the remainder can easily be worked out by the zealous reader. The basic idea is to show that any homomorphism from $L(\text{Pete})$ to $L(\widehat{X})$ (and vice versa) must identify a certain number of vertices, and that identifying any choice of this number of vertices creates an odd wheel.

Theorem 4.3.3. *The line graphs $L(\text{Pete})$ and $L(\widehat{X})$ are incomparable.*

Proof. We must show that $L(\widehat{X}) \not\rightarrow L(\text{Pete})$ and $L(\text{Pete}) \not\rightarrow L(\widehat{X})$. We will start with the former.

First, note that Pete is edge transitive (it is in fact 3-arc transitive), and therefore $L(\text{Pete})$ is vertex transitive. Thus it is easy to check that the neighborhood of any vertex in $L(\text{Pete})$ induces two disjoint K_2 's, and so $L(\text{Pete})$ does not contain W_4 as a subgraph. However, $L(\widehat{X})$ clearly does contain a W_4 subgraph. Therefore $L(\widehat{X})$ is not a subgraph of $L(\text{Pete})$.

So any homomorphism from $L(\widehat{X})$ to $L(\text{Pete})$ must identify two nonadjacent vertices. We will show that identifying any two nonadjacent vertices of $L(\widehat{X})$ creates an odd wheel. Note that it suffices to show that the union of the neighborhoods of any two nonadjacent vertices contains an odd cycle. Also note that the symmetry of the graph can be used to reduce the number of cases to check.

Let x^* be the vertex in the middle of the above drawing of $L(\widehat{X})$, i.e., the vertex whose neighborhood induces a 4-cycle. If y is one of the vertices of $L(\widehat{X})$ of degree three, then the union of the neighborhoods of y and x^* is all of the five remaining vertices of $L(\widehat{X})$. It is easy to see that there is a 5-cycle on these vertices and thus we cannot identify y and x^* . The vertex x^* is adjacent to all of the vertices of $L(\widehat{X})$ that have degree four other than itself, so it cannot be identified with any vertex by a homomorphism to $L(\text{Pete})$.

Let y again be a vertex of $L(\widehat{X})$ of degree three, and let z be a vertex adjacent to x^* which is not adjacent to y . By symmetry it does not matter which one we choose. Again, the union of the neighborhoods of y and z is the remaining five vertices of $L(\widehat{X})$ and it is easy to see that there is a triangle among these vertices. Therefore y and z cannot be identified by any homomorphism to $L(\text{Pete})$.

The only remaining case is if z and z' are nonadjacent neighbors of x^* . Again, the union of their neighborhoods is the remaining five vertices of $L(\widehat{X})$ and it is easy to see that these vertices form a 5-cycle.

So we have shown that $L(\widehat{X}) \not\rightarrow L(\text{Pete})$.

The following proof that $L(\text{Pete}) \not\rightarrow L(\widehat{X})$ was suggested by Claude Tardif as an alternative to the author's original proof which was longer.

Note that every edge, and thus every pair of adjacent vertices, of $L(\text{Pete})$ is contained in a triangle. However, the edge between the two vertices of degree three in $L(\widehat{X})$ is not contained in any triangle. Therefore, no homomorphism from $L(\text{Pete})$ can map a pair of adjacent vertices to the two vertices of degree three in $L(\widehat{X})$. This implies that any homomorphism from $L(\text{Pete})$ to $L(\widehat{X})$ is also a homomorphism from $L(\text{Pete})$ to $L(\widehat{X})$ with that edge removed. However, removing that edge from $L(\widehat{X})$ causes the graph to become 3-colorable, and $L(\text{Pete})$ is not 3-colorable. Therefore there does not exist any homomorphism from $L(\text{Pete})$ to $L(\widehat{X})$. \square

The fact that the edge between the two vertices of degree three in $L(\widehat{X})$ is not contained in a triangle is crucial to the above proof. This is exemplified by the fact that $L(\text{Pete})$ does admit a homomorphism to the line graph of the graph \widehat{Y} in Figure 4.4 below, which is simply \widehat{X} plus one more edge. We leave it as an exercise for the reader to find the homomorphism.




Figure 4.4: The graph \widehat{Y} (left) and its line graph (right).

There are a few more things that we can say about the interval $[K_3, K_4]_{\mathcal{L}}$, but they apply more generally to \mathcal{L} , and so we will not see them until later sections.

4.4 Density

Recall that the interval $[K_2, K_3)_{\mathcal{L}}$ contained infinitely many gaps. In fact, we saw that every element of $[K_2, K_3)_{\mathcal{L}}$ other than K_2 is in some gap with another element of $[K_2, K_3)_{\mathcal{L}}$. On the other hand, K_2 is not in a gap with any other element of $[K_2, K_3)_{\mathcal{L}}$. Indeed, if $Y \in [K_2, K_3)_{\mathcal{L}}$ is not K_2 , then $Y \equiv C_{2n-1}$ for some $n \geq 3$, and we have $K_2 < C_{2n+1} < Y$. We refer to this property of \mathcal{L} as being *dense from above at K_2* . The natural question to ask is whether \mathcal{L} is dense from above (or below) at any other point. The main result of this chapter will be showing that \mathcal{L} is dense from above at every complete graph K_d for $d \geq 2$.

The main idea of the proof is as follows: Consider a graph Y such that $K_d < L(Y)$. If $\Delta(Y) > d$, then $K_d < L(X) < K_{d+1} \leq L(Y)$ for any class II graph X with $\Delta(X) = d$. Otherwise, $\Delta(Y) = d$ and Y is class II. In this case, we repeatedly replace the maximum degree vertices of Y with gadgets to obtain a new graph X , which has retained the property of being class II and is guaranteed to satisfy $L(X) \rightarrow L(Y)$. We then show that the process

used to obtain X has ensured that any subgraph of X with fewer than $|E(Y)| + 1$ edges necessarily has a d -edge-coloring, and thus $L(Y)$ cannot have a homomorphism to $L(X)$.

We first define the gadgets used for the construction of the graph X mentioned above. Note that a *perfect matching* of a graph Y is a matching M for which any vertex of Y is incident to some edge in M .

Definition. For $d \geq 3$ odd, define G_d to be the complete graph K_d . We will normally assume that the vertex set is $\{1, 2, \dots, d\}$. For $d \geq 2$ even, define G_d to be K_d with a perfect matching removed and an additional vertex that is adjacent to all the other vertices. Here we will usually assume that the vertex set is $\{1, 2, \dots, d, d+1\}$ where $d+1$ is the vertex of degree d .

Note that for all $d \geq 2$, the vertices $\{1, 2, \dots, d\}$ all have degree $d-1$. Also note that for d even, the graph G_d can equivalently be described as K_{d+1} with a maximum matching removed, since a maximum matching in K_{d+1} for d even contains edges incident to all but one vertex. Since K_{d+1} is maximum matching transitive, this construction does not depend on which maximum matching we remove. In order to prove that the construction outlined above works, we will need the following two lemmas concerning G_d .

Lemma 4.4.1. *The edges of G_d are d -colorable.*

Proof. For d odd, the maximum degree of G_d is $d-1$, therefore its edges are d -colorable by Vizing's Theorem. For d even, note that the edges of K_{d+1} are $(d+1)$ -colorable (again by Vizing theorem). Since any matching in K_{d+1} can be extended to a maximum matching, we can choose the maximum matching we remove from K_{d+1} to obtain G_d to contain a color class in an $(d+1)$ -coloring of K_{d+1} . This gives us an d -coloring of G_d . \square

In an edge coloring of a graph, we say that a vertex v "misses" a color if there are no edges of that color incident to v . We will also say that the color misses vertex v .

Lemma 4.4.2. *For every proper d -coloring of the edges of G_d , the vertices $\{1, 2, \dots, d\}$ all miss exactly one color. Furthermore, the colors missed by these vertices are distinct.*

Proof. For d odd, G_d is just the complete graph K_d . It is probably common knowledge that the above holds for this graph, but we will give a proof here for the reader's convenience. Every vertex of K_d has degree $d-1$, and since all the edges incident to a vertex must receive distinct colors, each vertex must miss exactly one color. Now note that there are $\binom{d}{2} = \frac{d(d-1)}{2}$ edges in K_d . For d odd, the maximum size of a matching in K_d is $\frac{d-1}{2}$. Since we are coloring the edges with d colors, it must be the case that every color class is a maximum matching, since otherwise not enough edges will have been colored. This means

that every color misses exactly one vertex, therefore each vertex misses a different color. Note that this means that G_d cannot be $(d - 1)$ -edge-colored.

For d even, the vertices $\{1, 2, \dots, d\}$ all have degree $d - 1$ and so they must miss exactly one color. Recall that in this case G_d is K_{d+1} with a maximum matching removed. Therefore the maximum size of a matching in G_d is $\frac{d}{2}$, so the most edges we can color with d colors is $\frac{d^2}{2}$. Coincidentally, G_d contains exactly $\frac{(d+1)d}{2} - \frac{d}{2} = \frac{d^2}{2}$ edges, and therefore each color class is a maximum matching. Therefore, every color misses exactly one vertex, and thus no two vertices in $\{1, 2, \dots, d\}$ can miss the same color. Note that G_d is not $d - 1$ edge-colorable in this case since it contains a vertex of degree d . \square

The construction of the graph X from the graph Y as mentioned above is actually several applications of a simpler construction, which we now define.

Definition. Let Y be a graph of maximum degree d , and y a vertex of Y of maximum degree whose neighbors are y_1, y_2, \dots, y_d . Define $T(Y, y)$ to be the graph constructed by removing y from Y , adding G_d , and an edge between i and y_i for all $i \in \{1, 2, \dots, d\}$.

An example of the procedure of constructing $T(Y, y)$ from Y is depicted in Figure 4.5 below. Note that $E(T(Y, y)) \cap E(Y) = E(Y) \setminus \{yy_i : i \in [d]\}$.




Figure 4.5: A portion of the graphs Y (left) and $T(Y, y)$ (right).

Note that this operation does not change the maximum degree of a graph, and in fact does not change the degree of any vertex of $V(Y) \setminus \{y\}$. We will also consider $T(Y, S)$ where S is a subset of vertices of Y which all have maximum degree. This will be defined in the obvious way: we simply remove all vertices of S and replace each of them with a copy of G_d as above. Note that this is equivalent to applying the operation to all of the vertices of S in sequence rather than in parallel, thus $T(Y, \{y_1, y_2\})$ is isomorphic to $T(T(Y, y_1), y_2)$, and similarly for larger S . We will also define $T(Y)$ to be $T(Y, S)$ for S equal to the set

of all maximum degree vertices of Y . We will use the notation $T^2(Y) = T(T(Y))$ and similarly for larger exponents.

The following lemma is key to showing that the graph X that we construct from Y will satisfy $L(X) \rightarrow L(Y)$.

Lemma 4.4.3. *For any graph Y with vertex of maximum degree y , there exists a homomorphism from $L(T(Y, y))$ to $L(Y)$.*

Proof. Though we are concerned with giving a homomorphism from $L(T(Y, y))$ to $L(Y)$, it is equivalent to map edges of $T(Y, y)$ to edges of Y while preserving incidence. Phrasing things in this way makes the argument easier to follow.

Define $f : E(T(Y, y)) \rightarrow E(Y)$ to be identity on the elements of $E(T(Y, y)) \cap E(Y)$, and let $f(iy_i) = yy_i$ for all $i \in \{1, 2, \dots, d\}$. For the remaining edges of $T(Y, y)$, recall from the proofs of Lemmas 4.4.1 and 4.4.2 that the edges of G_d can be partitioned into d maximum matchings each missing exactly one distinct vertex of $\{1, \dots, d\}$. Let M_i be the maximum matching that misses vertex i . Let $f(e) = yy_i$ for all $e \in M_i$ and $i \in \{1, \dots, d\}$. The function f is a homomorphism from $L(T(Y, y))$ to $L(Y)$. \square

The above lemma immediately gives us the following corollary:

Corollary 4.4.4. *For any graph Y and subset S of vertices of Y of maximum degree, there exists a homomorphism from $L(T(Y, S))$ to $L(Y)$.* \square

As we will need the graph X that we construct to satisfy $K_d < L(X)$, we must have that $L(X) \not\rightarrow K_d$, i.e. that X is a class II graph. Since Y is necessarily class II, it suffices to show that the function T preserves this property. The following lemma proves this.

Lemma 4.4.5. *A graph Y with vertex y of maximum degree is class I if and only if $T(Y, y)$ is class I.*

Proof. If Y is class I, then since $T(Y, y)$ has the same maximum degree, it is also class I by Lemma 4.4.3. Conversely, if $T(Y, y)$ is class I, then its edges can be colored with $d = \Delta(Y)$ colors. However, this means that the copy of G_d in $T(Y, y)$ is colored with d colors, and so by Lemma 4.4.2, the edge iy_i must be colored differently for each $i \in \{1, \dots, d\}$. However, if we color the edges of Y by keeping the same colors on the edges of $E(Y) \cap E(T(Y, y))$, and coloring edge yy_i in Y the same color as iy_i in $T(Y, y)$, then we will have properly colored the edges of Y with d colors. \square

As with Lemma 4.4.3, this lemma also proves the result for $T(Y, S)$:

Corollary 4.4.6. *Let Y be a graph and S a subset of vertices of Y of maximum degree. Then Y is class I if and only if $T(Y, S)$ is class I.* \square

We are almost ready to prove our main result, but first we need the following lemma.

Lemma 4.4.7. *Let y_1, y_2, \dots, y_d be the vertices of $T^m(K_{1,d})$ adjacent to the vertices of degree one. Then the distance from y_i to y_j for $i \neq j$ is at least m .*

Proof. We must first check that the y_i are well defined. Let z_1, z_2, \dots, z_d be the vertices of $K_{1,d}$ of degree one. Since applying T does not change the degree of any of the degree one vertices of $K_{1,d}$ and only creates more vertices of degree d , the vertices z_1, z_2, \dots, z_d are exactly the vertices of $T^m(K_{1,d})$ of degree one. Furthermore, after one application of T , the degree one vertices are adjacent to distinct vertices and no further applications of T can cause two of them to become adjacent to the same vertex. For the remainder of the proof, let y_i be the neighbor of z_i .

We now proceed by induction on m . For $m = 1$, the y_i correspond to the vertices $1, 2, \dots, d$ of G_d and thus are distinct and therefore are at distance at least 1. It now suffices to show that each application of T strictly increases the distance from y_i to y_j .

Suppose that $m \geq 2$ and consider a shortest path $y_i = x_0, x_1, \dots, x_n = y_j$ for $i \neq j$ in $T^m(K_{1,d})$. Let $f : V(T^m(K_{1,d})) \rightarrow V(T^{m-1}(K_{1,d}))$ be the function which fixes the vertices of degree one and maps a vertex of degree d in $T^m(K_{1,d})$ to the vertex of $T^{m-1}(K_{1,d})$ it was created from. Note that if $x \sim y$ in $T^m(K_{1,d})$, then either $f(x) = f(y)$ or $f(x) \sim f(y)$ in $T^{m-1}(K_{1,d})$. Therefore, after replacing consecutive strings of the same vertex with a single occurrence of that vertex, the sequence $f(x_0), f(x_1), \dots, f(x_n)$ gives a walk of length at most n from the neighbor of z_i in $T^{m-1}(Y)$ to the neighbor of z_j in $T^{m-1}(Y)$. To prove that this walk is strictly shorter than n , consider the vertex x_0 . By the definition of T , all of the neighbors of x_0 other than z_i must have been created from $f(x_0)$. In particular, since P was a shortest path, we have that $f(x_1) = f(x_0)$. This implies that at least one term of $f(x_0), f(x_1), \dots, f(x_n)$ must be removed to make it a walk, and therefore the walk is strictly shorter than n . \square

We are now ready to prove the main theorem.

Theorem 4.4.8. *For all $d \geq 2$, if Y is a graph such that $K_d < L(Y)$, then there exists a graph X such that $K_d < L(X) < L(Y)$.*

Proof. As noted above, if $\Delta(Y) > d$, then $K_d < L(X) < K_{d+1} \leq L(Y)$ for any class II graph X of maximum degree d . Therefore, we may assume that $\Delta(Y) = d$, and thus Y is class II. Let S be the set of all maximum degree vertices in Y . Now consider $X = T^m(Y)$,

where $m = |E(Y)| + 1$. By Corollary 4.4.6 we have that $K_d < L(X)$, and by Corollary 4.4.4 we have that $L(X) \leq L(Y)$. So we only need to show that there is no homomorphism from $L(Y)$ to $L(X)$. We will show, equivalently, that there exists no function from the edges of Y to the edges of X which preserves edge incidence. We will do this by contradiction, showing that the image of any such function must have a d -edge-coloring.

Consider a vertex y of Y with maximum degree. After applying T once, y has been replaced by either d or $d + 1$ vertices, all of degree d . We then apply T again and replace all of these with even more vertices of degree d . This continues until we obtain X , which contains either d^m or $(d + 1)^m$ vertices of degree d all originating from the vertex y . Let us refer to the set of these vertices as V_y and the set of edges in the subgraph induced by V_y as E_y (see Figure 4.6). There exist d vertices of V_y which correspond to the vertices y_1, y_2, \dots, y_d from Lemma 4.4.7. Let us refer to these vertices of V_y as y_1, y_2, \dots, y_d as well. Note that these are the only vertices of V_y incident to an edge not in E_y . Note that since the operation T does not change the degree of any vertex of degree less than d , any vertex of degree d in X must be contained in V_y for some $y \in S$.

Now suppose that $f : E(Y) \rightarrow E(X)$ preserves the incidence of edges. Let $E' \subseteq E(X)$ be the image of f . Note that $|E'| \leq |E(Y)|$. We will describe how to d -color the edges of E' .

First, consider the graph X with the edges of $E^* = \cup_{y \in S} E_y$ removed. In this graph, any vertex in some V_y has degree either zero or one, and all other vertices have degree equal to their degree in Y , which must be strictly less than d . By Vizing's theorem, this graph can be d -edge-colored. Fix some d -edge-coloring, g , of this graph.

We must now show that we can extend g to $E^* \cap E'$. Since the V_y are pairwise disjoint, the E_y are pairwise disjoint and no edge from one of them is incident to an edge of another. Therefore, it suffices to show that we can extend g to $E_y \cap E'$ for some $y \in S$.

Note that since the subgraph induced by V_y is a subgraph of $T^m(K_{1,d})$ and $K_{1,d}$ is d -edge-colorable, the edges of E_y are d -colorable as well. Let E_0 be the edges of $E_y \cap E'$ that are not reachable along edges of E' from any of the vertices y_1, y_2, \dots, y_d described above, and are thus not reachable along edges of E' from any edge of $E(X) \setminus E^*$. We can extend g to E_0 by simply d -coloring its edges according to some d -coloring of E_y . Now let E_i be the set of edges in $E_y \cap E'$ reachable from y_i using only edges in $E_y \cap E'$. Note that E_i cannot contain any edges incident to y_j for $j \neq i$, since by Lemma 4.4.7 the shortest path from y_i to y_j contained in E_y has length at least $m = |E(Y)| + 1 > |E'|$. This implies that there is exactly one edge of $E(X) \setminus E^*$ which is incident to any edge of E_i , specifically it is the edge incident to y_i that is not contained in E_y . Let us call this edge e . Since y_i is incident to only $d - 1$ edges of E_y , we can d -color the edges of E_i so that none of the




Figure 4.6: Example of V_y and corresponding E_i for $i \in \{0, 1, \dots, d\}$ for $d = 5$.

edges of E_i incident to y_i are assigned the color g assigns to e . This allows us to extend g to each E_i and thus $E' \cap E_y$ and thus E' .

Since g is a d -coloring of E' , the map $g \circ f$ is a d -edge-coloring of Y , a contradiction. Therefore, there is no homomorphism from $L(Y)$ to $L(X)$, and we are done. \square

It is worth noting that we could actually reduce the number of applications of T needed above to something on the order of $\log_2(|E(Y)|)$, but we did not feel it necessary.

4.5 Gaps

The previous section can be thought of as a look at how the homomorphism order of line graphs behaves just above the complete graphs. This section then is a look at the other side of the complete graphs. In contrast to the density result above, here we will see examples of connected graphs X such that $L(X) < K_{d+1}$ and there exists no connected graph Y such that $L(X) < L(Y) < K_{d+1}$. We will refer to such graphs as d -maximal graphs, the number d referring to the degree of X . Note that X being d -maximal does not necessarily imply that $(L(X), K_{d+1})$ is a gap in \mathcal{L} , since there may be a disconnected Y such that $L(X) < L(Y) < K_{d+1}$. However, X being d -maximal is equivalent to saying that $(L(X), K_{d+1})$ is a gap in the restriction of \mathcal{L} to connected graphs, which we denote by \mathcal{L}_c .

Below we will see that the graph K_{d+1} is d -maximal for all even $d \geq 4$. To prove this, we will need the following simple lemma regarding the cliques of line graphs.

Lemma 4.5.1. *Let X be a graph with maximum degree $d \geq 4$. If S and T are two distinct maximum cliques of $L(X)$, then $|S \cap T| \leq 1$.*

Proof. Recall from the discussion following Lemma 4.1.4 that any maximum clique in X must be the set of edges incident to some degree d vertex of X . Let x_S and x_T be vertices whose sets of incident edges are equal to S and T respectively. Since S and T are distinct, so are x_S and x_T . If S and T do not share any vertices, then we are done. Otherwise, any vertex of $L(X)$ contained in both S and T must be an edge of X incident to both x_S and x_T , which of course implies that it is the edge $x_S x_T$. Therefore there can be no two distinct vertices of $L(X)$ contained in both S and T . \square

The power of the above lemma is that it shows that if X and Y are graphs of maximum degree $d \geq 4$ and $\varphi : L(X) \rightarrow L(Y)$ is a homomorphism such that the images of two cliques in $L(X)$ share two or more vertices, then they must in fact coincide completely. We explain this idea rigorously and apply it in the following proof.

Theorem 4.5.2. *For even $d \geq 4$, the graph K_{d+1} is d -maximal.*

Proof. We must show that there exists no connected graph X such that $L(K_{d+1}) < L(X) < K_{d+1}$. First note that any such graph X must have maximum degree $d \geq 4$, and therefore the maximum cliques of $L(X)$ are exactly the sets of edges incident to some vertex of degree d of X by the discussion following Lemma 4.1.4.

Now suppose that $\varphi : L(K_{d+1}) \rightarrow L(X)$ is an injective homomorphism. Since K_{d+1} is regular of degree d , we have that $L(K_{d+1})$ is regular of valency $2d - 2$. Since φ is injective, $L(K_{d+1})$ is isomorphic to a subgraph Y of $L(X)$. Furthermore, since the maximum degree of $L(X)$ is at most $2d - 2$, the subgraph Y must be an induced subgraph of $L(X)$. This is because adding any edge to Y would create a vertex of degree greater than $2d - 2$, a contradiction. This also implies that Y must be a component of $L(X)$, since if any vertex of Y was adjacent to a vertex of $L(X)$ not contained in Y , then it would necessarily have degree strictly greater than $2d - 2$, a contradiction. Since $L(X)$ is connected, $Y = L(X)$ and therefore $L(X)$ is isomorphic to $L(K_{d+1})$ and we do not have $L(K_{d+1}) < L(X)$.

Now suppose that $\varphi : L(K_{d+1}) \rightarrow L(X)$ is a homomorphism that identifies two vertices of $L(K_{d+1})$. We may assume that $V(K_{d+1}) = \{y_i : i \in [d + 1]\}$, and thus $V(L(K_{d+1})) = \{y_i y_j : i, j \in [d + 1], i \neq j\}$. Let us refer to the set of edges of K_{d+1} incident to vertex i with the notation E_i . Without loss of generality, the homomorphism φ identifies $y_1 y_2$ and $y_3 y_4$. Note that since $\omega(L(K_{d+1})) = d = \omega(L(X))$, the homomorphism φ must map

any maximum clique of $L(K_{d+1})$ to a maximum clique of $L(X)$ bijectively. Let Y_1 and Y_3 be the cliques of $L(X)$ that φ maps E_1 and E_3 to respectively. We must have that $Y_3 \ni \varphi(y_3y_4) = \varphi(y_1y_2) \in Y_1$ as well as $\varphi(y_1y_3) \in Y_1 \cap Y_3$. However, y_1y_3 and y_3y_4 are adjacent in $L(K_{d+1})$ and thus $\varphi(y_1y_3) \neq \varphi(y_3y_4)$. Therefore Y_1 and Y_3 share two distinct vertices of $L(X)$ and thus $Y_1 = Y_3$ by Lemma 4.5.1.

Now consider E_i for $i \neq 1, 3$. Let Y_i be the clique of $L(X)$ that φ maps E_i to. We have that $\varphi(y_iy_1) \in Y_i \cap Y_1$ and $\varphi(y_iy_3) \in Y_i \cap Y_3 = Y_i \cap Y_1$. Since y_iy_1 and y_iy_3 are adjacent in $L(K_{d+1})$, we have that $\varphi(y_iy_1) \neq \varphi(y_iy_3)$, and therefore Y_i and Y_1 share two distinct vertices and thus $Y_i = Y_1$ by Lemma 4.5.1.

Therefore, every E_i gets mapped to Y_1 by φ . Since every vertex of $L(K_{d+1})$ is contained in some E_i and Y_1 is a clique of size d , we have that φ is a d -coloring of $L(K_{d+1})$, a contradiction. \square

Note that d being even is necessary in the above since K_{d+1} is not d -maximal for odd d as it is d -edge-colorable and thus $L(K_{d+1}) \equiv K_d$ in this case. The existence of a d -edge-coloring of K_{d+1} for d odd is also where the proof above fails, since we would not reach the contradiction required.

The 3-maximal graph, M_3 , we will see below is constructed by gluing together three copies of the graph \widehat{Y} from Section 4.3 on the vertex of degree one. To prove that M_3 is 3-maximal, we will need a lemma concerning homomorphisms from $L(\widehat{Y})$. Lemma 4.5.3 will show that any homomorphism from $L(\widehat{Y})$ to a line graph strictly less than K_4 must be injective. We will refer to this property as being *weakly 3-maximal*. In general, we will say that a connected graph X is *weakly d -maximal* if $K_d \leq L(X) < K_{d+1}$, and any homomorphism from $L(X)$ to a line graph strictly less than K_{d+1} is injective. Note that removing the $K_d \leq L(X)$ restriction would simply have the effect of including the graphs $K_{1,k}$ for $k < d$ in the definition. This does not really make a difference for our purposes, but we prefer to exclude these graphs. Also note that it is not obvious that a d -maximal graph is necessarily weakly d -maximal. For instance, if X were a d -maximal graph such that $L(X)$ were not a core, then $L(X)$ would have a non-injective homomorphism to itself, which would imply that it is not weakly d -maximal. However, Theorem 4.6.6 will show that $L(X)$ must be a core for any d -maximal graph X , and this will allow us to see that all d -maximal graphs are weakly d -maximal.

Lemma 4.5.3. *Let X be a graph such that $L(X) < K_4$. If $\varphi : L(\widehat{Y}) \rightarrow L(X)$ is a homomorphism, then φ is injective and $L(X)$ contains $L(\widehat{Y})$ as an induced subgraph.*

Proof. The proof of Theorem 4.3.3 shows that no such homomorphism can identify any two vertices contained in the $L(\widehat{X})$ subgraph of $L(\widehat{Y})$, therefore we need only show that φ

cannot identify the degree two vertex of $L(\widehat{Y})$ with any of its other vertices.

Let y^* be the vertex of $L(\widehat{Y})$ of degree two, and let x^* be the vertex in the top middle of the drawing of $L(\widehat{Y})$ in Figure 4.4. The union of the neighborhoods of x^* and y^* is the remaining six vertices of $L(\widehat{Y})$ which contains a 5-cycle. Therefore x^* and y^* cannot be identified by φ .

The only remaining case is identifying y^* with some neighbor z of x^* . However, for any choice of z , the union of the neighborhoods of y^* and z contains all vertices of $L(\widehat{Y})$ other than y^* , z , and the neighbor of x^* opposite of z in Figure 4.4. These five vertices induce a 5-cycle and thus y^* and z cannot be identified by φ .

Therefore φ must be injective and $L(X)$ must contain $L(\widehat{Y})$ as a subgraph. This subgraph must be induced since there is only one vertex of $L(\widehat{Y})$ which does not have degree 4, and therefore no edge can be added without increasing the degree of some vertex to 5. However, X must be of maximum degree at most 3, and therefore $L(X)$ has maximum degree at most 4, a contradiction. \square

We are now ready to show that the graph M_3 appearing in Figure 4.7 is a 3-maximal graph. The lemma above allows us to assume that any potential homomorphism from $L(M_3)$ is injective on each of the three $L(\widehat{Y})$ subgraphs it contains.




Figure 4.7: The graph M_3 (left) and its line graph (right).

Theorem 4.5.4. *There exists no connected graph X such that $L(M_3) < L(X) < K_4$.*

Proof. Let us use Z_1 , Z_2 , and Z_3 to denote the three disjoint copies of $L(\widehat{Y})$ contained in $L(M_3)$, and let us use z_i to denote the vertex of degree 2 in Z_i for $i \in \{1, 2, 3\}$. Suppose that X is as stated in the theorem and $\varphi : L(M_3) \rightarrow L(X)$ is a homomorphism. If φ is injective, then $L(X)$ contains $L(M_3)$ as subgraph, and this subgraph must be a component of $L(X)$ since every vertex of $L(M_3)$ has degree 4. Since X is connected, $L(M_3) \cong L(X)$ which is a contradiction to $L(M_3) < L(X)$.

Now suppose that φ is not injective. By Lemma 4.5.3, the restriction of φ to Z_i must be injective, and there must be an induced subgraph X_i of $L(X)$ such that φ acts as an isomorphism from Z_i to X_i for each $i \in \{1, 2, 3\}$. Let x_i be the degree two vertex of X_i . Note that $\varphi(z_i) = x_i$ for all $i \in \{1, 2, 3\}$. Also note that, unless $L(X) = X_i$, the vertex x_i is a cut vertex of $L(X)$. This is because the vertices in $V(X_i) \setminus \{x_i\}$ have degree 4 in X_i and thus are not adjacent to any vertices of $L(X)$ outside of X_i .

Since φ is not injective, without loss of generality φ must identify some vertex of Z_1 with some vertex of Z_2 , i.e. X_1 and X_2 must share some vertex. Since the X_i are connected and x_1 is the only vertex of X_1 adjacent to vertices of $L(X)$ outside of X_1 , we must have that x_1 is contained in X_2 . Similarly, x_2 is contained in X_1 , and since $z_1 \sim z_2$, the vertices x_1 and x_2 must be distinct.

Now suppose that X_1 and X_2 do not share all of their vertices. Then there exists a vertex x' of X_2 not contained in X_1 . However, removing x_1 from $L(X)$ will disconnect all of the vertices of X_1 , including x_2 , from the vertices not in X_1 , including x' . Since x_2 , x' , and x_1 are contained in X_2 , it must be that x_1 is a cut vertex of X_2 . But this is a contradiction since X_2 has no cut vertex.

Therefore, X_2 and X_1 must share all of their vertices. Since they are induced, this implies that the only degree two vertex of X_1 must also be the only degree two vertex of X_2 , i.e. $x_1 = x_2$ which contradicts the fact that φ is a homomorphism. \square

Note that though we have given an infinite number of examples of maximal graphs, we have not shown that any interval contains more than one. This raises the question of whether it is possible for there to be more than one d -maximal graph. Note that for $d = 2$, the answer to this question is “no” since C_5 is the only 2-maximal graph. However, we will see that for larger d the answer can be “yes”. In particular, we will show that the graph M_4 shown in Figure 4.9 is 4-maximal. Along with K_5 , this gives two examples of 4-maximal graphs. Let \widehat{Z} be the graph in Figure 4.8 produced by subdividing one edge of K_5 . We need the following lemma.

Lemma 4.5.5. *Let X be a graph such that $L(X) < K_5$. If $\varphi : L(\widehat{Z}) \rightarrow L(X)$ is a homomorphism, then φ is injective and $L(X)$ contains $L(\widehat{Y})$ as an induced subgraph.*




Figure 4.8: The graph \widehat{Z} .

Proof. First note that \widehat{Z} is not 4-edge-colorable. This can be seen by noting that the only matchings of \widehat{Z} of size 3 contain an edge incident to the vertex z . Only two such edges exist so the number of edges in 4 disjoint matchings is at most $10 < |E(\widehat{Z})|$. For a vertex $v \in V(\widehat{Z})$, let $C(v)$ denote the set of edges incident to v . We will use Lemma 4.5.1 to show that no two vertices of $L(\widehat{Z})$ can be identified by a homomorphism φ to a line graph strictly below K_5 .

We will think in terms of identifying edges of \widehat{Z} as opposed to vertices of $L(\widehat{Z})$.

First, suppose φ identifies y_1z with x_1x_2 . Then the images of $C(y_1)$ and $C(x_1)$ have at least two edges in common and thus must have all edges in common. Similarly the image of $C(x_2)$ coincides with the image of $C(y_1)$ and $C(x_1)$. Since the images of $C(x_1)$ and $C(x_2)$ coincide, they also coincide with $C(x_3)$ and $C(y_2)$. Since every edge of \widehat{Z} is contained in at least one of $C(x_1), C(x_2), C(y_1), C(y_2), C(y_3)$, the map φ must be a 4-edge-coloring of \widehat{Z} , a contradiction.

Now suppose that y_1z is identified by φ with x_1y_2 . In this case the images of $C(y_1)$ and $C(x_1)$ have at least two edges in common and thus must coincide. Let S be the shared image of $C(y_1)$ and $C(x_1)$. It follows that the images of $C(x_2)$ and $C(x_3)$ are both S . It then follows that the image of $C(y_2)$ is S and thus φ is a 4-edge-coloring of \widehat{Z} , a contradiction.

Suppose that φ identifies x_1y_1 and x_2x_3 . This implies that the images of $C(x_1)$ and $C(x_2)$ share at least two elements and thus coincide. Refer to their images as S . It follows that the images of $C(y_1), C(x_3)$, and $C(y_2)$ are all S and thus φ is a 4-edge-coloring of \widehat{Z} , a contradiction.

Finally, suppose φ identifies x_2y_1 and x_3y_2 . This implies that the images of $C(x_2)$

and $C(x_3)$ share at least two elements and thus coincide. This implies that the images of $C(x_1), C(y_1)$, and $C(y_2)$ coincide with the images of $C(x_2)$ and $C(x_3)$ and thus φ is a 4-edge-coloring of \widehat{Z} , a contradiction.

Up to symmetry, these are the only possible choices for identifying pairs of edges of \widehat{Z} . Therefore, φ must be injective and therefore $L(X)$ contains $L(\widehat{Z})$ as a subgraph. This subgraph must be induced since $L(\widehat{Z})$ contains only two vertices of degree less than 6 and they are adjacent. \square

We are now able to show the following theorem.




Figure 4.9: The graph M_4 .

Theorem 4.5.6. *The graph M_4 is 4-maximal.*

Proof. Let us use W_1 and W_2 to denote the two disjoint copies of $L(\widehat{Z})$ contained in $L(M_4)$.

Suppose that φ is a homomorphism from $L(M_4)$ to $L(X)$ where X is connected with maximum degree 4, and there does not exist a homomorphism from $L(X)$ to $L(M_4)$, i.e., $L(M_4) < L(X) < L(K_5)$. If φ is injective, then $L(X)$ contains $L(M_4)$ as subgraph, and this subgraph must be a component of $L(X)$ since every vertex of $L(M_4)$ has degree 6. Since X is connected, $L(M_4) \cong L(X)$ which is a contradiction to $L(M_4) < L(X)$.

Now suppose that φ is not injective. By Lemma 4.5.5, the restriction of φ to W_i must be injective, and there must be an induced subgraph X_i of $L(X)$ such that φ acts as an isomorphism from W_i to X_i for each $i \in \{1, 2\}$. Let u_i and v_i be the degree two vertices of X_i for each $i \in \{1, 2\}$. Note that unless $L(X) = X_i$, the set $\{u_i, v_i\}$ is a cut set of $L(X)$ of size 2. Furthermore, this is the only subset of vertices of X_i which is a cut set of $L(X)$ of size 2.

Since φ is not injective, at least one vertex of X_1 is also a vertex of X_2 . If X_2 contains a vertex not in X_1 , then removing $\{u_1, v_1\}$ from $L(X)$ disconnects X_2 , a contradiction since X_2 has no cut set of size two or less. Therefore X_1 and X_2 have the same vertex set. Of

course this implies that $\{u_1, v_1\} = \{u_2, v_2\}$. However, this is impossible since the vertex of W_1 mapped to u_1 is adjacent to the two vertices of W_2 mapped to u_2 and v_2 . Therefore φ cannot identify any two vertices of $L(M_4)$ but also cannot be injective and therefore cannot exist. \square

4.6 d -Maximal Graphs

In the previous section we introduced the notions of d -maximality and weak d -maximality, and we saw some specific examples of graphs with these properties. Here we will focus on how these two notions relate to one another. In particular, we will see that a graph is d -maximal if and only if it is weakly d -maximal and regular. Along the way to proving this, we will also see that if X is a d -maximal graph, then $L(X)$ must be a core.

One simple observation about d -maximal graphs is the following lemma:

Lemma 4.6.1. *If X and Y are d -maximal graphs, then either $L(X) \equiv L(Y)$ or $L(X) \parallel L(Y)$.*

Proof. Suppose that $L(X)$ and $L(Y)$ are not incomparable. Then without loss of generality, we have that $L(X) \leq L(Y)$. However, since X is d -maximal, we cannot have $L(X) < L(Y)$, and therefore we must have $L(X) \equiv L(Y)$. \square

We will revisit this lemma at the end of this section where we will see that we can significantly strengthen it using the results we present below.

We noted in the previous section that it is not clear that if X is d -maximal, then it must be weakly d -maximal. We will eventually show that this is true, but for now we will see that given an extra assumption on X it is easy to prove that it must be weakly d -maximal.

Lemma 4.6.2. *If X is d -maximal and $L(X)$ is a core, then X is weakly d -maximal.*

Proof. Suppose that X is d -maximal and $L(X)$ is a core. Furthermore, let Y be a graph such that $L(Y) < K_{d+1}$, and suppose that $\varphi : L(X) \rightarrow L(Y)$. Since X is connected, the image of φ is a connected subgraph Y' of Y . Note that $L(Y') < K_{d+1}$. By the definition of d -maximality, we must have that $L(X) \equiv L(Y')$, and therefore $L(X)$ is the core of $L(Y')$. This implies that any homomorphism from $L(X)$ to $L(Y')$ must be injective, and in particular φ is injective. \square

In the other direction, it is also possible to add an assumption to being weakly d -maximal that will guarantee d -maximality:

Lemma 4.6.3. *If X is weakly d -maximal and regular, then X is d -maximal.*

Proof. Suppose that X is weakly d -maximal and regular. Note that this implies that X is regular of degree d and therefore $L(X)$ is regular of degree $2d - 2$. Let Y be a connected graph such that $L(Y) < K_{d+1}$. Note that Y has maximum degree at most d and therefore $L(Y)$ has maximum degree at most $2d - 2$. Suppose that φ is a homomorphism from $L(X)$ to $L(Y)$. Since X is weakly d -maximal, φ must be injective, and therefore $L(X)$ is a subgraph of $L(Y)$. Of course, since $L(Y)$ has maximum degree at most $2d - 2$, this subgraph must be a component of $L(Y)$. Since Y is connected, this component is in fact all of $L(Y)$, and therefore $L(X) \cong L(Y)$. This trivially implies that $L(X) \equiv L(Y)$ and thus there exists no connected line graph strictly between $L(X)$ and K_{d+1} , i.e. X is d -maximal. \square

We remark that the above also implies that if X is weakly d -maximal and regular, and Y is a graph such that $L(X) < L(Y) < K_{d+1}$, then $L(Y)$ must contain $L(X)$ as a component. The converse of the above lemma also holds, giving us a necessary and sufficient condition for d -maximality in terms of weak d -maximality. However, to prove this we will need to show that if X is d -maximal, then it is regular and $L(X)$ is a core. In order to do this we will require the following lemma.

Lemma 4.6.4. *Let $X \not\cong K_3$ be a connected regular graph such that $L(X)$ is a core. Then $L(X)$ is the only connected line graph in $\mathcal{H}(L(X))$.*

Proof. Let d be the degree of X . It is easy to see that there is no regular graph, other than K_3 , whose line graph is a complete graph. Since $L(X)$ is a core, this implies that $L(X) \not\cong K_d$. Therefore, $K_d < L(X) < K_{d+1}$. Since X is regular of degree d , we have that $L(X)$ is regular of degree $2d - 2$.

Now suppose that $L(Y)$ is connected and homomorphically equivalent to $L(X)$. This implies that the core of $L(Y)$ is isomorphic to the core of $L(X)$. However, $L(X)$ is its own core and therefore $L(Y)$ contains a copy, Z , of $L(X)$ as an *induced* subgraph. Since $L(X)$ is regular of degree $2d - 2$, Z must be a component of $L(Y)$. However, $L(Y)$ is connected and therefore $L(X) \cong L(Y)$. \square

As an interesting corollary to the above lemma, we obtain a special case of our density result, Theorem 4.4.8, with a much simpler proof.

Corollary 4.6.5. *If Y is a d -regular graph and $L(Y)$ is a core, then there exists a graph X such that $K_d < L(X) < L(Y)$.*

Proof. We may assume that Y is connected since otherwise letting X be the graph obtained by removing one component of Y proves the result.

Let y be any vertex of Y , and let $X = T(Y, y)$. By the results of Section 4.4, we have that $K_d < L(X) \leq L(Y)$. So we only need to show that $L(X) \not\cong L(Y)$. However, by Lemma 4.6.4, $L(Y)$ is the only connected line graph in $\mathcal{H}(L(Y))$. Therefore, since $L(X)$ is connected, it cannot be homomorphically equivalent to $L(Y)$. \square

Our main theorem of this section says that if X is d -maximal, then X is d -regular and $L(X)$ is a core. The above lemma will allow us to assume that $L(X)$ is a core and prove that X must be d -regular in that case. The proof of this is quite long and requires several cases, but the basic idea is not difficult. We first suppose that X is d -maximal, $L(X)$ is a core, and X is not regular. Then, based on the particular case we are in, we add some edges and vertices to X to obtain a new graph X' that has the same maximum degree. Since X' contains X as a subgraph, we have that $L(X) \rightarrow L(X')$. However, since X is d -maximal, we must have that $L(X) \equiv L(X')$. Because $L(X)$ is a core, it must be the core of $L(X')$, and therefore there exists a retraction from $L(X')$ to $L(X)$ which fixes the vertices of $L(X)$. We then show that, based on the construction of X' , no such retraction can exist. This gives us our desired contradiction.

Theorem 4.6.6. *If X is d -maximal, then X is d -regular and $L(X)$ is a core.*

Proof. Note that we can assume that $d \geq 3$ since the only 2-maximal graph is C_5 which satisfies the theorem. We will first show that if Y is d -maximal and $L(Y)$ is a core, then Y is d -regular. We will then use Lemma 4.6.4 to prove the theorem for general d -maximal graphs.

Suppose that Y is a d -maximal graph such that $L(Y)$ is a core. Since Y is d -maximal, it has maximum degree d . If Y is not regular, then there exists a vertex $y \in V(Y)$ that has degree $d' < d$. We have several cases.

Suppose that $d' \geq 3$. Construct Y' from Y by adding a new vertex y' and an edge between y and y' . Clearly, Y' still has maximum degree d and is connected. Furthermore, $L(Y)$ is a subgraph of $L(Y')$ and thus $L(Y) \leq L(Y') < K_{d+1}$. Since Y is d -maximal, we must have that $L(Y) \equiv L(Y')$. However, since $L(Y)$ is a core, it is a core of $L(Y')$, and therefore there exists a retraction $\rho : L(Y') \rightarrow L(Y)$ which fixes the vertices of $L(Y)$. Thinking of ρ as a map from $E(Y')$ to $E(Y)$, we see that ρ must map the edge yy' to an edge of Y that is incident to all of the edges incident to y in Y . However, since $d' \geq 3$, no such edge can exist, a contradiction.

Suppose that $d' = 2$ and $d \geq 4$. Let x_1 and x_2 be the two neighbors of y in Y . Construct Y' from Y by adding two new vertices z_1 and z_2 and the edges yz_1 and yz_2 .

Since $d \geq d' + 2$, the maximum degree of Y' is still d . By the same argument above, there must exist a retraction $\rho : L(Y') \rightarrow L(Y)$ which fixes the edges of Y . Then ρ must map yz_1 to an edge of Y that is incident to both yx_1 and yx_2 . Only one such edge can possibly exist: an edge between x_1 and x_2 . However, even if this edge does exist, ρ must also map yz_2 to it by the same argument. But yz_1 and yz_2 are incident in Y' and therefore cannot be mapped to the same edge, a contradiction.

Suppose that $d' = 2$ and $d = 3$. Let x_1 and x_2 be the two neighbors of y and suppose that $x_1 \not\sim x_2$. Construct Y' from Y by adding a new vertex y' adjacent to y . As above, a retraction ρ must fix the edges of Y and map yy' to an edge of Y incident to both yx_1 and yx_2 . But since $x_1 \not\sim x_2$, no such edge exists, a contradiction.

By the arguments above, any vertex not of maximum degree in Y must have degree either 1 or 2, and if it has degree 2 then $d = 3$ and its two neighbors are adjacent.

Suppose that there are two vertices, y_1 and y_2 , of degree 2 in Y . Note that this implies that $d = 3$. We have two subcases.

Suppose that $y_1 \sim y_2$. Let x be the other neighbor of y_1 in Y . By the above, the two neighbors of y_1 must be adjacent, and therefore $y_2 \sim x$ as well. Since Y is connected and not K_3 , the vertex x must have some other neighbor, say z , in Y . The map which fixes the elements of $E(Y) \setminus \{y_1y_2\}$ and maps y_1y_2 to xz is a proper endomorphism of $L(Y)$, a contradiction to the fact that $L(Y)$ is a core.

Suppose that $y_1 \not\sim y_2$. Let x_{1a} and x_{1b} be the neighbors of y_1 , and let x_{2a} and x_{2b} be the neighbors of y_2 . Note that none of these four vertices can be equal to either y_1 or y_2 . By the above, we have that $x_{1a} \sim x_{1b}$ and $x_{2a} \sim x_{2b}$, but note that no other pairs of these four vertices are necessarily distinct. Construct Y' by adding the edge y_1y_2 . Since y_1 and y_2 had degree 2 in Y , the graph Y' still has maximum degree d . As above, there must exist a retraction ρ which fixes the edges of Y and maps the edge y_1y_2 to some edge of Y . Since y_1y_2 is incident to both y_1x_{1a} and y_1x_{1b} , it must be mapped by ρ to the edge $x_{1a}x_{1b}$. However, it also must be mapped to the edge $x_{2a}x_{2b}$ by the same argument. This is only possible if $x_{1a}x_{1b}$ and $x_{2a}x_{2b}$ are the same edge. This implies that without loss of generality $x_{1a} = x_{2a}$ and $x_{1b} = x_{2b}$. Therefore we have that x_{1a} is adjacent to the distinct vertices x_{1b}, y_1 , and y_2 . Similarly x_{1b} is adjacent to distinct vertices x_{1a}, y_1 , and y_2 . Since Y is connected and has max degree 3 and y_1 and y_2 have degree 2 in Y by assumption, the subgraph induced by the vertices y_1, y_2, x_{1a} , and x_{1b} is in fact all of Y . However, this graph is not even class II, and thus cannot be d -maximal.

Therefore we see that Y can have at most one vertex of degree two. Suppose that y is such a vertex and x_1 and x_2 are its neighbors in Y . By the above we have that $x_1 \sim x_2$. Note that this implies that any other vertex of Y must have degree either 1 or 3. Since x_1

is adjacent to both y and x_2 , it does not have degree 1 and thus has degree 3. Similarly, x_2 must have degree 3. Note that the edges yx_1 and yx_2 have degree 3 in $L(Y)$. Furthermore, these are the only vertices of $L(Y)$ with degree 3 since any edge of Y which is incident to exactly 3 edges must be incident to one vertex of degree 3 and one vertex of degree 2, i.e. it must be incident to y . Construct a graph Z as follows: First, take Y and a disjoint copy Y' of Y . Let y', x'_1, x'_2 be the copies of y, x_1, x_2 in Y' . Add a new vertex z which is adjacent to both y and y' . Again, there must exist a retraction $\rho : L(Z) \rightarrow L(Y)$ which fixes the edges of Y . Furthermore, by Lemma 2.5.2, the restriction of ρ to $L(Y')$ must be an isomorphism. Since yx_1 and yx_2 are the only vertices of $L(Y)$ with degree 3, and similarly for $y'x'_1$ and $y'x'_2$ in $L(Y')$, we must have that ρ maps the set $\{y'x'_1, y'x'_2\}$ bijectively to the set $\{yx_1, yx_2\}$. Since zy' is incident to both $y'x'_1$ and $y'x'_2$, it must be mapped by ρ to an edge of Y incident to both yx_1 and yx_2 , which can only be x_1x_2 . However, zy must also be mapped to x_1x_2 , which is a contradiction since zy and zy' are incident to each other.

So we have seen that Y cannot have any vertices of degree two. Therefore, if Y has any vertices which are not of maximum degree, then they must have degree one. We have two subcases.

Suppose that Y contains two vertices y_1 and y_2 of degree 1 which are not adjacent to the same vertex. Note that y_1 and y_2 are not adjacent since otherwise $Y \cong K_2$ or Y is disconnected, which is not the case. Construct Y' by identifying y_1 and y_2 , i.e. removing y_1 and y_2 and adding a new vertex adjacent to each of the neighbors of y_1 and y_2 . Let us refer to the edges that were incident to y_1 and y_2 in Y as e_1 and e_2 respectively. We consider the two edges incident to the new vertex of Y' as being the same e_1 and e_2 , even though they have different endpoints now. Therefore, we can view $L(Y')$ as having the same vertex set as $L(Y)$. The only difference is that e_1 and e_2 are adjacent in $L(Y')$ but not in $L(Y)$. We again must have a retraction $\rho : L(Y') \rightarrow L(Y)$ which fixes the vertices of $L(Y)$. However, this does not preserve the adjacency of e_1 and e_2 , a contradiction.

We are now left with the final case: The only non-maximum degree vertices of Y have degree 1 and are all adjacent to the same vertex. Let y_1, y_2, \dots, y_m be the vertices of degree 1 in Y and let x be their common neighbor. Note that x must have neighbors other than the y_i since otherwise $Y \cong K_{1,d}$ which is not d -maximal. This implies that $m < d$. Also note that the edges xy_i for $i \in [m]$ are the only vertices of $L(Y)$ with degree strictly less than $2d - 2$. Construct a new graph Z as follows: Let Y_1, Y_2, \dots, Y_d be d disjoint copies of Y . Identify all d copies of y_1 into a single vertex (giving it degree d), and do similarly for the d copies of each of the other y_i 's. Note that Z still has maximum degree d and is connected. Also note that Z contains d edge disjoint, but not vertex disjoint, copies of Y . We will still use Y_1, Y_2, \dots, Y_d to refer these copies of Y in Z . We will refer to the copy of x in Y_j as x_j for all $j \in [d]$. Since we identified each copy of y_i , we can refer to the vertex

they formed as y_i . Again, there must exist a retraction $\rho : L(Z) \rightarrow L(Y_1)$ which fixes the edges of Y_1 . As in a case above, the restriction of ρ to $L(Y_j)$ must be an isomorphism for each $j \in [d]$. Furthermore, since $x_j y_1, x_j y_2, \dots, x_j y_m$ are the only vertices not of maximum degree in $L(Y_j)$, they must be mapped bijectively to $x_1 y_1, x_1 y_2, \dots, x_1 y_m$ for all $j \in [d]$. In particular, this implies there exists a function $f : [d] \rightarrow [m]$ such that $\rho(x_j y_1) = x_1 y_{f(j)}$. However, since $d > m$, there must exist distinct $j, j' \in [d]$ such that $f(j) = f(j')$. This implies that $\rho(x_j y_1) = \rho(x_{j'} y_1)$, which is a contradiction since these two edge are incident in Z .

Therefore, we have shown that Y must be regular.

Now suppose that X is any d -maximal graph. Since the core of a line graph is a line graph, there exists a graph Y such that $L(Y)$ is the core of $L(X)$. Since $L(X)$ and $L(Y)$ are homomorphically equivalent and the core of a connected graph is connected, we have that Y must also be d -maximal. By the above, this implies that Y is d -regular. However, by Lemma 4.6.4, this implies that $L(Y)$ is not homomorphically equivalent to any connected line graph not isomorphic to itself. Therefore, $L(X) \cong L(Y)$ and thus $X \cong Y$ by Whitney's theorem (see Theorem 4.6.9 below). \square

Note that the necessary condition for X to be d -maximal given by the above theorem is not sufficient. To see this consider the Petersen graph. This graph is 3-regular and it is not too difficult to see that $L(\text{Pete})$ is a core. Indeed, since Pete is edge transitive, $L(\text{Pete})$ is vertex transitive. Therefore, by Theorem 2.9.3, its core must contain 1, 3, 5, or 15 vertices. It clearly is not homomorphically equivalent to K_1 , which is the only core on one vertex. Furthermore, the only cores on 3 or 5 vertices are K_3 , C_5 , and K_5 . Since Pete is not 3-edge-colorable, $L(\text{Pete})$ is not equivalent to K_3 . Since $L(\text{Pete})$ contains a triangle but no K_5 , it is not equivalent to either C_5 or K_5 . Therefore, it must be its own core. However, we have already seen that $L(\text{Pete})$ has a homomorphism to $L(\widehat{Y})$, and is thus not 3-maximal.

On the other hand, we are able to give a sufficient condition for d -maximality in terms of weak d -maximality.

Theorem 4.6.7. *A graph X is d -maximal if and only if it is weakly d -maximal and regular.*

Proof. We have already seen the backwards direction as Lemma 4.6.3, so it only remains to show the forward direction. If X is d -maximal, then by Theorem 4.6.6 it is regular and $L(X)$ must be a core. Since $L(X)$ is a core, by Lemma 4.6.2 the graph X is weakly d -maximal. \square

The above theorem finally allows us to say that all d -maximal graphs are also weakly d -maximal. Another consequence of Theorem 4.6.6 is the following:

Lemma 4.6.8. *Let X be d -maximal and Y be such that $L(Y) < K_{d+1}$. If $L(X) \rightarrow L(Y)$, then $L(Y)$ contains $L(X)$ as a component.*

Proof. Since X must be weakly d -maximal, $L(Y)$ must contain $L(X)$ as a subgraph. Furthermore, since X is regular of degree d , we have that $L(X)$ is regular of degree of $2d - 2$ and thus it must in fact be a component of $L(Y)$. \square

We can somewhat strengthen the above by using Whitney's theorem [51] which we state without proof below.

Theorem 4.6.9. *Let X and Y be connected graphs not isomorphic to either K_3 or $K_{1,3}$. Then $X \cong Y$ if and only if $L(X) \cong L(Y)$.* \square

Applying this to the above lemma we obtain:

Lemma 4.6.10. *Let X be d -maximal and Y be such that $L(Y) < K_{d+1}$. If $L(X) \rightarrow L(Y)$, then Y contains X as a component.*

Proof. Neither K_3 nor $K_{1,3}$ are d -maximal graphs. \square

A consequence of this lemma is that if X is a graph such that $(L(X), K_{d+1})$ is a gap in \mathcal{L} , then X must contain every d -maximal graph as components. Otherwise, letting Y be the disjoint union of X and one of the d -maximal graphs X does not contain would imply $L(X) < L(Y) < K_{d+1}$. In particular, this implies that if there are infinitely many gaps below K_{d+1} in \mathcal{L}_c , then there is no gap below K_{d+1} in \mathcal{L} . We will revisit the connection between gaps below K_{d+1} in \mathcal{L} and in \mathcal{L}_c in Section 4.8.

Recalling Lemma 4.6.1 from the beginning of this section, we can apply Theorems 4.6.6 and 4.6.9 to obtain the following strengthening

Lemma 4.6.11. *If X and Y are non-isomorphic d -maximal graphs, then $L(X) \parallel L(Y)$.*

Proof. If $L(X)$ and $L(Y)$ are not incomparable, then by Lemma 4.6.1 we have that $L(X) \equiv L(Y)$. However, by Theorem 4.6.6, both $L(X)$ and $L(Y)$ are cores and therefore by Lemma 2.5.1, $L(X) \cong L(Y)$. Applying Whitney's theorem shows that $X \cong Y$. \square

4.7 Another Gap

We have seen that \mathcal{L}_c contains gaps in infinitely many of the intervals $[K_d, K_{d+1}]$, and we have also seen that it can contain more than one gap in a given interval of this form.

However, all of the examples we have seen thus far have had a complete graph as the larger of the two graphs forming the gap. Therefore, the obvious question is: Does \mathcal{L}_c contain any gaps not of this form? In this section, we will see that the answer to this question is “yes”, by showing that $(L(\widehat{X}), L(\widehat{Y}))$ is a gap in \mathcal{L}_c .

Theorem 4.7.1. *There does not exist a connected graph Z such that $L(\widehat{X}) < L(Z) < L(\widehat{Y})$.*

Proof. Suppose that Z is such a graph. Since \widehat{X} is weakly 3-maximal, $L(Z)$ must contain $L(\widehat{X})$ as a subgraph. As $L(\widehat{X})$ only has two vertices of degree less than 4, and these two vertices are adjacent, $L(Z)$ must in fact contain $L(\widehat{X})$ as an induced subgraph. Therefore, by Lemma 4.1.1, Z contains a subgraph W whose line graph is isomorphic to $L(\widehat{X})$. By Whitney’s theorem this implies that W is isomorphic to \widehat{X} . Since $L(\widehat{X}) < L(Z)$, we must have that W is not all of Z . However, \widehat{X} contains only one vertex of degree less than 3, and therefore Z must contain a vertex adjacent to the degree 2 vertex of W , giving us a copy of \widehat{Y} contained in Z . Therefore $L(\widehat{Y}) \rightarrow L(Z)$, contradicting our assumption. \square

It is not too difficult to see that the above proof technique can be generalized to other weakly d -maximal graphs whose low-degree edges are all incident to one another. For instance let X be a K_5 with one edge subdivided, i.e. one half of M_4 , and let Y be X plus an extra vertex adjacent to the degree two vertex of X . Then, an identical proof to the one above shows that $(L(X), L(Y))$ is a gap in \mathcal{L}_c . Note that we actually need to show that $L(X) < L(Y)$ first, but this is easy to check.

Though $(L(\widehat{X}), L(\widehat{Y}))$ is a gap in \mathcal{L}_c , this does not imply that it is a gap in \mathcal{L} . Indeed, if we let Z be the disjoint union of \widehat{X} and the Petersen graph, then we have that $L(\widehat{X}) \rightarrow L(Z) \rightarrow L(\widehat{Y})$ since \widehat{X} is a subgraph of Z and both $L(\widehat{X})$ and $L(\text{Pete})$ have homomorphisms to $L(\widehat{Y})$. Furthermore, $L(\widehat{Y}) \not\rightarrow L(Z) \not\rightarrow L(\widehat{X})$ since $L(\text{Pete}) \not\rightarrow L(\widehat{X})$ and $L(\widehat{Y})$ does not have a homomorphism to either of $L(\widehat{X})$ or $L(\text{Pete})$. Therefore $(L(\widehat{X}), L(\widehat{Y}))$ is not a gap in \mathcal{L} .

4.8 Remarks and Open Questions

The main results of this chapter are that the order \mathcal{L} is dense from above at the complete graphs, and that the order \mathcal{L}_c contains gaps below the even complete graphs as well as K_3 . One of the most intriguing things about these results is that they seem to mirror similar results for circular chromatic numbers of line graphs. In particular, they show

in [1], that there exists no line graph whose circular chromatic number lies in the interval $(11/3, 4)$. It has also been asked by Xuding Zhu whether there exists $\varepsilon = \varepsilon(n) > 0$ such that there exists no line graph with circular chromatic number in $(n - \varepsilon, n)$ for all $n \in \mathbb{N}$. In contrast, it was shown in [35] that for any rational number $r \in [3, 1/3)$, there exists a line graph with circular chromatic number equal to r . It has also been shown [52] that for $r \in [2k + 1, 2k + 1 + \frac{1}{4}]$, there exists a line graph with circular chromatic number r . Together, these results seem to suggest that the homomorphism order of line graphs becomes “richer” as you approach the complete graphs from above, and conversely that it becomes “sparse” as you approach the complete graphs from below. This idea also seems to fit well with our results on density and gaps in the orders \mathcal{L} and \mathcal{L}_c .

There are of course many questions concerning \mathcal{L} and \mathcal{L}_c left unanswered by our results. We take some time here to address those which we feel are the most relevant to our work.

Perhaps the most obvious question is whether there are any other places where \mathcal{L} is dense. The technique for showing that \mathcal{L} is dense from above at the complete graphs does not seem to easily generalize to other places in \mathcal{L} . One difficulty in particular is that the gadgets we used to construct the interpolating graph very much depended on the fact that one of our graphs was the complete graph. Thus if we are given graphs X and Z such that $L(X) < L(Z)$, and we want to construct a graph Y such that $L(X) < L(Y) < L(Z)$ in the same manner as in Section 4.4, then the gadgets we use would need to depend on X in some way. In particular, if we let G be the gadget we use, and $E = \{e_1, e_2, \dots, e_d\}$ be the d edges of G which correspond to the d edges incident to the vertex G is replacing, then we would need that any incidence preserving map from the edges of G to the edges of X must map E bijectively to the set of edges incident to some degree d vertex of X . This does not appear to be easy to guarantee for a general graph X .

Another question related to the work above is whether or not \mathcal{L} contains any gaps other than those in the interval $[K_2, K_3]$. We noted above in Section 4.6, that if there are infinitely many d -maximal graphs, then there is no gap below K_{d+1} in \mathcal{L} . In fact, there is a gap below K_{d+1} in \mathcal{L} if and only if there is a finite number of d -maximal graphs, and every line graph strictly less than K_{d+1} has a homomorphism to the line graph of some d -maximal graph. It is straightforward to show that any line graph less than K_{d+1} has a homomorphism to some weakly d -maximal graph. However, it is not clear how to show that the line graph of any weakly d -maximal graph has a homomorphism to some d -maximal graph. It does not seem too unlikely that there can only be finitely many d -maximal graphs for a given d . However, it is not obvious how one would prove this. Possibly one could show that for a given $d \in \mathbb{N}$, a d -maximal graph can have at most some fixed number of vertices.

One property of \mathcal{G} that we discussed in Chapter 2 but did not address above, is that of \mathcal{G} being a lattice. It is natural to ask if \mathcal{L} is a lattice as well, but we are unsure of the answer. Obviously, since the disjoint union of two line graphs is a line graph, any two elements of \mathcal{L} have a join in \mathcal{L} . However, it is not clear that any two elements of \mathcal{L} have a meet in \mathcal{L} . Recall that the meet of two elements of \mathcal{G} corresponded to the categorical product of graphs. The problem with applying the same construction to define a meet operation on \mathcal{L} is that the categorical product of two line graphs is not necessarily a line graph. To see this note that the minimum eigenvalue of the adjacency matrix of a line graph is at least -2 [27], and the eigenvalues of the categorical product of two graphs are the products of their individual eigenvalues. Of course, this does not mean that the core of the categorical product of two line graphs is not a line graph. If this is the case then \mathcal{L} is in fact a lattice with the same meet and join operations as \mathcal{G} . On the other hand, if there exists a pair of line graphs such that the core of their categorical product is not a line graph, this still does not necessarily mean that \mathcal{L} is not a lattice. It may be possible that \mathcal{L} is a lattice but has a different meet operation than \mathcal{G} . As an illustration of this possibility, note that if \mathcal{L}_c does have a join operation, then it cannot be the same as \mathcal{G} , since the disjoint union of graphs is not connected. However, the graphs $L(\widehat{X})$ and $L(\text{Pete})$ of \mathcal{L}_c do in fact have a meet in this poset: $L(\widehat{Y})$. Indeed, if Z is a connected graph such that $L(\text{Pete}) \leq L(Z)$ and $L(\widehat{X}) \leq L(Z)$, then Z must contain \widehat{X} , but also cannot be simply isomorphic to \widehat{X} . Therefore, Z contains \widehat{Y} and thus $L(\widehat{Y}) \leq L(Z)$.

Through private communication we have recently learned that a result of Robert Šámal [49] implies the existence of an infinite family of cubic graphs whose line graphs are pairwise incomparable. However, we have not had time to verify this yet and so we have not included it in this thesis.

Chapter 5

Quantum Information Background

The purpose of this chapter is to acquaint the uninitiated reader with the basics of quantum information theory. As this is background material, none of it represents original work by the author. The topics covered will include quantum states, measurements, superposition, shared quantum systems, and entanglement. This material serves as background for the results on quantum homomorphisms discussed in the next chapter. However, our approach to quantum homomorphisms is such that the reader needs as little knowledge of quantum theory as possible. The material of this chapter will only be directly used in Section 6.5, where we give the original motivation for considering quantum homomorphisms. Though the exposition below does not follow any particular source directly, the book [39] was a helpful reference.

We begin by discussing how to mathematically describe the state of a classical system. We then introduce the mathematical description of “pure states”, the simplest type of quantum states. We discuss the possible valid quantum operations and define the simplest type of quantum measurement, the full basis measurement. The next section defines density matrices which are used to describe general quantum states. We also define general quantum measurements and give an example of performing such a measurement on a quantum state. The last section investigates shared quantum systems. In particular, we describe what happens when one party performs a measurement on their part of a shared system. We end with an example of two parties performing a measurement on a shared state which guarantees that their outcomes are always the same.

5.1 Introduction to Quantum Theory

Consider a physical system which has n possible states, for instance, the upward facing face of an n -sided die. We can represent the die having the face with i on it up with the i^{th} standard basis vector e_i . If someone else were to roll the die and not let us see it, how would we describe its state? If the probability of rolling an i is $p_i \in [0, 1]$, then this is equivalent to saying that there is probability p_i that the die is in state i . A natural way to describe this mathematically is by saying the die is in state

$$\sum_{i=1}^n p_i e_i = (p_1, \dots, p_n).$$

Obviously, the specific values of the p_i can vary, but since they are probabilities, they must all be in the interval $[0, 1]$ and satisfy $\sum_{i=1}^n p_i = 1$. Conversely, any vector whose entries satisfy these properties can be viewed as a vector representing the state of some system. Vectors of this type are sometimes known as *stochastic vectors*.

Now that we know what the possible vectors representing states of our system are, we would like to know what kind of operations map state vectors to state vectors. If P is a matrix which maps stochastic vectors to stochastic vectors, then it is not difficult to see that the entries of P must all be nonnegative, and the sums of each column must be equal to one. Such a matrix is known as a *left stochastic matrix*.

As an example, suppose Alice flips a fair coin and does not let Bob know the outcome. If e_1 and e_2 represent heads and tails respectively, then Bob would describe the state of the coin as $1/2(e_1 + e_2)$. If Alice then tells Bob that if the coin had landed on heads she switched it to tails with probability $1/4$, and if the coin had landed on tails she switched it to heads with probability $1/6$, then how would Bob describe the state of the coin? The probability it ended on heads is equal to the probability it landed on heads and Alice did not switch it ($1/2 \cdot 3/4 = 3/8$) plus the probability that it landed on tails and Alice did switch it ($1/2 \cdot 1/6 = 1/12$) which gives $3/8 + 1/12 = 11/24$. The probability of tails must be $1 - 11/24 = 13/24$. So the state vector of the coin is

$$\begin{pmatrix} 11/24 \\ 13/24 \end{pmatrix}$$

However, we can think of the actions made by Alice after the flip as a stochastic matrix. Switching heads to tails with probability $1/4$ corresponds to sending the vector e_1 to the

vector $(3/4, 1/4)^T$. Switching tails to heads with probability $1/6$ corresponds to sending e_2 to $(1/6, 5/6)^T$. So the matrix associated with this action is

$$\begin{pmatrix} 3/4 & 1/6 \\ 1/4 & 5/6 \end{pmatrix}.$$

It is not hard to see that

$$\begin{pmatrix} 3/4 & 1/6 \\ 1/4 & 5/6 \end{pmatrix} \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix} = \begin{pmatrix} 11/24 \\ 13/24 \end{pmatrix}.$$

The state of a quantum system is also described by vectors. As in the probabilistic case, the deterministic states coincide with the standard basis vectors. Unlike in the probabilistic case, the other possible states are given by the following physical principle known as *superposition*:

Principle of Superposition: If a quantum system can be in orthogonal states ψ_1 and ψ_2 , then it can be in any state $\alpha_1\psi_1 + \alpha_2\psi_2$ such that $\alpha_1, \alpha_2 \in \mathbb{C}$ and $|\alpha_1|^2 + |\alpha_2|^2 = 1$.

From this principle it is easy to see that the possible quantum states are the complex unit vectors in the ℓ_2 -norm. (Actually, to describe the most general quantum states we will need a different mathematical object, but we will not introduce this until the next section. The states which can be described by vectors are referred to as *pure states*.) Since the vectors which represent quantum states are different from those describing classical states, the operations mapping quantum states to quantum states will be different as well. Suppose that U is a matrix which maps quantum states in \mathbb{C}^n to quantum states in \mathbb{C}^n . Clearly, U must be norm-preserving. Therefore, for any unit vector ψ ,

$$1 = \|U\psi\| = (U\psi)^*(U\psi) = \psi^*(U^*U)\psi,$$

where $*$ denotes the conjugate transpose. This implies that every nonzero vector is an eigenvector of U^*U with eigenvalue 1, and thus $U^*U = I$. As U is a square matrix, we have that $UU^* = I$ as well. Such a matrix U is known as a *unitary matrix*. Any unitary matrix can be easily shown to be norm preserving, and so the valid quantum operations are unitaries.

We will introduce general quantum measurements in Section 5.2, but here we introduce the simplest type of measurement: the full basis measurement. Given an (ordered) orthonormal basis $\mathcal{B} = (u_1, u_2, \dots, u_n)$, the measurement with respect to \mathcal{B} is the tuple

$$\mathbf{M} = (M_1, M_2, \dots, M_n) = (u_1u_1^*, u_2u_2^*, \dots, u_nu_n^*).$$

More generally, we may use sets other than $[n]$, for instance the vertex set of a graph, to index the elements of a measurement. If we perform the above measurement \mathbf{M} on a quantum state $\psi \in \mathbb{C}^n$, we will obtain an outcome $i \in [n]$, and the new state of our system will be u_i . The state of a quantum system after performing a measurement is typically referred to as a *post-measurement state*. The outcome of the measurement is in general not determined, and will be i with probability $\psi^* M_i \psi = |\psi^* u_i|^2$. An important consequence of this is that outcome i occurs with zero probability if ψ is orthogonal to u_i . Also note that multiplying ψ by a complex number of modulus one does not have any effect on the probability of an outcome of a measurement.

As an example, suppose that our quantum system is in state

$$\psi = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

If we perform the measurement $(e_1 e_1^*, e_2 e_2^*)$, then we will obtain outcome 1 with probability

$$\left| \begin{pmatrix} 1 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right|^2 = 1,$$

and our post-measurement state will be e_1 . However, if we were to measure ψ with respect to the basis

$$\left(\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix} \right),$$

then we would obtain outcome 1 with probability

$$\left| \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \right|^2 = \frac{1}{2},$$

and the post-measurement state would be $\frac{1}{\sqrt{2}}(1, 1)^T$. Similarly, we would obtain outcome 2 with probability

$$\left| \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 \end{pmatrix} \begin{pmatrix} 1 \\ -1 \end{pmatrix} \right|^2 = \frac{1}{2},$$

and the post-measurement state would be $\frac{1}{\sqrt{2}}(1, -1)^T$.

5.2 General States and Measurements

In the previous section we considered the example of a die which was rolled but whose outcome was not revealed. In order to describe the state of the die we used a convex

combination of the vectors describing the deterministic states. Now let us consider a quantum analog of this scenario. Suppose Alice prepares a quantum state for Bob in a way such that with probability $1/2$ she prepares the state $(1, 0)^T$, and with probability $1/2$ she prepares the state $\frac{1}{\sqrt{2}}(1, 1)^T$. If Bob knows these probabilities but not the exact state Alice has prepared, how can he describe this state mathematically? Taking a convex combination no longer makes sense because the result is not guaranteed to be a unit vector, i.e. a valid quantum state. Such a state is called *mixed*, and is described mathematically by a *density matrix*. A density matrix is simply a positive semidefinite matrix which has trace equal to one. The density matrix corresponding to a pure state ψ is the matrix $\psi\psi^*$. Note that if ψ was a valid quantum pure state, then $\psi\psi^*$ will have trace one. Using density matrices, we can describe the state prepared for Bob by

$$\frac{1}{2} \left(\begin{pmatrix} 1 \\ 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \end{pmatrix} \right) + \frac{1}{2} \left(\frac{1}{2} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \end{pmatrix} \right) = \frac{1}{4} \begin{pmatrix} 3 & 1 \\ 1 & 1 \end{pmatrix}.$$

Note that this is positive semidefinite with trace one and thus a valid density matrix as expected. More generally, if pure state ψ_i is prepared for Bob with probability p_i for $i \in [k]$, then Bob can describe the state he has with the density matrix

$$\sum_{i=1}^k p_i \psi_i \psi_i^*.$$

Since these more general quantum states are described by a different type of object, we need to know how to mathematically describe the action of quantum operations and measurements on such states. We will mainly focus on measurements, but we note that if we apply a quantum operation given by a unitary U to a general quantum state ρ , then the resulting state is $U\rho U^*$.

If ρ is a density matrix representing the state of some quantum system, and we perform measurement $\mathbf{M} = (u_1 u_1^*, u_2 u_2^*, \dots, u_n u_n^*)$ on this system, then we will obtain outcome i with probability

$$\text{Tr}(u_i u_i^* \rho),$$

and the post-measurement state will be $u_i u_i^*$, which can be equivalently described as simply u_i since it is a pure state. Note that if ρ corresponds to a pure state ψ , then $\rho = \psi\psi^*$ and thus

$$\text{Tr}(u_i u_i^* \rho) = \text{Tr}(u_i u_i^* \psi\psi^*) = \text{Tr}(\psi^* u_i u_i^* \psi) = |\psi^* u_i|^2,$$

which coincides with the expression for the probability of obtaining outcome i given in the previous section.

Recalling the state

$$\frac{1}{4} \begin{pmatrix} 3 & 1 \\ 1 & 1 \end{pmatrix}$$

from above, if we were to measure this state with $(e_1 e_1^*, e_2 e_2^*)$, then we would obtain outcome 1 with probability

$$\text{Tr} \left(\frac{1}{4} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 3 & 1 \\ 1 & 1 \end{pmatrix} \right) = \text{Tr} \left(\frac{1}{4} \begin{pmatrix} 3 & 1 \\ 0 & 0 \end{pmatrix} \right) = \frac{3}{4},$$

and our post-measurement state would be

$$\begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

Similarly, we would obtain outcome 2 with probability

$$\text{Tr} \left(\frac{1}{4} \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 3 & 1 \\ 1 & 1 \end{pmatrix} \right) = \text{Tr} \left(\frac{1}{4} \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix} \right) = \frac{1}{4},$$

and our post-measurement state would be

$$\begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}.$$

Now recall that the density matrix

$$\frac{1}{4} \begin{pmatrix} 3 & 1 \\ 1 & 1 \end{pmatrix}$$

corresponded to a system which was in state e_1 with probability $1/2$ and in state $\frac{1}{\sqrt{2}}(1, 1)^T$ with probability $1/2$. If it is in state e_1 , then measuring with measurement $(e_1 e_1^*, e_2 e_2^*)$ yields outcome 1 with probability 1 as noted previously. Similarly, if it is in state $\frac{1}{\sqrt{2}}(1, 1)^T$, then performing this measurement results in outcome 1 with probability

$$\left| \frac{1}{\sqrt{2}} (1 \ 1) \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right|^2 = \frac{1}{2}.$$

So the total probability of obtaining outcome 1 is

$$\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{2} = \frac{3}{4}.$$

This coincides with the computation of this probability using the density matrix as one would want. This will always hold in general, thus validating the use of density matrices to describe a general quantum state.

In general, a quantum measurement does not have to be a full basis measurement which we discussed in the previous section. A more general type of measurement is known as a *projective measurement*. A tuple

$$\mathbf{M} = (M_1, M_2, \dots, M_n)$$

is a projective measurement if M_i is a projector for all $i \in [n]$, and

$$\sum_{i=1}^n M_i = I,$$

where I is the identity matrix. Here, we say that M_i is a *projector* if $M_i^2 = M_i$ and M_i is Hermitian, which is elsewhere known as an *orthogonal* projector. We use the former terminology because in Chapter 6 we refer to two projectors being *orthogonal to each other*, and we believe that using orthogonal in both ways could cause confusion.

Performing a projective measurement $\mathbf{M} = (M_1, M_2, \dots, M_n)$ on a quantum state ρ is similar to performing a full basis measurement. The outcome i will occur with probability $\text{Tr}(M_i\rho)$, and the post-measurement state will be

$$\frac{1}{\text{Tr}(M_i\rho)} M_i\rho M_i.$$

The scalar in front is just a normalization factor to ensure that the resulting matrix has trace one. Note that since the M_i are projectors, they are positive semidefinite and thus $\text{Tr}(M_i\rho)$ is nonnegative. Furthermore, since the M_i sum to identity, the probabilities $\text{Tr}(M_i\rho)$ sum to one as required. Also note that if $\rho = \psi\psi^*$ for some pure state ψ , then

$$M_i\rho M_i = M_i\psi\psi^* M_i = (M_i\psi)(M_i\psi)^*,$$

which (after normalization) is the density matrix corresponding to the normalized projection of ψ onto the column space of M_i . Thus projective measurements take pure states to pure states.

Projective measurements will be important to the work of Chapter 6, but this is not the most general quantum measurement. The most general quantum measurement is a tuple

$$\mathbf{M} = (M_1, \dots, M_n)$$

where the M_i are any matrices satisfying

$$\sum_{i=1}^n M_i^* M_i = I.$$

Note that projective measurements meet this condition since if M_i is a projector, then $M_i^* M_i = M_i M_i = M_i$. Measuring a quantum state ρ with a general measurement \mathbf{M} results in outcome i with probability

$$\text{Tr}(M_i \rho M_i^*) = \text{Tr}(M_i^* M_i \rho),$$

and the post-measurement state will be

$$\frac{1}{\text{Tr}(M_i \rho M_i^*)} M_i \rho M_i^*.$$

Note that the outcome probabilities for measurement \mathbf{M} only depend on the matrices $H_i = M_i^* M_i$ and not on the M_i . The tuple (H_1, \dots, H_n) is referred to as the *positive-operator valued measure* (POVM), corresponding to the measurement \mathbf{M} . The matrices H_i for $i \in [n]$ are positive semidefinite and sum to identity by the definition of measurement. Conversely, any tuple of positive semidefinite matrices which sum to identity is the POVM corresponding to some measurement. However, a POVM does not completely specify a measurement, as there are multiple measurements which correspond to the same POVM. If we are only concerned with the outcome probabilities, as we will be in Section 6.5, then it suffices to specify the POVM, rather than the measurement.

As an example of performing a general measurement, consider the state

$$\rho = \frac{1}{2} \psi_1 \psi_1^* + \frac{1}{2} \psi_2 \psi_2^* = \begin{pmatrix} \frac{1}{4} & 0 \\ 0 & \frac{3}{4} \end{pmatrix},$$

where

$$\psi_1 = \begin{pmatrix} -1/2 \\ \sqrt{3}/2 \end{pmatrix}, \quad \psi_2 = \begin{pmatrix} -1/2 \\ -\sqrt{3}/2 \end{pmatrix}.$$

Also, let $\mathbf{M} = (M_1, M_2, M_3)$ be the measurement where

$$M_1 = \sqrt{\frac{2}{3}} \psi_1 \psi_1^*, \quad M_2 = \sqrt{\frac{2}{3}} \psi_2 \psi_2^*, \quad M_3 = \sqrt{\frac{2}{3}} e_1 e_1^*.$$

Note that

$$M_1^* M_1 = \frac{2}{3} \begin{pmatrix} \frac{1}{4} & \frac{-\sqrt{3}}{4} \\ \frac{-\sqrt{3}}{4} & \frac{3}{4} \end{pmatrix}, \quad M_2^* M_2 = \frac{2}{3} \begin{pmatrix} \frac{1}{4} & \frac{\sqrt{3}}{4} \\ \frac{\sqrt{3}}{4} & \frac{3}{4} \end{pmatrix}, \quad M_3^* M_3 = \frac{2}{3} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}.$$

If we perform measurement \mathbf{M} on state ρ , then we will obtain outcome 1 with probability

$$\text{Tr} \left(\frac{2}{3} \begin{pmatrix} \frac{1}{4} & \frac{-\sqrt{3}}{4} \\ \frac{-\sqrt{3}}{4} & \frac{3}{4} \end{pmatrix} \begin{pmatrix} \frac{1}{4} & 0 \\ 0 & \frac{3}{4} \end{pmatrix} \right) = \frac{5}{12},$$

and the post-measurement state will be

$$\frac{12}{5} \cdot \frac{2}{3} \begin{pmatrix} \frac{1}{4} & \frac{-\sqrt{3}}{4} \\ \frac{-\sqrt{3}}{4} & \frac{3}{4} \end{pmatrix} \begin{pmatrix} \frac{1}{4} & 0 \\ 0 & \frac{3}{4} \end{pmatrix} \begin{pmatrix} \frac{1}{4} & \frac{-\sqrt{3}}{4} \\ \frac{-\sqrt{3}}{4} & \frac{3}{4} \end{pmatrix} = \frac{1}{4} \begin{pmatrix} 1 & -\sqrt{3} \\ -\sqrt{3} & 3 \end{pmatrix} = \psi_1 \psi_1^*.$$

We will obtain outcome 2 with probability

$$\text{Tr} \left(\frac{2}{3} \begin{pmatrix} \frac{1}{4} & \frac{\sqrt{3}}{4} \\ \frac{\sqrt{3}}{4} & \frac{3}{4} \end{pmatrix} \begin{pmatrix} \frac{1}{4} & 0 \\ 0 & \frac{3}{4} \end{pmatrix} \right) = \frac{5}{12},$$

and the post-measurement state will be

$$\frac{12}{5} \cdot \frac{2}{3} \begin{pmatrix} \frac{1}{4} & \frac{\sqrt{3}}{4} \\ \frac{\sqrt{3}}{4} & \frac{3}{4} \end{pmatrix} \begin{pmatrix} \frac{1}{4} & 0 \\ 0 & \frac{3}{4} \end{pmatrix} \begin{pmatrix} \frac{1}{4} & \frac{\sqrt{3}}{4} \\ \frac{\sqrt{3}}{4} & \frac{3}{4} \end{pmatrix} = \frac{1}{4} \begin{pmatrix} 1 & \sqrt{3} \\ \sqrt{3} & 3 \end{pmatrix} = \psi_2 \psi_2^*.$$

Finally, we will obtain outcome 3 with probability

$$\text{Tr} \left(\frac{2}{3} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \frac{1}{4} & 0 \\ 0 & \frac{3}{4} \end{pmatrix} \right) = \frac{1}{6},$$

and the post-measurement state will be

$$6 \cdot \frac{2}{3} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \frac{1}{4} & 0 \\ 0 & \frac{3}{4} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = e_1 e_1^*.$$

5.3 Shared Systems

Suppose that now we have two dice, and we want to describe the state of both of them as a single system. If the first die has face i up with probability p_i and the second die has face j up with probability q_j , then, assuming the dice are independent of each other, the probability that the first die has face i up and the second has face j up is $p_i q_j$. If the vectors describing states of each of the two dice are p and q respectively, then the tensor product $p \otimes q$ has the values $p_i q_j$ as its entries. Thus we can use the tensor product (a.k.a. Kronecker

product) of vectors describing the states of two classical systems to describe the state of both of them as a single system.

The same method works for quantum systems: if Alice has a quantum system in pure state $\psi_A \in \mathbb{C}^{d_A}$ and Bob has a system whose state is $\psi_B \in \mathbb{C}^{d_B}$, then the state of their combined system can be described as $\psi_A \otimes \psi_B \in \mathbb{C}^{d_A} \otimes \mathbb{C}^{d_B}$. Note, however, that not all unit vectors in $\mathbb{C}^{d_A} \otimes \mathbb{C}^{d_B}$ can be written as the tensor product of a unit vector in \mathbb{C}^{d_A} and a unit vector in \mathbb{C}^{d_B} . But by the principle of superposition, it must be possible for Alice and Bob's joint system to be in these non-product states. If $\psi \in \mathbb{C}^{d_A} \otimes \mathbb{C}^{d_B}$ is such a unit vector, then we say that ψ is an *entangled state*. An important example of such state is

$$\Phi_d = \frac{1}{\sqrt{d}} \sum_{i=1}^d e_i \otimes e_i,$$

where we have assumed that $d_A = d_B = d$. This state is known as the *canonical maximally entangled state* and is often used in quantum information to accomplish tasks which are not possible classically. The canonical maximally entangled state will be of importance to us in Section 6.5.

Now suppose that Alice and Bob's combined system is in some general state $\rho \in \mathbb{C}^{d_A \times d_A} \otimes \mathbb{C}^{d_B \times d_B}$, where $\mathbb{C}^{d \times d}$ is the space of $d \times d$ complex matrices. We would like to be able to describe the "local" state of just Alice's system. In order to do this, we need the following definition.

Definition. Given $d_A, d_B \in \mathbb{N}$, define the *partial trace over A*, denoted Tr_A , to be the unique linear map from $\mathbb{C}^{d_A \times d_A} \otimes \mathbb{C}^{d_B \times d_B}$ to $\mathbb{C}^{d_B \times d_B}$ such that

$$\text{Tr}_A(M_A \otimes M_B) = \text{Tr}(M_A)M_B$$

for all $M_A \in \mathbb{C}^{d_A \times d_A}$, $M_B \in \mathbb{C}^{d_B \times d_B}$. The partial trace over B is defined analogously.

Using this definition, we can describe the local state of Alice's system in terms of the state of Alice and Bob's shared system. If the state of their combined system is $\rho \in \mathbb{C}^{d_A \times d_A} \otimes \mathbb{C}^{d_B \times d_B}$, then Alice's part of the system is in state

$$\rho_A := \text{Tr}_B(\rho),$$

and Bob's part is in state

$$\rho_B := \text{Tr}_A(\rho).$$

The states ρ_A and ρ_B are typically referred to as the *reduced states* of ρ . It is important to note that though ρ_A and ρ_B fully describe the local states of Alice's and Bob's respective

parts of their shared system, they do not give a full description of the system as a whole. In particular, it is not always the case that

$$\rho = \rho_A \otimes \rho_B.$$

This is in some sense the essence of entanglement, in that it is possible for Alice and Bob's shared system to exhibit properties that are not merely combinations of local properties.

Suppose that Alice performs a quantum operation given by unitary U on her system, how does this affect her and Bob's combined system? Since she performed an operation only on her system, Bob's system should not be affected, so we would want the reduced states after Alice's operation to be $U\rho_A U^*$ and ρ_B respectively. It is not hard to check that

$$\text{Tr}_A((U \otimes I)\rho(U^* \otimes I)) = \text{Tr}_A(\rho)$$

and

$$\text{Tr}_B((U \otimes I)\rho(U^* \otimes I)) = U \text{Tr}_B(\rho) U^*.$$

Thus Alice performing U on her system corresponds to $U \otimes I$ being performed on the combined system. More generally, if Alice performs U and Bob performs V , then this corresponds to $U \otimes V$ being performed on the combined system.

Now suppose that Alice performs a measurement $\mathbf{M} = (M_1, \dots, M_n)$ on her system. We want to show that this is equivalent to performing the measurement $\mathbf{M}' = (M_1 \otimes I, \dots, M_n \otimes I)$ on the combined system. To do this we must show that

$$\text{Tr}_B((M_i \otimes I)\rho(M_i^* \otimes I)) = M_i(\text{Tr}_B \rho)M_i^*. \quad (5.1)$$

The above are the unnormalized post-measurement states on Alice's system. If they are equal, then the probability of obtaining outcome i in the two measurements above are equal since

$$\text{Tr}((M_i \otimes I)\rho(M_i^* \otimes I)) = \text{Tr}(\text{Tr}_B((M_i \otimes I)\rho(M_i^* \otimes I))).$$

To show that (5.1) holds, we first write ρ as a sum of tensor products:

$$\rho = \sum_{j=1}^k E_j \otimes F_j,$$

where the E_i and F_j are not necessarily positive semidefinite. Now we have that

$$\begin{aligned}
\mathrm{Tr}_B((M_i \otimes I)\rho(M_i^* \otimes I)) &= \sum_{j=1}^k \mathrm{Tr}_B((M_i E_j M_i^*) \otimes F_j) \\
&= \sum_{j=1}^k M_i E_j M_i^* \mathrm{Tr}(F_j) \\
&= M_i \left(\sum_{j=1}^k E_j \mathrm{Tr}(F_j) \right) M_i^* \\
&= M_i \mathrm{Tr}_B(\rho) M_i^*.
\end{aligned}$$

The above shows that performing \mathbf{M} on Alice's system or \mathbf{M}' on the combined system are equivalent to Alice, but what about Bob? Clearly, Alice performing \mathbf{M} on her system should have no effect on the local state of Bob's system, but at first glance it appears that performing \mathbf{M}' on the combined system will alter the state of Bob's system, since

$$\mathrm{Tr}_A((M_i \otimes I)\rho(M_i^* \otimes I)) = \sum_{j=1}^k \mathrm{Tr}(M_i E_j M_i^*) F_j$$

is not guaranteed to be equal to

$$\mathrm{Tr}_A(\rho) = \sum_{j=1}^k \mathrm{Tr}(E_j) F_j.$$

However, if Alice simply performed a measurement on her system, then Bob would not know the outcome. If we assume that \mathbf{M}' is performed on the combined system, but Bob is not told the outcome, then from Bob's perspective the state of his system is (probabilities

and normalization factors canceling):

$$\begin{aligned}
\mathrm{Tr}_A \left(\sum_{i=1}^n (M_i \otimes I) \rho (M_i^* \otimes I) \right) &= \sum_{j=1}^k \sum_{i=1}^n \mathrm{Tr}_A ((M_i E_j M_i^*) \otimes F_j) \\
&= \sum_{j=1}^k \sum_{i=1}^n \mathrm{Tr} (M_i E_j M_i^*) F_j \\
&= \sum_{j=1}^k F_j \left(\sum_{i=1}^n \mathrm{Tr} (M_i^* M_i E_j) \right) \\
&= \sum_{j=1}^k F_j \mathrm{Tr} \left(\sum_{i=1}^n M_i^* M_i E_j \right) \\
&= \sum_{j=1}^k F_j \mathrm{Tr} (E_j) \\
&= \mathrm{Tr}_A (\rho).
\end{aligned}$$

On the other hand, if Alice informed Bob of her measurement outcome, then the state of Bob's system would change. This may seem counterintuitive, but it is possible to construct an analogous classical scenario.

Suppose that Charlie flips a fair coin and then places coins in some box A and box B such that they have the same side facing up as the flipped coin. Now suppose that Charlie gives Alice and Bob boxes A and B respectively, and tells them about the correlation between their coins but not which side they are on. Then from Alice and Bob's perspectives, the state of each of their systems is $(1/2, 1/2)^T$. If Alice opens her box then she will know which side her coin is on and therefore what side Bob's coin is on. However, from Bob's perspective his coin is still in state $(1/2, 1/2)^T$. But once Alice tells him what side her coin was on, the state of his system will change from his perspective.

If both Alice and Bob perform measurements $\mathbf{M} = (M_1, \dots, M_k)$ and $\mathbf{N} = (N_1, \dots, N_\ell)$ on their respective parts of a shared system, then this corresponds to performing the measurement $(M_i \otimes N_j)_{i \in [k], j \in [\ell]}$ on the combined system.

We will end with an example of performing measurements on a shared system. Suppose that Alice and Bob share a system which is in state

$$\varphi = \frac{1}{\sqrt{d}} \sum_{i=1}^d u_i \otimes \bar{u}_i \in \mathbb{C}^d \otimes \mathbb{C}^d$$

where $\{u_1, \dots, u_d\}$ is an orthonormal basis for \mathbb{C}^d , and \bar{u}_i denotes the complex conjugate of u_i . Suppose that Alice performs measurement $(u_1 u_1^*, \dots, u_d u_d^*)$ and Bob performs measurement $(\bar{u}_1 \bar{u}_1^*, \dots, \bar{u}_d \bar{u}_d^*)$. The probability of Alice obtaining outcome k and Bob obtaining outcome ℓ is

$$\begin{aligned} \varphi^*(u_k u_k^* \otimes \bar{u}_\ell \bar{u}_\ell^*) \varphi &= \frac{1}{d} \left(\sum_{i=1}^d u_i^* \otimes \bar{u}_i^* \right) (u_k u_k^* \otimes \bar{u}_\ell \bar{u}_\ell^*) \left(\sum_{j=1}^d u_j \otimes \bar{u}_j \right) \\ &= \begin{cases} \frac{1}{d} & \text{if } k = \ell \\ 0 & \text{if } k \neq \ell \end{cases} . \end{aligned}$$

Therefore, Alice and Bob always obtain the same outcome. In fact, it turns out that this consequence does not rely on the choice of orthonormal basis used in the measurements for Alice and Bob. As long as Alice and Bob's bases are complex conjugates of each other, they will always obtain the same outcome. We will not prove this in detail, but it follows from the fact that the state φ in the example above is in fact equal to

$$\Phi_d = \frac{1}{\sqrt{d}} \sum_{i=1}^d e_i \otimes e_i,$$

for any choice of orthonormal basis $\{u_1, \dots, u_d\}$.

Chapter 6

Quantum Homomorphisms

In this chapter we will introduce the notion of quantum homomorphisms and give several results concerning them. The study of quantum homomorphisms was originally motivated by concerns in the field of quantum information. Initially only defining quantum colorings, the definition was an operational one, meaning that a graph X was said to have a quantum n -coloring if there existed a “quantum strategy” for accomplishing a certain task. This task was winning a game in which two players who cannot communicate attempt to convince a referee that they have an n -coloring of X . A “quantum strategy” refers to one in which the players share some state and can perform measurements on their individual parts. Players using classical strategies, which will be formally defined in Section 6.4, can win this game if and only if there exists an n -coloring of X , thus inspiring the term “quantum coloring” and eventually the quantum chromatic number, which is the smallest n for which a graph admits a quantum n -coloring. In [8], Cameron et al. showed that the existence of a winning quantum strategy to the coloring game is equivalent to the existence of a set of projectors satisfying certain orthogonalities determined by the graph in question. Later in [42], Mančinska and Roberson generalized the coloring game to a homomorphism game, and noted that the result of [8] still held for this game.

Here we will approach quantum homomorphisms from a more mathematical approach. One of the things we mean by this is that we will define quantum homomorphisms in terms of the existence of the projectors mentioned above, and then show this is equivalent to the existence of a quantum strategy for the homomorphism game, rather than the other way around. Obviously, it does not matter which of these two ways they are defined, but we will also more generally focus less on the game based definition, and all of our proofs will use the definition in terms of projectors.

One of the underlying goals of this chapter will be to compare the behavior of quantum homomorphisms to that of homomorphisms. To this end we will show that certain theorems about homomorphisms can be adapted in a natural way into theorems about quantum homomorphisms. We will also be able to use the notion of quantum homomorphisms to define various quantum analogs of graph parameters which have definitions in terms of homomorphisms. This extends the idea of the quantum chromatic number to include things such as quantum clique number or quantum odd girth. Comparing a parameter to its quantum analog gives us another method for comparing homomorphisms with quantum homomorphisms. As with homomorphisms, one can use quantum homomorphisms to construct a partial order. The quantum homomorphism order serves as another tool for comparing homomorphisms to quantum homomorphisms, though we only have a few results in this direction.

The rest of this chapter is outlined as follows. In Section 6.1, we introduce some notation and some basic linear algebra tools that we will use throughout the remaining sections. Next we give the definition of quantum homomorphisms in Section 6.2, and we discuss a few simple properties. Section 6.3 introduces the notion of the measurement graph of a graph, which can be used to formulate quantum homomorphisms in terms of homomorphisms. In Section 6.4, we formally define the homomorphism game described above and explain why the existence of classical strategies for this game is equivalent to the existence of homomorphisms. Next we introduce a general quantum strategy for the homomorphism game in Section 6.5, and we go through a proof of the Cameron et al. result for the homomorphism game.

Section 6.6 explores some of the basic properties of quantum homomorphisms through three lemmas. Of particular importance is that quantum homomorphisms are transitive. In Section 6.7 we introduce vector colorings, as well as both strict and rigid vector colorings, along with their corresponding chromatic numbers. We note that all three of these parameters are related to the Lovász ϑ function, and that the strict vector chromatic number is in fact equal to Lovász ϑ of the complement. We go on to prove that these three variants of chromatic number are quantum homomorphism monotone. Next we introduce the notion of the homomorphic product in Section 6.8, which encodes the orthogonality constraints of a quantum homomorphism into adjacencies in a graph. We also define projective packings in this section, and show that there exists a quantum homomorphism from X to Y if and only if there exists a projective packing of their homomorphic product of a certain value. We also prove that the projective packing number of the complement is quantum homomorphism monotone and upper bounded by the vector chromatic number.

In Section 6.9 we use quantum homomorphisms to define some quantum analogs of graph parameters such as odd girth and independence number. We briefly discuss why it is

clear that these quantum parameters are quantum homomorphism monotone. Section 6.10 focuses on the quantum chromatic number, showing that this parameter is lower bounded by the rigid vector chromatic number. We also investigate a class of graphs that has been of particular importance in the study of quantum colorings, showing that the known upper bound on their quantum chromatic number holds with equality. We study the quantum independence number in Section 6.11, showing that it is upper bounded by the projective packing number of Section 6.8. We also prove a quantum analog of the fact that X has a homomorphism to Y if and only if the homomorphic product of X with Y has an independent set of size $|V(X)|$. This allows us to construct a graph with larger quantum independence number than quantum independence number from any graphs X and Y such that X admits a quantum homomorphism but no homomorphism to Y . In Section 6.12 we discuss a counterexample to a possible quantum analog of the no-homomorphism lemma.

We define projective representations in Section 6.13, along with their associated parameter, the projective rank. Projective representations and projective rank are closely related to projective packings and the projective packing number. We make this relationship explicit by showing that the projective rank of a graph is always at least its number of vertices divided by its projective packing number. We furthermore show that equality holds in the case where the graph is vertex transitive, analogously to the relationship between fractional chromatic number and independence number. We also note in this section that the projective rank lower bounds the quantum chromatic number. Section 6.14 continues the study of projective rank, showing that it is quantum homomorphism monotone and lower bounded by the rigid vector chromatic number. In Section 6.15 we consider the projective rank of some special classes of graphs such as Kneser graphs and odd cycles. We compute the projective rank exactly for both of these classes of graphs which allows us to establish that the quantum odd girth of any Kneser graph is equal to its odd girth.

In the next two sections we focus on the order theoretic properties of quantum homomorphisms. We begin in Section 6.16 by showing that the categorical product and disjoint union of graph act as the meet and join in the quantum homomorphism order, as they did for the homomorphism order. Section 6.17 formally defines the quantum homomorphism order and focuses on comparing it to the homomorphism order. In particular we note that there exists a lattice homomorphism from the homomorphism order to the quantum homomorphism order, and that the quantum homomorphism order is isomorphic to the suborder of the homomorphism order of infinite graphs induced by the measurement graphs. We end the chapter with Section 6.18 which summarizes our results, and with Section 6.19 which outlines some open questions of interest to us. Some of the results in this chapter are from a joint work with Laura Mančinska [42].

6.1 Some Linear Algebra

In this section we present some basic linear algebra facts that will be frequently used in the study of quantum homomorphisms. The reader will likely be familiar with some or all of these, so this mainly serves as a notice to the reader to keep these facts in the front of their mind, since we will be using them often and without announcement. The results discussed in this section are previously known and not the work of the author.

We say that two projectors A and B are *orthogonal* (to each other) if $AB = 0$. Note that this is equivalent to $\text{Tr}(AB) = 0$ which is equivalent to the column spaces of A and B being orthogonal. This remains true even if we only require that A and B be positive semidefinite. In the next section we will define quantum homomorphisms in terms of projectors, and thus the following will prove useful in our study of them:

If $P_1, P_2, \dots, P_k \in \mathbb{C}^{d \times d}$ are pairwise orthogonal projectors, then

- $\sum_{i=1}^k P_i$ is a projector;
- $\text{rk}\left(\sum_{i=1}^k P_i\right) = \sum_{i=1}^k \text{rk}(P_i)$;
- $\sum_{i=1}^k \text{rk}(P_i) = d$ if and only if $\sum_{i=1}^k P_i = I_d$.

Here I_d denotes the $d \times d$ identity, though we will often drop the subscript when the dimension is implicit or not of particular relevance. We will also use the fact that if A and B are projectors, then the tensor product $A \otimes B$ is a projector with rank $\text{rk}(A) \text{rk}(B)$.

A less trivial result that we will need for the proof of Theorem 6.5.1 is the following:

Lemma 6.1.1. *Let $A, B \in \mathbb{C}^{d \times d}$ be positive semidefinite matrices such that $I - A$ is also positive semidefinite. Then $\text{Tr}(AB) = \text{Tr}(B)$ implies that the column space of B is contained in the 1-eigenspace of A .*

Proof. Using spectral decomposition, we have

$$A = \sum_{i=1}^d \lambda_i v_i v_i^*,$$

where v_i is an eigenvector of A for eigenvalue λ_i and $\{v_i : i \in [d]\}$ is an orthonormal basis for \mathbb{C}^d . Note that since $I - A$ is positive semidefinite, we have that $\lambda_i \leq 1$ for all $i \in [d]$.

Using this, we see that

$$\mathrm{Tr}(AB) = \sum_{i=1}^d \lambda_i \mathrm{Tr}(v_i v_i^* B) = \sum_{i=1}^d \lambda_i v_i^* B v_i \leq \sum_{i=1}^d v_i^* B v_i = \mathrm{Tr}(B).$$

Since $\mathrm{Tr}(AB) = \mathrm{Tr}(B)$, equality must hold in the inequality above. This implies that $\lambda_i = 1$ whenever $v_i^* B v_i \neq 0$. Let $N = \{i \in [d] : v_i^* B v_i = 0\}$. Since B is positive semidefinite, $N = \{i \in [d] : B v_i = 0\}$ and thus $\mathrm{span}(\{v_i : i \in N\})$ is a subspace of $\mathrm{null}(B)$. Therefore,

$$\mathrm{span}(\{v_i : i \notin N\}) = \mathrm{span}(\{v_i : i \in N\})^\perp \supseteq \mathrm{null}(B)^\perp = \mathrm{col}(B),$$

where $^\perp$ denotes the orthogonal complement. By the above, $i \notin N$ implies that $\lambda_i = 1$ and thus we have proven the lemma. \square

The last tool we will need is known as the *Schmidt decomposition*, which is just a particular way of writing a vector in a tensor product space.

Theorem 6.1.2. *Suppose that $v \in \mathbb{C}^m \otimes \mathbb{C}^n$ and $m \leq n$. Then there exist orthonormal sets $\{\alpha_1, \dots, \alpha_m\} \subseteq \mathbb{C}^m$ and $\{\beta_1, \dots, \beta_m\} \subseteq \mathbb{C}^n$ such that*

$$v = \sum_{i=1}^m \sqrt{\lambda_i} \alpha_i \otimes \beta_i,$$

where the λ_i are nonnegative and uniquely determined, as a set, by v . \square

We will not give a proof of this here, but we note that it is actually equivalent to singular value decomposition.

6.2 Definition of Quantum Homomorphism

Though quantum homomorphisms were originally defined by Mančinska and the author via a game played between two players and a referee [42], here we will present a more “mathematical” definition. This definition will, hopefully, be more palatable to mathematicians, and not too offensive to physicists. Of course, in Section 6.5 we will see that the two definitions are in fact equivalent.

Definition. For graphs X and Y , we say that X has a *quantum homomorphism* to Y , and write $X \xrightarrow{q} Y$, if there exists a $d \in \mathbb{N}$ and projectors $E_{xy} \in \mathbb{C}^{d \times d}$ for $x \in V(X)$, $y \in V(Y)$,

such that

$$\sum_{y \in V(Y)} E_{xy} = I, \quad \forall x \in V(X); \quad (6.1)$$

$$E_{xy}E_{x'y'} = 0, \quad \text{if } (x = x' \ \& \ y \neq y') \text{ or } (x \sim x' \ \& \ y \not\sim y'). \quad (6.2)$$

We will say that the “ E_{xy} are projectors that give a quantum homomorphism from X to Y ”.

Note that $y \not\sim y'$ includes the case of $y = y'$. Also note that Condition 6.1 actually implies that $E_{xy}E_{x'y'} = 0$ for $x = x'$ and $y \neq y'$, and so this part of Condition 6.2 is redundant. However, we wanted to point out all of the necessary orthogonalities explicitly. For a fixed $x \in V(X)$, the matrices E_{xy} for $y \in V(Y)$ form a projective measurement whose outcomes are indexed by the vertices of Y . This is not merely a coincidence, and we will see the reason behind this in Section 6.5.

There is nothing stopping us from applying the definition of quantum homomorphism to graphs X and Y which are potentially infinite. Keep in mind though, that if Y is infinite, then for a given $x \in V(X)$ there will be at most finitely many nonzero E_{xy} . This is because the rank of a matrix is integer valued, and so it is not possible for more than d nonzero projectors to sum up to the $d \times d$ identity.

The above definition may seem a bit arbitrary, but as we mentioned above, we will see the physical motivation for it in the next section. Of course, we would like homomorphisms and quantum homomorphisms to be related in some way, and the following lemma, which is joint work with Laura Mančinska, shows that this is in fact the case.

Lemma 6.2.1. *If $X \rightarrow Y$, then $X \xrightarrow{q} Y$.*

Proof. Let φ be a homomorphism from X to Y . For $x \in V(X)$ and $y \in V(Y)$, define E_{xy} to be the identity matrix if $\varphi(x) = y$, and be the zero matrix otherwise. We claim that the E_{xy} give a quantum homomorphism from X to Y . For $x \in V(X)$, we see that

$$\sum_{y \in V(Y)} E_{xy} = I$$

since φ maps each vertex of X to exactly one vertex of Y . Furthermore, if $x \sim x'$ and $y \not\sim y'$, then it must be the case that at least one of $\varphi(x) = y$ and $\varphi(x') = y'$ is false. Therefore, in this case we have that

$$E_{xy}E_{x'y'} = 0.$$

These are all of the conditions for the E_{xy} to give a quantum homomorphism from X to Y , and so we are done. \square

The dimension of the matrices used in the above proof could be anything, even one. In fact, it is not hard to see that X has a homomorphism to Y if and only if X has a quantum homomorphism to Y using projectors in $\mathbb{C}^{1 \times 1}$.

In Section 6.7, we will investigate the relationship between quantum homomorphisms and some graph parameters which are defined via assignments of real vectors to vertices. The work in that section will require us to create an assignment of real vectors for a graph X from an assignment of real vectors for a graph Y assuming that $X \xrightarrow{q} Y$. In order to ensure that the vectors we assign to X are real, we need the following lemma which shows that real matrices suffice for quantum homomorphisms. This lemma is joint work with Laura Mančinska.

Lemma 6.2.2. *If $X \xrightarrow{q} Y$, then there exist real projectors E_{xy} for $x \in V(X)$, $y \in V(Y)$ that give a quantum homomorphism from X to Y .*

Proof. Suppose that $X \xrightarrow{q} Y$, and let $F_{xy} \in \mathbb{C}^{d \times d}$ for $x \in V(X)$, $y \in V(Y)$ be projectors which give a quantum homomorphism from X to Y . Let $R : \mathbb{C}^{d \times d} \rightarrow \mathbb{R}^{2d \times 2d}$ be the map defined by

$$R(A) = \begin{pmatrix} \Re(A) & \Im(A) \\ -\Im(A) & \Re(A) \end{pmatrix}$$

where $\Re(A)$ and $\Im(A)$ are the real and imaginary parts of A respectively. It is routine to check that $R(A + B) = R(A) + R(B)$, $R(AB) = R(A)R(B)$, and that R takes Hermitian matrices to symmetric matrices. Since $R(I) = I$ and $R(0) = 0$, the matrices $E_{xy} = R(F_{xy})$ for $x \in V(X)$, $y \in V(Y)$ are real projectors that give a quantum homomorphism from X to Y . \square

As we saw with the no-homomorphism lemma, the symmetry of vertex transitive graphs can sometimes be exploited to obtain information about homomorphisms to these graphs. The following lemma shows that this idea works for quantum homomorphisms as well.

Lemma 6.2.3. *Let X and Y be graphs and further let Y be vertex transitive. If $X \xrightarrow{q} Y$, then there exist projectors E_{xy} for $x \in V(X)$, $y \in V(Y)$ which give a quantum homomorphism from X to Y and all have the same rank.*

Proof. Suppose that $X \xrightarrow{q} Y$. Then there exist $d \in \mathbb{N}$ and projectors $F_{xy} \in \mathbb{C}^{d \times d}$ which give a quantum homomorphism from X to Y . Now fix some ordering of the elements of $\text{Aut}(Y)$. For $\sigma \in \text{Aut}(Y)$, let M_σ denote the matrix whose rows and columns are indexed

by the elements of $\text{Aut}(Y)$, such that its (σ, σ) -entry is 1 and all other entries are 0. For $x \in V(X)$ and $y \in V(Y)$, define

$$E_{xy} = \sum_{\sigma \in \text{Aut}(Y)} F_{x\sigma(y)} \otimes M_\sigma$$

We will show that the E_{xy} give a quantum homomorphism from X to Y and all have the same rank. First we must check that they are indeed projectors. Since the F_{xy} and M_σ are projectors, each of the terms $F_{x\sigma(y)} \otimes M_\sigma$ is a projector. Since $M_\sigma M_{\sigma'} = 0$ for distinct $\sigma, \sigma' \in \text{Aut}(X)$, the terms in the sum above are pairwise orthogonal and therefore E_{xy} is a projector.

Now we must check that Condition 6.1 from the definition of quantum homomorphism holds for the E_{xy} . To see this, first note that, for a given $\sigma \in \text{Aut}(Y)$,

$$\sum_{y \in V(Y)} F_{x\sigma(y)} = \sum_{y \in V(Y)} F_{xy} = I,$$

since σ is a bijection on $V(Y)$. Now we have that

$$\begin{aligned} \sum_{y \in V(Y)} E_{xy} &= \sum_{y \in V(Y)} \sum_{\sigma \in \text{Aut}(Y)} (F_{x\sigma(y)} \otimes M_\sigma) \\ &= \sum_{\sigma \in \text{Aut}(Y)} \left(\left(\sum_{y \in V(Y)} F_{x\sigma(y)} \right) \otimes M_\sigma \right) \\ &= \sum_{\sigma \in \text{Aut}(Y)} (I \otimes M_\sigma) \\ &= I \otimes \sum_{\sigma \in \text{Aut}(Y)} M_\sigma \\ &= I \otimes I = I. \end{aligned}$$

To check Condition 6.2, let $x, x' \in V(X)$ and $y, y' \in V(Y)$ be such that $x \sim x'$ and $y \not\sim y'$.

Note that $\sigma(y) \not\sim \sigma(y')$, and therefore $F_{x\sigma(y)}F_{x'\sigma(y')} = 0$ for all $\sigma \in \text{Aut}(Y)$. We have that

$$\begin{aligned}
E_{xy}E_{x'y'} &= \left(\sum_{\sigma \in \text{Aut}(Y)} F_{x\sigma(y)} \otimes M_\sigma \right) \left(\sum_{\sigma' \in \text{Aut}(Y)} F_{x'\sigma'(y')} \otimes M_{\sigma'} \right) \\
&= \sum_{\sigma, \sigma' \in \text{Aut}(Y)} F_{x\sigma(y)} F_{x'\sigma'(y')} \otimes M_\sigma M_{\sigma'} \\
&= \sum_{\sigma \in \text{Aut}(Y)} F_{x\sigma(y)} F_{x'\sigma(y')} \otimes M_\sigma \\
&= 0.
\end{aligned}$$

Now all that is left to show is that the E_{xy} all have the same rank. For some $x \in V(X)$ and $y \in V(Y)$, we have that

$$\begin{aligned}
\text{rk}(E_{xy}) &= \sum_{\sigma \in \text{Aut}(Y)} \text{rk}(F_{x\sigma(y)} \otimes M_\sigma) \\
&= \sum_{\sigma \in \text{Aut}(Y)} \text{rk}(F_{x\sigma(y)}) \text{rk}(M_\sigma) \\
&= \sum_{\sigma \in \text{Aut}(Y)} \text{rk}(F_{x\sigma(y)}) \\
&= \sum_{\sigma \in \text{Aut}(Y)} \text{Tr}(F_{x\sigma(y)}) \\
&= \text{Tr} \left(\sum_{\sigma \in \text{Aut}(Y)} F_{x\sigma(y)} \right),
\end{aligned}$$

since the rank of a projector is equal to its trace. However, since Y is vertex transitive, the last sum above contains each $F_{xy'}$ for $y' \in V(Y)$ the same number of times. More specifically, it contains each of them exactly $m := |\text{Aut}(Y)|/|V(Y)|$ times. Therefore, we have that

$$\sum_{\sigma \in \text{Aut}(Y)} F_{x\sigma(y)} = mI.$$

Thus, $\text{rk}(E_{xy}) = md$ for all $x \in V(X)$ and $y \in V(Y)$. \square

Recalling that the sum of pairwise orthogonal projectors in $\mathbb{C}^{d \times d}$ is equal to the identity if and only if the sum of their ranks is equal to d , we see that in the case where Y is vertex transitive, Condition 6.1 can be replaced by the requirement that all of the E_{xy} have some

constant rank r such that $d = r|V(Y)|$. Note however, that in this case the $E_{xy}E_{x'y'} = 0$ for $x = x'$ and $y \neq y'$ condition is no longer redundant.

As an application of the above lemma, we give yet another proof of the clique-coclique bound:

Lemma 2.7.2. *If Y is a vertex transitive graph, then*

$$\alpha(Y)\omega(Y) \leq |V(Y)|.$$

Proof. Let $m = \omega(Y)$. Then $K_m \rightarrow Y$ and thus by Lemma 6.2.1 $K_m \xrightarrow{q} Y$. By Lemma 6.2.3, there exist rank r projectors $E_{iy} \in \mathbb{C}^{d \times d}$ for $i \in [m]$, $y \in V(Y)$ that give a quantum homomorphism from K_m to Y , where $d = r|V(Y)|$. Let S be an independent set of Y of size $\alpha(Y)$. For any $i, j \in [m]$ and $y, y' \in S$ such that $(i, y) \neq (j, y')$, either $i = j$ and $y \neq y'$, or $i \sim j$ and $y \not\sim y'$. Therefore, the projectors E_{iy} for $i \in [m]$ and $y \in S$ are pairwise orthogonal, and thus

$$d \geq \text{rk} \left(\sum_{i \in [m], y \in S} E_{iy} \right) = \sum_{i \in [m], y \in S} \text{rk}(E_{iy}) = r\alpha(Y)\omega(Y).$$

Since $d = r|V(Y)|$, we have that $\alpha(Y)\omega(Y) \leq |V(Y)|$. □

Note that the above proof actually shows something stronger, namely that if $K_m \xrightarrow{q} Y$ for vertex-transitive Y , then $m\alpha(Y) \leq |V(Y)|$. What this means is that we can replace the clique number of Y in the clique-coclique bound with the “quantum clique number” of Y , and it still holds true. We will not formally define quantum clique number until Section 6.9, but the reader is welcome to skip ahead. Later we will prove a “quantum clique-coclique bound” where both $\omega(Y)$ and $\alpha(Y)$ are replaced with their quantum analogs.

6.3 Measurement Graphs

In this section we introduce the notion of a measurement graph. This allows us to relate the existence of a quantum homomorphism from X to Y to the existence of a homomorphism from X to the measurement graph of Y .

Definition. For a potentially infinite graph Y and $d \in \mathbb{N}$, define $M(Y, d)$ to be the infinite graph whose vertices are the tuples $\mathbf{E} = (E_y)_{y \in V(Y)}$ such that $E_y \in \mathbb{C}^{d \times d}$ is a projector for all $y \in V(Y)$ and $\sum_{y \in V(Y)} E_y = I$. Two such tuples, \mathbf{E} and \mathbf{E}' , are adjacent if $E_y E'_{y'} = 0$ for all $y \not\sim y'$. We refer to the graph $M(Y, d)$ as the *measurement graph of Y in dimension d* .

Note that the vertices of $M(Y, d)$ are exactly the projective measurements in dimension d whose outcomes are indexed by the vertices of Y . The following theorem gives the motivation behind the definition of the measurement graphs.

Theorem 6.3.1. *If X and Y are (possibly infinite) graphs, then $X \xrightarrow{q} Y$ if and only if $X \rightarrow M(Y, d)$ for some $d \in \mathbb{N}$.*

Proof. Suppose that $\varphi : X \rightarrow M(Y, d)$ for some $d \in \mathbb{N}$. Let us denote $\varphi(x)$ with \mathbf{E}^x , for all $x \in V(X)$. Then, for all $x \in V(X)$ and $y \in V(Y)$, let $E_{xy} = E_y^x$. It is straightforward to check that the E_{xy} satisfy Conditions 6.1 and 6.2 and thus $X \xrightarrow{q} Y$.

Conversely, if $X \xrightarrow{q} Y$, then there exists $d \in \mathbb{N}$ and projectors $E_{xy} \in \mathbb{C}^{d \times d}$ which give a quantum homomorphism from X to Y . For each $x \in V(X)$, define $\mathbf{E}^x = (E_y^x)_{y \in V(Y)}$ such that $E_y^x = E_{xy}$ for all $y \in V(Y)$. It is easy to see that the function $\varphi : V(X) \rightarrow V(M(Y, d))$ defined by $\varphi(x) = \mathbf{E}^x$ is a homomorphism from X to $M(Y, d)$. \square

The above theorem lets us translate questions regarding quantum homomorphisms into questions regarding homomorphisms, albeit to infinite graphs. This is probably not of much practical use, but it is interesting to be able to do so.

Note that we must consider the existence of homomorphisms from X to an infinite number of graphs in the theorem above. However, we may only want to consider the existence of a homomorphism from X to a single graph. We can achieve this by simply using the disjoint union of graphs:

Define the graph $\mathfrak{M}(Y)$ to be the disjoint union of the $M(Y, d)$ for all $d \in \mathbb{N}$. We refer to $\mathfrak{M}(Y)$ as simply *the measurement graph of Y* . We can now replace “ $X \rightarrow M(Y, d)$ for some $d \in \mathbb{N}$ ” in the above theorem with “ $X \rightarrow \mathfrak{M}(Y)$ ”:

Theorem 6.3.2. *If X and Y are graphs, then $X \xrightarrow{q} Y$ if and only if $X \rightarrow \mathfrak{M}(Y)$.*

Proof. If $X \xrightarrow{q} Y$ then by Theorem 6.3.1 there exists $d \in \mathbb{N}$ such that $X \rightarrow M(Y, d)$. Since $M(Y, d)$ is a subgraph of $\mathfrak{M}(Y)$, we have that $X \rightarrow \mathfrak{M}(Y)$.

Now suppose that $\varphi : X \rightarrow \mathfrak{M}(Y)$. For all $x \in V(X)$ let $\mathbf{E}^x := \varphi(x)$, and $E_{xy} := \mathbf{E}_y^x$. Furthermore, let X_1, X_2, \dots, X_c be the components of X . Since X_i is connected, the image of X_i under φ is connected and thus is contained in $M(Y, d_i)$ for some $d_i \in \mathbb{N}$. Let $d \in \mathbb{N}$ be a common multiple of the d_i , and let $d'_i = d/d_i$. For all $x \in V(X)$ and $y \in V(Y)$, define

$$F_{xy} = E_{xy} \otimes I_{d'_i} \text{ if } x \in V(X_i).$$

It is easy to check that the E_{xy} give a quantum homomorphism from X to Y . \square

Note that the above proof does not work for infinite graphs, since we are not guaranteed that there exists a finite common multiple of an infinite number of integers. We will investigate measurement graphs further in Section 6.17, where we will see that for a graph X , the graph $M(X, d)$ is “quantum homomorphically equivalent” to X for all $d \in \mathbb{N}$.

6.4 Homomorphism Game

Here we will introduce the motivation behind the definition of quantum homomorphisms: the homomorphism game. Initially, only a special case of the homomorphism game, the graph coloring game, was introduced in [17, 9]. However, in [42] Mančinska and Roberson introduced the homomorphism game, which generalizes the coloring game in the same way that homomorphisms generalize colorings.

For graphs X and Y , the (X, Y) -homomorphism game is played between two players, usually referred to as Alice and Bob, and a referee. The general idea is that Alice and Bob are attempting to convince the referee that they know a homomorphism from X to Y . More specifically, the game is played as follows: The referee sends Alice and Bob vertices $x_A, x_B \in V(X)$ respectively, and they each reply to the referee with vertices $y_A, y_B \in V(Y)$ accordingly. Though they are allowed to agree on a strategy beforehand, Alice and Bob are not allowed to communicate during the game. In order for Alice and Bob to win the (X, Y) -homomorphism game, the following conditions must be met:

$$\text{if } x_A = x_B, \text{ then } y_A = y_B; \tag{6.3}$$

$$\text{if } x_A \sim x_B, \text{ then } y_A \sim y_B. \tag{6.4}$$

The first condition is usually referred to as the consistency condition, while the second condition is referred to as the adjacency condition as it corresponds to the adjacency-preserving property of homomorphisms. Typically, we assume that the inputs the referee sends are chosen uniformly at random, but we really only need that every possible pair of inputs (which can result in a loss of the game) is sent with some nonzero probability. The game is only played for one round, i.e., only one pair of inputs and one pair of outputs are involved. When say that Alice and Bob “can win the (X, Y) -homomorphism game”, we will mean that they have a strategy for playing the game that wins with probability one. In the special case of Y being a complete graph, the (X, Y) -homomorphism game reduces to the graph coloring game for X .

Let X and Y be graphs and let $\varphi : X \rightarrow Y$ be a homomorphism. Suppose Alice and Bob play the (X, Y) -homomorphism game as follows: Alice responds with $\varphi(x_A)$ when she

receives $x_A \in V(X)$, and Bob responds with $\varphi(x_B)$ when he receives $x_B \in V(X)$. It is easy to see that this is a winning strategy since if $x_A = x_B$ then $\varphi(x_A) = \varphi(x_B)$, and if $x_A \sim x_B$ then $\varphi(x_A) \sim \varphi(x_B)$. Conversely, suppose that Alice and Bob can win the (X, Y) -homomorphism game using some deterministic strategy. Since it is deterministic, we can define functions $f_A, f_B : V(X) \rightarrow V(Y)$ such that $f_A(x)$ is the vertex of Y Alice responds with when sent vertex $x \in V(X)$, and f_B is defined for Bob symmetrically. It is easy to see that $f_A = f_B$ must be true in order to satisfy the consistency condition, and furthermore that f_A must be a homomorphism in order to meet the adjacency condition. Therefore, Alice and Bob can win the (X, Y) -homomorphism game with a deterministic strategy if and only if $X \rightarrow Y$.

In general, “classical” players are not limited to deterministic strategies. They may use a strategy which relies on some shared randomness, such as a shared random n -bit 01-string. So after they receive their inputs, Alice and Bob can query some or all bits of their shared random string and then respond with an output which depends on both the value of their input and the values of the bits they queried. A strategy of this type is typically referred to as a “probabilistic strategy”. Given such a strategy, we can define functions $f_A, f_B : \{0, 1\}^n \times V(X) \rightarrow V(Y)$ such that $f_A(r, x_A)$ is the vertex Alice responds with upon receiving x_A given that the shared random string took on value r , and f_B defined similarly for Bob. Note that $f_A(r, x) = f_B(r, x)$ must hold for all r which occur with nonzero probability in order to guarantee that the consistency condition is always satisfied. We can then define $f_r : V(X) \rightarrow V(Y)$ for all $r \in \{0, 1\}^n$ by letting $f_r(x) = f_A(r, x)$. It is easy to see that for any r which occurs with nonzero probability, the function f_r must constitute a winning deterministic strategy for the (X, Y) -homomorphism game. In this way any winning probabilistic strategy can be viewed as a probability distribution over deterministic strategies, in which any deterministic strategy which occurs with nonzero probability must be a winning one. Therefore, classical players can win the (X, Y) -homomorphism game if and only if $X \rightarrow Y$.

In the next section we will introduce the general “quantum strategy” for the (X, Y) -homomorphism game, and show that there exists a winning quantum strategy for the (X, Y) -homomorphism game if and only if $X \xrightarrow{q} Y$. Later, in Sections 6.10 and 6.11, we will see examples of graphs X and Y such that $X \xrightarrow{q} Y$ but $X \not\rightarrow Y$.

6.5 The Quantum Strategy

For graphs X and Y , the general quantum strategy for the (X, Y) -homomorphism game is as follows: Upon receiving input $x \in V(X)$, Alice performs POVM $\mathbf{E}_x = (E_{xy})_{y \in V(Y)}$

on her part of a shared state $\psi \in \mathbb{C}^{d_A} \otimes \mathbb{C}^{d_B}$ and obtains some outcome $y \in V(Y)$, which she sends to the verifier as her answer. Bob acts similarly, except that he uses POVMs $\mathbf{F}_x = (F_{xy})_{y \in V(Y)}$ for $x \in V(X)$. The probability that Alice outputs $y \in V(Y)$ and Bob outputs $y' \in V(Y)$ upon receiving inputs $x, x' \in V(X)$ respectively is given by $\psi^*(E_{xy} \otimes F_{x'y'})\psi$. For this to be a winning strategy, the probability of the consistency and adjacency conditions not holding must be zero, i.e. we must have that

$$\begin{aligned} \psi^*(E_{xy} \otimes F_{x'y'})\psi &= 0 \text{ for } y \neq y'; \\ \psi^*(E_{xy} \otimes F_{x'y'})\psi &= 0 \text{ for } x \sim x' \text{ and } y \not\sim y'. \end{aligned}$$

Having to choose POVMs for both Alice and Bob as well as a shared entangled state makes thinking about quantum strategies quite cumbersome. Thankfully, the authors of [8] proved the following theorem for the graph coloring game which states that if there is a winning quantum strategy, then there is one which only depends on Alice's POVMs. Laura Mančinska and I showed that the same proof works more generally for the homomorphism game, and we have also filled in many details absent from the proof given in [8].

Theorem 6.5.1. *If the (X, Y) -homomorphism game can be won by a quantum strategy, then it can be won by a quantum strategy such that*

1. For some $d \in \mathbb{N}$, the E_{xy} and F_{xy} are $d \times d$ projectors;
2. $E_{xy} = F_{xy}^T$, for all $x \in V(X)$, $y \in V(Y)$;
3. $\psi = \frac{1}{\sqrt{d}} \sum_{i=1}^d e_i \otimes e_i$.

Proof. The proof is structured as follows: First we show that we can assume without loss of generality that the shared state has full Schmidt rank. Then we show that, assuming this, the first and second claims above must hold. Finally, we show that using a maximally entangled state still gives us a winning strategy.

Let $\psi' \in \mathbb{C}^{d_A} \otimes \mathbb{C}^{d_B}$ be the state that Alice and Bob share and for $x \in V(X)$, $y \in V(Y)$, let E'_{xy}, F'_{xy} be the measurement operators used by Alice and Bob respectively. We can write ψ' in its Schmidt decomposition as

$$\psi' = \sum_{i=1}^d \sqrt{\lambda_i} \alpha_i \otimes \beta_i$$

for some orthonormal sets $\{\alpha_1, \dots, \alpha_d\} \subseteq \mathbb{C}^{d_A}$, $\{\beta_1, \dots, \beta_d\} \subseteq \mathbb{C}^{d_B}$, and all $\lambda_i > 0$. Let $e_i \in \mathbb{C}^d$ be the i^{th} standard basis vector for $i \in [d]$. Now define

$$P_A := \sum_{i=1}^d e_i \alpha_i^*, \quad P_B = \sum_{i=1}^d e_i \beta_i^*,$$

and note that they satisfy $P_A P_A^* = I = P_B P_B^*$. Using P_A and P_B , we can define new POVMs and shared state for Alice and Bob as follows:

$$\begin{aligned} \psi &:= (P_A \otimes P_B) \psi' = \sum_{i=1}^d \sqrt{\lambda_i} e_i \otimes e_i, \\ E_{xy} &:= P_A E'_{xy} P_A^* \quad \forall x \in V(X), \quad y \in V(Y), \\ F_{xy} &:= P_B F'_{xy} P_B^* \quad \forall x \in V(X), \quad y \in V(Y). \end{aligned}$$

It is easy to check that the $\{E_{xy}\}_{y \in V(Y)}$ and $\{F_{xy}\}_{y \in V(Y)}$ are indeed measurements and ψ is a valid quantum state (has unit norm). Furthermore, we have that

$$\begin{aligned} (P_A^* \otimes P_B^*) \psi &= \sum_{i=1}^d \sqrt{\lambda_i} P_A^* e_i \otimes P_B^* e_i \\ &= \sum_{i=1}^d \sqrt{\lambda_i} \left(\sum_{j=1}^d \alpha_j e_j^* e_i \right) \otimes \left(\sum_{k=1}^d \beta_k e_k^* e_i \right) \\ &= \sum_{i=1}^d \sqrt{\lambda_i} \alpha_i \otimes \beta_i \\ &= \psi'. \end{aligned}$$

We must verify that these new measurements and shared state still constitute a winning strategy for the (X, Y) -homomorphism game. The following calculation shows this by proving that replacing the shared state and measurement operators with their unprimed versions does not change the probability of Alice and Bob responding with y and y' given they were sent x and x' respectively:

$$\begin{aligned} \psi^* (E_{xy} \otimes F_{x'y'}) \psi &= \psi^* ((P_A E'_{xy} P_A^*) \otimes (P_B F'_{x'y'} P_B^*)) \psi \\ &= \psi^* (P_A \otimes P_B) (E'_{xy} \otimes F'_{x'y'}) (P_A^* \otimes P_B^*) \psi \\ &= \psi'^* (E'_{xy} \otimes F'_{x'y'}) \psi' \end{aligned}$$

since $(P_A^* \otimes P_B^*)\psi = \psi'$. Clearly ψ has full Schmidt rank, so we have proven the desired claim.

Now we will show that the E_{xy} 's and F_{xy} 's are projectors. Let $D := \text{diag}(\sqrt{\lambda_i})$. This means that $\text{vec}(D) = \psi$, where $\text{vec}(D)$ is the vector obtained from D by stacking its rows on top of each other with the first row on top. In the following, the support of a matrix E , denoted $\text{supp}(E)$, refers to the column space of E . However, we will also sometimes use $\text{supp}(E)$ to refer to the orthogonal projection onto the column space of E . It will be up to the reader to distinguish which we mean from the context.

Using that $(A \otimes B) \text{vec}(X) = \text{vec}(AXB^T)$ and $\text{vec}(A)^* \text{vec}(B) = \text{Tr}(A^*B)$ for all matrices A, B, X such that these products are defined, we get that

$$\begin{aligned} \psi^*(E_{xy} \otimes F_{x'y'})\psi &= \text{vec}(D)^*(E_{xy} \otimes F_{x'y'}) \text{vec}(D) \\ &= \text{Tr}(D^* E_{xy} D F_{x'y'}^T) \\ &= \text{Tr}(E_{xy} D F_{x'y'}^T D) \end{aligned}$$

since $D^* = D$. From the consistency conditions we get that $\text{Tr}(E_{xy} D F_{x'y'}^T D) = 0$ for $y \neq y'$. Using this and the fact that $\sum_{y' \in V(Y)} E_{xy'} = I$, we obtain

$$\text{Tr}(D F_{xy}^T D) = \sum_{y' \in V(Y)} \text{Tr}(E_{xy'} D F_{xy}^T D) = \text{Tr}(E_{xy} D F_{xy}^T D).$$

By Lemma 6.1.1, this implies that the 1-eigenspace of E_{xy} contains the support of $D F_{xy}^T D$. All 1-eigenspaces of the $E_{xy'}$ are mutually orthogonal since $\sum_{y' \in V(Y)} E_{xy'} = I$ and the $E_{xy'}$ are all positive semidefinite. This implies that $\text{supp}(D F_{xy}^T D) \perp \text{supp}(D F_{xy'}^T D)$ for $y \neq y'$, where \perp denotes orthogonality of subspaces. Also, since D^2 has full rank and

$$\sum_{y' \in V(Y)} D F_{xy'}^T D = D^2$$

we have that

$$\bigoplus_{y' \in V(Y)} \text{supp}(D F_{xy'}^T D) = \mathbb{C}^d,$$

where \bigoplus denotes the direct sum of subspaces. Recall that for $y \neq y'$ we have that $\text{Tr}(E_{xy} D F_{xy'}^T D) = 0$. However, this is equivalent to $\text{supp}(E_{xy}) \perp \text{supp}(D F_{xy'}^T D)$ for $y \neq y'$. Together these two facts give

$$\text{supp}(E_{xy}) \subseteq \left(\bigoplus_{y' \neq y} \text{supp}(D F_{xy'}^T D) \right)^\perp = \text{supp}(D F_{xy}^T D).$$

Combining this with the fact that $\text{supp}(DF_{xy}^T D)$ is contained in the 1-eigenspace of E_{xy} , we get that $\text{supp}(E_{xy})$ is equal to the 1-eigenspace of E_{xy} and thus E_{xy} is a projector. This of course implies that $E_{xy} = \text{supp}(DF_{xy}^T D)$. Similarly, we can conclude that the F_{xy} are projectors and that $F_{xy} = \text{supp}(DE_{xy}^T D)$.

Now we show that D commutes with both E_{xy} and F_{xy} , and use this to obtain that $E_{xy} = F_{xy}^T$. Recall that for any two positive semidefinite matrices A and B , $\text{Tr}(AB) = 0$ if and only if $\text{Tr}(A \text{supp}(B)) = 0$ if and only if $AB = 0$. Using this and the fact that D has full rank, we get that

$$0 = \text{Tr}(DE_{xy}DF_{xy'}^T) = \text{Tr}((DE_{xy}D) \text{supp}(DE_{xy'}D)) \Leftrightarrow 0 = E_{xy}D^2E_{xy'}.$$

Therefore,

$$D^2 = \sum_{y,y' \in V(Y)} E_{xy}D^2E_{xy'} = \sum_{y \in V(Y)} E_{xy}D^2E_{xy}.$$

Multiplying the first and last expressions above on either side by $E_{xy'}$ gives that

$$D^2E_{xy'} = E_{xy'}D^2E_{xy'} = E_{xy'}D^2.$$

Since the E_{xy} commute with D^2 , they also commute with D . Similarly the F_{xy} commute with D . Using this we get that

$$E_{xy} = \text{supp}(DF_{xy}^T D) = \text{supp}(F_{xy}^T D^2) = \text{supp}(F_{xy}^T) = F_{xy}^T$$

since D is invertible. So we have proven the second claim.

Now using the fact that D commutes with F_{xy}^T , we see that

$$0 = \psi^*(E_{xy} \otimes F_{x'y'})\psi = \text{Tr}(E_{xy}(DF_{x'y'}^T D))$$

if and only if

$$E_{xy}DF_{x'y'}^T D = 0$$

if and only if

$$E_{xy}F_{x'y'}^T = 0.$$

This means that we can replace ψ with any state which will give an invertible D that commutes with the E_{xy} and F_{xy} . If we use the state $\Psi = \frac{1}{\sqrt{d}} \sum_i e_i \otimes e_i$, then D is a scalar multiple of the identity matrix and thus commutes with all other matrices. Therefore we can assume without loss of generality that we use the state Ψ and so we have proven the final claim. \square

This theorem allowed Mančinska and the author to obtain the following corollary from [42]. The same corollary was also noted in [8], though in different terms and for quantum colorings.

Corollary 6.5.2. *There exists a winning quantum strategy for the (X, Y) -homomorphism game if and only if $X \xrightarrow{q} Y$.*

Proof. Suppose that there exists a winning quantum strategy for the (X, Y) -homomorphism game. Then there exists a winning quantum strategy of the form given in Theorem 6.5.1. Suppose that E_{xy} and F_{xy} for $x \in V(X)$, $y \in V(Y)$ are Alice's and Bob's measurement operators for this strategy. Recall from the last paragraph of the proof above that the probability of Alice and Bob responding with $y, y' \in V(Y)$ after being sent $x, x' \in V(X)$ respectively is zero if and only if $0 = E_{xy}F_{x'y'}^T = E_{xy}E_{x'y'}$. This implies that $E_{xy}E_{x'y'} = 0$ whenever $x = x'$ and $y \neq y'$, or $x \sim x'$ and $y \not\sim y'$. Furthermore, since $(E_{xy})_{y \in V(Y)}$ is a measurement for all $x \in V(X)$, we have that

$$\sum_{y \in V(Y)} E_{xy} = I.$$

These are exactly the conditions required for the E_{xy} to give a quantum homomorphism from X to Y , and thus $X \xrightarrow{q} Y$.

Conversely, suppose that E_{xy} for $x \in V(X)$, $y \in V(Y)$ are projectors which give a quantum homomorphism from X to Y . It is easy to check that if Alice and Bob use the E_{xy} and E_{xy}^T as their respective measurement operators and share the state $\frac{1}{\sqrt{d}} \sum_{i=1}^d e_i \otimes e_i$, then they will win the (X, Y) -homomorphism game. \square

6.6 Basic Properties of Quantum Homomorphisms

In this section we will prove some basic properties that quantum homomorphisms have in common with homomorphisms. The most important result here is that quantum homomorphisms are transitive, as this property will be used ubiquitously in the remainder of this chapter.

Lemma 6.6.1. *The relation \xrightarrow{q} is transitive: if $X \xrightarrow{q} Y$ and $Y \xrightarrow{q} Z$, then $X \xrightarrow{q} Z$.*

Proof. Suppose that E_{xy} for $x \in V(X)$, $y \in V(Y)$ and F_{yz} for $y \in V(Y)$, $z \in V(Z)$ give quantum homomorphisms from X to Y and from Y to Z respectively. We will show that

$$P_{xz} = \sum_{y \in V(Y)} E_{xy} \otimes F_{yz}$$

for $x \in V(X)$, $z \in V(Z)$ give a quantum homomorphism from X to Z . We must first show that the P_{xz} are projectors, i.e. are Hermitian idempotents. The fact that they are Hermitian follows from the fact that they are the sum of tensor products of Hermitian matrices. To see that they are idempotent we must consider the product of P_{xz} with itself, but we will first write the product of an arbitrary pair since we will need to consider this later in the proof:

$$\begin{aligned} P_{xz}P_{x'z'} &= \left(\sum_{y \in V(Y)} E_{xy} \otimes F_{yz} \right) \left(\sum_{y' \in V(Y)} E_{x'y'} \otimes F_{y'z'} \right) \\ &= \sum_{y, y' \in V(Y)} E_{xy} E_{x'y'} \otimes F_{yz} F_{y'z'}. \end{aligned}$$

If $x = x'$, then $E_{xy}E_{x'y'} = 0$ whenever $y \neq y'$. Therefore

$$P_{xz}^2 = \sum_{y \in V(Y)} E_{xy} E_{xy} \otimes F_{yz} F_{yz} = \sum_{y \in V(Y)} E_{xy} \otimes F_{yz} = P_{xz}.$$

So we have shown that the P_{xz} are projectors. Now we will show that $(P_{xz})_{z \in V(Z)}$ is a measurement for all $x \in V(X)$.

$$\begin{aligned} \sum_{z \in V(Z)} P_{xz} &= \sum_{z \in V(Z)} \sum_{y \in V(Y)} E_{xy} \otimes F_{yz} \\ &= \sum_{y \in V(Y)} \sum_{z \in V(Z)} E_{xy} \otimes F_{yz} \\ &= \sum_{y \in V(Y)} E_{xy} \otimes \left(\sum_{z \in V(Z)} F_{yz} \right) \\ &= \sum_{y \in V(Y)} E_{xy} \otimes I \\ &= I \otimes I = I. \end{aligned}$$

Now we only have left to show that the P_{xz} satisfy the appropriate orthogonality conditions.

Recalling the expression for $P_{xz}P_{x'z'}$ given above, if $x \sim x'$ and $z \not\sim z'$, then we have

$$\begin{aligned} P_{xy}P_{x'y'} &= \sum_{y,y' \in V(Y)} E_{xy}E_{x'y'} \otimes F_{yz}F_{y'z'} \\ &= \sum_{y \sim y' \in V(Y)} E_{xy}E_{x'y'} \otimes F_{yz}F_{y'z'} \\ &= 0. \end{aligned}$$

In the above, the second equality follows from the orthogonalities of the E_{xy} 's, and the third equality follows from the orthogonalities of the F_{yz} 's. We have shown that the P_{xz} give a quantum homomorphism from X to Z and thus $X \xrightarrow{q} Z$. \square

The above lemma and proof holds equally well for infinite graphs. We can avoid an infinite sum in the definition of the P_{xz} by only summing over y such that E_{xy} is nonzero.

The next lemma we give shows that quantum homomorphisms cannot map larger complete graphs to smaller ones. It can be shown that if one has a strategy to win the (K_m, K_n) -homomorphism game for $m > n$, then one can construct a strategy to win the (X, K_n) -homomorphism game for any graph X . Because of this, the following lemma can be seen as evidence that quantum homomorphisms do not behave too strangely in comparison to homomorphisms. Note that this result was originally given in slightly different terms in [8].

Lemma 6.6.2. *For $m, n \in \mathbb{N}$, we have $K_m \xrightarrow{q} K_n$ if and only if $m \leq n$.*

Proof. If $m \leq n$, then $K_m \rightarrow K_n$ and therefore $K_m \xrightarrow{q} K_n$ by Lemma 6.2.1. For the converse, recall the proof of the clique-coclique bound from Section 6.2. There we showed that if Y is vertex transitive, then $K_m \xrightarrow{q} Y$ implies that $m\alpha(Y) \leq |V(Y)|$. Therefore, if $K_m \xrightarrow{q} K_n$, then $m = m\alpha(K_n) \leq |V(K_n)| = n$. \square

Next we show that if X is a connected graph and Y is such that $X \xrightarrow{q} Y$, then there must exist some component of Y to which X has a quantum homomorphism. For homomorphisms this is fairly straightforward, since they must map walks to walks and therefore must preserve connectedness. For quantum homomorphisms it is not so simple, since it is not even clear if there is something which could be called the image of a quantum homomorphism.

Lemma 6.6.3. *Suppose that X is a connected graph and Y is a graph with components Y_1, \dots, Y_n . If $X \xrightarrow{q} Y$, then there exists $k \in [n]$ such that $X \xrightarrow{q} Y_k$.*

Proof. If X is a single vertex then we are done, otherwise X has an edge. Suppose that $X \xrightarrow{q} Y$ and that E_{xy} for $x \in V(X), y \in V(Y)$ are projectors giving a quantum homomorphism from X to Y . For each $i \in [n]$ and $x \in V(X)$ define, E_{xi} as follows:

$$E_{xi} = \sum_{y \in V(Y_i)} E_{xy}.$$

Since $E_{xy}E_{xy'} = 0$ for $y \neq y'$, we have that the E_{xi} are projectors. Furthermore, since the $V(Y_i)$ partition $V(Y)$, we have that $E_{xi}E_{xj} = 0$ for $i \neq j$. Now suppose that $x', x'' \in V(X)$ are adjacent. Since $y \not\sim y'$ for $y \in V(Y_i), y' \in V(Y_j)$, and $i \neq j$, we have that

$$E_{x'i}E_{x''j} = 0 \text{ for all } i \neq j.$$

Since $\sum_{i \in [n]} E_{x'i} = \sum_{j \in [n]} E_{x''j} = I$, the above implies that $E_{x'i}(I - E_{x''i}) = 0$ and $(I - E_{x'i})E_{x''i} = 0$ for all $i \in [n]$. Therefore, $E_{x'i} = E_{x'i}E_{x''i} = E_{x''i}$. Since X is connected, we further obtain that, for each $i \in [n]$, there exists a projector E_i such that $E_{xi} = E_i$ for all $x \in V(X)$.

Since $\sum_{i \in [n]} E_i = I$, we can fix a $k \in [n]$ such that $E_k \neq 0$. Then for all $x \in V(X), y \in V(Y_k)$ we have that

$$\begin{aligned} \sum_{y \in V(Y_k)} E_{xy} &= E_k \\ E_{xy}E_{x'y'} &= 0 \text{ if } (x = x' \ \& \ y \neq y') \text{ or } (x \sim x' \ \& \ y \not\sim y'). \end{aligned}$$

Hence, when restricted to the image of E_k , the projectors E_{xy} for $x \in V(X), y \in V(Y_k)$ satisfy Conditions 6.1 and 6.2. Therefore $X \xrightarrow{q} Y_k$. \square

6.7 Vector Colorings

The main goal of this section is to introduce three graph parameters which are defined via assignments of vectors to vertices of graphs, and then to show that these three parameters are quantum homomorphism monotone. One of these parameters is equal to the well known Lovász theta function of the complement, while the other two are variations of this. One of the variations is the vector chromatic number of [32], though this was actually introduced many years earlier in both [47] and [37] independently. The other variation is due to Szegedy in [48].

We begin with the definition of a vector coloring which comes from [32]:

Definition. Let \mathcal{S}^d denote the unit sphere in \mathbb{R}^{d+1} . For a graph X , a map $\varphi : V(X) \rightarrow \mathcal{S}^d$ is called a *vector k -coloring* if whenever $x \sim x'$,

$$\varphi(x)^T \varphi(x') \leq -\frac{1}{k-1}.$$

The basic idea of the above definition is that adjacent vertices are assigned vectors which are “far apart”. A vector coloring can in fact be thought of as a homomorphism to the infinite graph whose vertices are the elements of \mathcal{S}^d such that vectors u and v are adjacent whenever $u^T v \leq -1/(k-1)$. Note that k need not be an integer.

A vector coloring, φ , such that $\varphi(x)^T \varphi(x') = -1/(k-1)$ for all $x \sim x'$ is called a *strict vector coloring*. If φ is a strict vector coloring which furthermore satisfies $\varphi(x)^T \varphi(x') \geq -1/(k-1)$ for all $x, x' \in V(X)$, then we say that φ is a *rigid vector coloring*. Note that both strict and rigid vector colorings can be defined in terms of homomorphisms similarly to how it was done for vector coloring above.

With these three types of colorings defined, we can define the corresponding chromatic numbers:

Vector chromatic number:

$$\bar{\vartheta}^-(X) := \text{minimum } k \text{ such that } X \text{ has a vector } k\text{-coloring.}$$

Strict vector chromatic number:

$$\bar{\vartheta}(X) := \text{minimum } k \text{ such that } X \text{ has a strict vector } k\text{-coloring.}$$

Rigid vector chromatic number:

$$\bar{\vartheta}^+(X) := \text{minimum } k \text{ such that } X \text{ has a rigid vector } k\text{-coloring.}$$

Though it is not obvious, the minimums above are in fact always attained for nonempty graphs. This is a consequence of the fact that they all can be written as semidefinite programs. For the empty graph, the minimum should be changed to an infimum which will have value 1 in that case.

The notation above originates from [34], in which Lovász defined his ϑ function. The strict vector chromatic number is equal to the Lovász ϑ function of the complement [32], hence the notation $\bar{\vartheta}(X) = \vartheta(\bar{X})$. Similarly, ϑ^- is the parameter introduced in [47] and [37], and $\bar{\vartheta}^-(X) = \vartheta^-(\bar{X})$. Furthermore, ϑ^+ is the parameter introduced in [48], and $\bar{\vartheta}^+(X) = \vartheta^+(\bar{X})$. The + and - superscripts are inspired by the fact that $\bar{\vartheta}^-(X) \leq \bar{\vartheta}(X) \leq \bar{\vartheta}^+(X)$

for all graphs X . Seeing this is not difficult as any rigid vector k -coloring is also a strict vector k -coloring, and any strict vector k -coloring is a vector k -coloring.

It is known that all three of these parameters are homomorphism monotone. As a consequence of this and the fact that $\bar{\vartheta}^-(K_n) = \bar{\vartheta}(K_n) = \bar{\vartheta}^+(K_n) = n$, we have the following inequalities:

$$\omega(X) \leq \bar{\vartheta}^-(X) \leq \bar{\vartheta}(X) \leq \bar{\vartheta}^+(X) \leq \chi(X).$$

We will later see a quantum version of this. The fact that these parameters are homomorphism monotone is not too surprising since they can all be defined in terms of homomorphisms. However, it is at least somewhat surprising that they are all quantum homomorphism monotone as well.

Theorem 6.7.1. *If $X \xrightarrow{q} Y$, then $\bar{\vartheta}^-(X) \leq \bar{\vartheta}^-(Y)$, $\bar{\vartheta}(X) \leq \bar{\vartheta}(Y)$, and $\bar{\vartheta}^+(X) \leq \bar{\vartheta}^+(Y)$.*

Proof. The proofs for all three are very similar, so we will give the proof for $\bar{\vartheta}^+$ and then note the changes needed for the other proofs.

Suppose that $X \xrightarrow{q} Y$ and that φ is a rigid vector k -coloring of Y . Let $v_y = \varphi(y)$ for all $y \in V(Y)$ and let $\alpha = -1/(k-1)$. To prove the theorem it suffices to show that X has a rigid vector k -coloring. Since $X \xrightarrow{q} Y$, Lemma 6.2.2 implies that there exist *real* projectors E_{xy} in some dimension d that give a quantum homomorphism from X to Y . For $x \in V(X)$, define vectors

$$u_x = \frac{1}{\sqrt{d}} \sum_{y \in V(Y)} v_y \otimes \text{vec}(E_{xy}).$$

Then we have

$$\begin{aligned} u_x^T u_{x'} &= \frac{1}{d} \left(\sum_{y \in V(Y)} v_y \otimes \text{vec}(E_{xy}) \right)^T \left(\sum_{y' \in V(Y)} v_{y'} \otimes \text{vec}(E_{x'y'}) \right) \\ &= \frac{1}{d} \sum_{y, y' \in V(Y)} v_y^T v_{y'} \text{Tr}(E_{xy} E_{x'y'}). \end{aligned}$$

Since $\text{Tr}(E_{xy}E_{xy'}) = 0$ for all $y \neq y'$, we see that

$$\begin{aligned} u_x^T u_x &= \frac{1}{d} \sum_{y \in V(Y)} v_y^T v_y \text{Tr}(E_{xy}E_{xy}) \\ &= \frac{1}{d} \text{Tr} \left(\sum_{y \in V(Y)} E_{xy} \right) \\ &= \frac{1}{d} \text{Tr}(I) = 1. \end{aligned}$$

So the u_x are unit vectors and now we just need to check that $u_x^T u_{x'} = \alpha$ whenever $x \sim x'$, and that $u_x^T u_{x'} \geq \alpha$ for all $x, x' \in V(X)$. If $x \sim x'$, then $\text{Tr}(E_{xy}E_{x'y'}) = 0$ whenever $y \not\sim y'$ and so

$$\begin{aligned} u_x^T u_{x'} &= \frac{1}{d} \sum_{y \sim y'} v_y^T v_{y'} \text{Tr}(E_{xy}E_{x'y'}) = \frac{\alpha}{d} \sum_{y \sim y'} \text{Tr}(E_{xy}E_{x'y'}) \\ &= \frac{\alpha}{d} \sum_{y, y' \in V(Y)} \text{Tr}(E_{xy}E_{x'y'}) = \frac{\alpha}{d} \text{Tr} \left(\left(\sum_{y \in V(Y)} E_{xy} \right) \left(\sum_{y' \in V(Y)} E_{x'y'} \right) \right) \quad (6.5) \\ &= \frac{\alpha}{d} \text{Tr}(I_d) = \alpha. \end{aligned}$$

In general, since $v_x^T v_{x'} \geq \alpha$ for all $x, x' \in V(X)$, and $\text{Tr}(AB) \geq 0$ for any positive semidefinite matrices A and B , we have that

$$u_x^T u_{x'} \geq \frac{1}{d} \sum_{y, y' \in V(Y)} \alpha \text{Tr}(E_{xy}E_{x'y'}) = \alpha$$

as in the above. Therefore $\psi : x \mapsto u_x$ is a rigid vector k -coloring.

For $\bar{\vartheta}$, the same exact proof works except that one does not need the last part showing that $u_x^T u_{x'} \geq \alpha$ for all $x, x' \in V(X)$. For $\bar{\vartheta}^-$, one also does not need this part, but the second '=' in Equation 6.5 must be replaced with a ' \leq '. Note that this added inequality uses the fact that $\text{Tr}(E_{xy}E_{x'y'}) \geq 0$. \square

Later we will see that this theorem can be used to bound certain parameters, such as quantum chromatic number, which are defined in terms of quantum homomorphisms. More generally, if $\bar{\vartheta}(X) > \bar{\vartheta}(Y)$ (or similarly for $\bar{\vartheta}^-$ or $\bar{\vartheta}^+$), then we can conclude that $X \not\rightarrow Y$. Considering that there is no known algorithm for determining whether a graph X admits a quantum homomorphism to a graph Y , this technique can be a significant help. This is made even more true by the fact that the parameters $\bar{\vartheta}^-$, $\bar{\vartheta}$, $\bar{\vartheta}^+$ can be written as semidefinite programs and thus we are able to compute them efficiently.

6.8 The Homomorphic Product and Projective Packings

The definition of a quantum homomorphism from a graph X to a graph Y is phrased in terms of an assignment of projectors to the set $V(X) \times V(Y)$ satisfying certain conditions. In this section we define a graph product which encodes the necessary orthogonality conditions of a quantum homomorphism into graph adjacencies. We then show that Condition 6.1 can be replaced with a condition on the sum of the ranks of the projectors giving a quantum homomorphism. This inspires the definition of a new graph parameter known as the projective packing number. Later we will see that this parameter is related to the quantum independence number which we will introduce in Section 6.9.

The graph product we require was originally defined in [26]:

Definition. For graphs X and Y , define their *homomorphic product*, denoted $X \times Y$, to be the graph with vertex set $V(X) \times V(Y)$ with distinct vertices (x, y) and (x', y') being adjacent if either $x = x'$, or $x \sim x'$ and $y \not\sim y'$.

Obviously, the adjacencies of the homomorphic product exactly correspond to the required orthogonalities of projectors that give a quantum homomorphism. We will see later that one can relate the existence of a homomorphism from X to Y to the existence of an independent set of size $|V(X)|$ in $X \times Y$.

When dealing with this product we will often use the following two properties:

$$K_{|V(Y)|} \rightarrow X \times Y \tag{6.6}$$

$$\overline{X \times Y} \rightarrow K_{|V(X)|}. \tag{6.7}$$

Both of these follow from the fact that the sets $V_x = \{(x, y) : y \in V(Y)\}$, for $x \in V(X)$, induce cliques of size $|V(Y)|$ and partition $V(X \times Y)$.

It is an easy observation that for $x \in V(X)$, $y \in V(Y)$, projectors $E_{xy} \in \mathbb{C}^{d \times d}$ give a quantum homomorphism from X to Y if and only if

$$\begin{aligned} E_{xy}E_{x'y'} &= 0, \text{ whenever } (x, y) \sim (x', y') \text{ in } X \times Y \text{ and} \\ \sum_{y \in V(Y)} E_{xy} &= I, \text{ for all } x \in V(X). \end{aligned}$$

Note that the second condition above implies that

$$\sum_{y \in V(Y)} \text{rk}(E_{xy}) = d, \text{ for all } x \in V(X),$$

and therefore

$$\sum_{x \in V(X), y \in V(Y)} \text{rk}(E_{xy}) = d|V(X)|.$$

Somewhat surprisingly, this condition on the sum of the ranks along with the above orthogonality conditions is sufficient to guarantee that the projectors give a quantum homomorphism.

Lemma 6.8.1. *If X and Y are graphs, then $X \xrightarrow{q} Y$ if and only if there exist projectors $E_{xy} \in \mathbb{C}^{d \times d}$ for $x \in V(X)$, $y \in V(Y)$ such that*

$$E_{xy}E_{x'y'} = 0, \text{ whenever } (x, y) \sim (x', y') \text{ in } X \times Y, \text{ and}$$

$$\sum_{x \in V(X), y \in V(Y)} \text{rk}(E_{xy}) = d|V(X)|.$$

Proof. We have already proven the only if direction, so it remains to show the if direction. For this it suffices to show that if E_{xy} for $x \in V(X)$, $y \in V(Y)$ satisfy the above two conditions, then

$$\sum_{y \in V(Y)} E_{xy} = I \text{ for all } x \in V(X).$$

For a fixed $x \in V(X)$, the vertices $\{(x, y) : y \in V(Y)\}$ are pairwise adjacent and thus the projectors E_{xy} for $y \in V(Y)$ are pairwise orthogonal. Due to this,

$$\sum_{y \in V(Y)} \text{rk}(E_{xy}) \leq d \text{ for all } x \in V(X).$$

However, we have assumed that

$$\sum_{x \in V(X), y \in V(Y)} \text{rk}(E_{xy}) = d|V(X)|,$$

and thus the above inequality must hold with equality for all $x \in V(X)$. But since the E_{xy} for fixed $x \in V(X)$ are pairwise orthogonal, this implies the desired condition. \square

The above lemma motivates the following definition:

Definition. An assignment of projectors in $\mathbb{C}^{d \times d}$ to the vertices of a graph X such that adjacent vertices receive orthogonal projectors is a *projective packing* of X . The *value* of a projective packing is the sum of the ranks of the projectors divided by d . The *projective packing number* of a graph X , denoted $\tilde{\alpha}(X)$ is the supremum of the values of projective packings of X .

Note that by the transformation given in the proof of Lemma 6.2.2, the value of $\tilde{\alpha}$ does not change if we restrict to real matrices. Also note that if X is a graph with an independent set S , then assigning the identity matrix to every vertex of S and the zero matrix to all other vertices gives a projective packing of value $|S|$. This implies that $\alpha(X) \leq \tilde{\alpha}(X)$ for all graphs X . We will be particularly interested in projective packings of $X \times Y$, and so it is useful to consider what the largest possible value of a projective packing of this graph is. For a fixed $x \in V(X)$, the set $V_x = \{(x, y) : y \in V(Y)\}$ is a clique in $X \times Y$ and thus the projectors assigned to these vertices in a projective packing must be pairwise orthogonal and thus the sum of their ranks is at most d . Therefore, the maximum value of any projective packing of $X \times Y$ is at most $|V(X)|$.

Using the terminology of projective packings, Lemma 6.8.1 can be succinctly written as: “ $X \xrightarrow{q} Y$ if and only if there exists a projective packing of $X \times Y$ of value $|V(X)|$.” Unfortunately we cannot say that the latter is equivalent to $\tilde{\alpha}(X \times Y) = |V(X)|$, since it is not clear that the supremum in the definition of $\tilde{\alpha}$ is always attained.

We will also consider the projective packing number of the complement of a graph X , which we denote as $\tilde{\omega}(X) := \tilde{\alpha}(\overline{X})$. In a projective packing of \overline{X} , we assign projectors $E_x \in \mathbb{C}^{d \times d}$ to the vertices of X such that distinct nonadjacent vertices receive orthogonal projectors. This implies that the projectors assigned the vertices of an independent set of X are pairwise orthogonal. Therefore,

$$\sum_{x \in S} \text{rk}(E_x) \leq d$$

for any independent set of X . If we replaced the projector E_x assigned to vertex x with the weight $w_x = \text{rk}(E_x)/d$ for all $x \in V(X)$, then we would have that

$$\sum_{x \in S} w_x \leq 1$$

for all independent set S of X . The reader may recall that this is exactly the condition required for a fractional clique of X . Furthermore, the value of this fractional clique is

$$\sum_{x \in V(X)} w_x = \sum_{x \in V(X)} \frac{\text{rk}(E_x)}{d},$$

which is simply the value of the projective packing of \overline{X} . This implies that $\tilde{\omega}(X) \leq \chi_f(X)$. However, we can actually do better than this. We will show that $\tilde{\alpha}(X) \leq \vartheta^-(X)$, but we need the following semidefinite programming definition of ϑ^- to do so. Below, J refers to

the all ones matrix, A refers to the adjacency matrix of X , while \bar{A} refers to the adjacency matrix of \bar{X} , and \circ denotes the entry-wise or Schur product. Also, $P \succeq 0$ denotes that P is positive semidefinite, while $P \geq 0$ denotes that every entry of P is nonnegative.

$$\begin{array}{rcc}
& \text{PRIMAL} & \text{DUAL} \\
\vartheta^-(X) & = \max & \text{Tr}(JP) & = \min & \lambda \\
& \text{s.t.} & P \circ A = 0 & \text{s.t.} & Q \circ I = (\lambda - 1)I \\
& & \text{Tr}(P) = 1 & & Q \circ \bar{A} \leq -\bar{A} \\
& & P \geq 0 & & Q \succeq 0. \\
& & P \geq 0 & &
\end{array}$$

The primal SDP above is the original definition given by Schrijver in [47]. There is not a straightforward way to see that this definition (after taking complements) is equivalent to the vector coloring definition of $\bar{\vartheta}^-$. However, Schrijver also gave the above dual semidefinite program and showed that strong duality holds. Using this, one can recover a vector coloring of \bar{X} from a solution to the dual above by normalizing the vectors for which Q is a Gram matrix.

We are now able to prove the following:

Lemma 6.8.2. *For any graph X ,*

$$\tilde{\alpha}(X) \leq \vartheta^-(X).$$

Proof. Suppose that assigning the projector $E_x \in \mathbb{C}^{d \times d}$ to x for each $x \in V(X)$ is a projective packing of X . We must find a solution to the above SDP which has objective value at least that of this projective packing. Define

$$E = \sum_{x \in V(X)} E_x.$$

Recall that the rank of a projector is equal to its trace and thus the value of the projective packing is

$$\frac{1}{d} \sum_{x \in V(X)} \text{rk}(E_x) = \frac{1}{d} \sum_{x \in V(X)} \text{Tr}(E_x) = \frac{1}{d} \text{Tr} \left(\sum_{x \in V(X)} E_x \right) = \frac{\text{Tr}(E)}{d}.$$

Consider the matrix M whose xx' -entry is

$$\text{vec}(E_x)^T \text{vec}(E_{x'}) = \text{Tr}(E_x E_{x'}).$$

We will show that some positive multiple of M is a feasible solution to the SDP for ϑ^- with objective value at least $\text{Tr}(E)/d$. First, since M is a Gram matrix, it is positive semidefinite. Second, since $\text{Tr}(AB) \geq 0$ for any positive semidefinite matrices A and B , we have that $M \geq 0$. Furthermore, for $x \sim x'$, we have that

$$M_{xx'} = \text{Tr}(E_x E_{x'}) = 0$$

by the definition of projective packing.

Note that multiplying by a positive constant does not change any of these three facts. The expression $\text{Tr}(JM)$ is equal to the sum of the entries of M , and therefore

$$\text{Tr}(JM) = \sum_{x, x' \in V(X)} \text{Tr}(E_x E_{x'}) = \text{Tr} \left(\sum_{x, x' \in V(X)} E_x E_{x'} \right) = \text{Tr} \left(\left(\sum_{x \in V(X)} E_x \right)^2 \right) = \text{Tr}(E^2).$$

Also,

$$\text{Tr}(M) = \sum_{x \in V(X)} \text{Tr}(E_x E_x) = \sum_{x \in V(X)} \text{Tr}(E_x) = \text{Tr}(E).$$

Thus, $P = M/\text{Tr}(E)$ is a feasible solution to the SDP, and it remains to check its objective value. Since $\text{Tr}(JM) = \text{Tr}(E^2)$, we have that

$$\text{Tr}(JP) = \frac{\text{Tr}(E^2)}{\text{Tr}(E)}.$$

We need that $\text{Tr}(JP) \geq \text{Tr}(E)/d$ and thus we must show that

$$\text{Tr}(E^2) \geq \frac{(\text{Tr}(E))^2}{d}.$$

Recall that the trace of a matrix is equal to the sum of its eigenvalues. If $\lambda_1, \lambda_2, \dots, \lambda_d$ are the eigenvalues of E including multiplicities, then $\lambda_1^2, \lambda_2^2, \dots, \lambda_d^2$ are the eigenvalues of E^2 . So it is equivalent to show that

$$\sum_{i=1}^d \lambda_i^2 \geq \frac{\left(\sum_{i=1}^d \lambda_i \right)^2}{d}.$$

But this inequality can be proved by applying Cauchy-Schwarz to the vector $(\lambda_1, \dots, \lambda_d)$ and the all ones vector. \square

The above lemma gives us the following corollary:

Corollary 6.8.3. *If $X \xrightarrow{q} Y$, then*

$$\vartheta^-(X \times Y) = |V(X)|.$$

Proof. First note that by property 6.7 of the homomorphic product, we have that

$$\vartheta^-(X \times Y) = \bar{\vartheta}^-(\overline{X \times Y}) \leq \chi(\overline{X \times Y}) \leq |V(X)|$$

for any graphs X and Y . Now suppose that $X \xrightarrow{q} Y$. By prior discussions we know that $\tilde{\alpha}(X \times Y) = |V(X)|$, and thus by Lemma 6.8.2 we have

$$\vartheta^-(X \times Y) \geq \tilde{\alpha}(X \times Y) = |V(X)|.$$

□

Obviously, the above lemma is weaker than the fact that $X \xrightarrow{q} Y$ implies that $\tilde{\alpha}(X \times Y) = |V(X)|$, but since ϑ^- is efficiently computable the above lemma is probably of more practical use.

The parameter $\tilde{\alpha}$ is not homomorphism monotone. This is most easily seen by adding isolated vertices to a graph and considering how this effects the value of $\tilde{\alpha}$. However, $\tilde{\omega}$ is homomorphism monotone, and is even quantum homomorphism monotone:

Theorem 6.8.4. *Let X and Y be graphs. If $X \xrightarrow{q} Y$ and \overline{X} has a projective packing of value γ , then \overline{Y} has a projective packing of value γ and thus*

$$\tilde{\omega}(X) \leq \tilde{\omega}(Y).$$

Proof. Suppose that $E_{xy} \in \mathbb{C}^{d_1 \times d_1}$ for $x \in V(X)$, $y \in V(Y)$ are projectors which give a quantum homomorphism from X to Y . Also, let $F_x \in \mathbb{C}^{d_2 \times d_2}$ for $x \in V(X)$ be a projective packing of \overline{X} with value γ . We will construct a projective packing of \overline{Y} with value γ . Define

$$P_y = \sum_{x \in V(X)} E_{xy} \otimes F_x$$

for all $y \in V(Y)$. For $y, y' \in V(Y)$ we have

$$\begin{aligned} P_y P_{y'} &= \left(\sum_{x \in V(X)} E_{xy} \otimes F_x \right) \left(\sum_{x' \in V(X)} E_{x'y'} \otimes F_{x'} \right) \\ &= \sum_{x, x' \in V(X)} E_{xy} E_{x'y'} \otimes F_x F_{x'}. \end{aligned}$$

Suppose that $y = y'$. If $x \neq x'$, then either $x \not\sim x'$ in X and thus $F_x F_{x'} = 0$, or $x \sim x'$ in X and $E_{xy} E_{x'y} = 0$. Therefore,

$$P_y P_y = \sum_{x \in V(X)} E_{xy} E_{xy} \otimes F_x F_x = P_y,$$

and thus P_y is a projector since it is clearly Hermitian. Note that this argument also showed that the P_y are the sum of pairwise orthogonal projectors, and thus has rank equal to the sum of their ranks.

Now suppose that $y \neq y'$ and $y \not\sim y'$ in Y . If $x = x'$ or $x \sim x'$ in X , then $E_{xy} E_{x'y'} = 0$, and if $x \neq x'$ and $x \not\sim x'$ in X , then $F_x F_{x'} = 0$. Therefore, for $y \sim y'$ in \bar{Y} , we have that $P_y P_{y'} = 0$. This shows that the P_y are a projective packing.

Considering the value of this projective packing we see that

$$\begin{aligned} \sum_{y \in V(Y)} \text{rk}(P_y) &= \sum_{y \in V(Y)} \text{rk} \left(\sum_{x \in V(X)} E_{xy} \otimes F_x \right) \\ &= \sum_{x \in V(X), y \in V(Y)} \text{rk}(E_{xy}) \text{rk}(F_x) \\ &= \sum_{x \in V(X)} \text{rk}(F_x) \sum_{y \in V(Y)} \text{rk}(E_{xy}) \\ &= \sum_{x \in V(X)} \text{rk}(F_x) d_1 \\ &= \gamma d_2 d_1. \end{aligned}$$

Therefore, the projective packing P_y has value $\gamma d_1 d_2 / d_1 d_2 = \gamma$, and we are done. \square

6.9 Defining Quantum Parameters

In Section 2.3 we saw several examples of parameters which can be defined in terms of homomorphisms, such as chromatic, clique, and independence numbers. We also saw that $\bar{\vartheta}^-$, $\bar{\vartheta}$, and $\bar{\vartheta}^+$ can be defined in terms of homomorphisms in Section 6.7. Now that we have defined quantum homomorphisms, it is natural to define quantum analogs of these parameters by simply replacing “homomorphism” with “quantum homomorphism” in the

definition. Here we consider the following parameters which were originally introduced by Mančinska and Roberson in [42]:

$$\begin{aligned}
\text{quantum chromatic number:} & \quad \chi_q(X) := \min\{n \in \mathbb{N} : X \xrightarrow{q} K_n\}; \\
\text{quantum clique number:} & \quad \omega_q(X) := \max\{n \in \mathbb{N} : K_n \xrightarrow{q} X\}; \\
\text{quantum independence number:} & \quad \alpha_q(X) := \omega_q(\overline{X}); \\
\text{quantum odd girth:} & \quad \text{og}_q(X) := \min\{n \in \mathbb{N}, n \text{ odd} : C_n \xrightarrow{q} X\}.
\end{aligned}$$

Note that the above definitions of quantum clique and independence numbers are different from those given in [4] and [12]. The quantum clique number of [4] and the various independence numbers of [12] are defined in terms of the amount of quantum or classical information one can send over a quantum channel, and are therefore more analogous to capacities than our notion of quantum independence number. On the other hand, the definition of quantum chromatic number given above is equivalent to the definition given in previous works such as [3, 8, 16, 45, 36]. There they defined χ_q in terms of the graph coloring game, but as we mentioned above, this is simply the (X, Y) -homomorphism game with Y being a complete graph.

In the next few sections we will investigate these parameters in greater detail, specifically the quantum chromatic and quantum independence numbers. Here we will discuss some of their basic properties, and how they relate to quantum homomorphisms.

If X and Y are graphs such that $X \xrightarrow{q} Y$, and $c = \chi_q(Y)$, then we have that $X \xrightarrow{q} Y \xrightarrow{q} K_c$ and thus $X \xrightarrow{q} K_c$. Therefore we have that $\chi_q(X) \leq \chi_q(Y)$. In other words, χ_q is quantum homomorphism monotone. Similarly, one can see that ω_q is quantum homomorphism monotone. Furthermore, og_q and α_q relate to quantum homomorphisms in the same way that their classical counterparts do.

Since $X \rightarrow Y$ implies that $X \xrightarrow{q} Y$, we have that $\omega(X) \leq \omega_q(X)$ and $\chi_q(X) \leq \chi(X)$ for all graphs X , and similarly for other quantum parameters. Furthermore, by Lemma 6.6.2 we can see that $\omega_q(X) \leq \chi_q(X)$, analogously to their non-quantum versions.

One might wonder why we have not defined quantum versions of $\bar{\vartheta}^-$, $\bar{\vartheta}$, and $\bar{\vartheta}^+$. We mentioned above that these parameters can be defined in terms of homomorphisms to certain classes of graphs, and thus we could define their quantum analogs by simply replacing “homomorphism” with “quantum homomorphism” as we did above. However, Theorem 6.7.1 is actually equivalent to the fact that $\bar{\vartheta}^-$, $\bar{\vartheta}$, and $\bar{\vartheta}^+$ are equal to their quantum analogs.

6.10 Quantum Chromatic Number

As the quantum chromatic number and quantum colorings were precursors to quantum homomorphisms, these special cases have already been relatively well-studied [3, 8, 16, 45, 36]. In particular, a family of graphs exhibiting an exponential separation between quantum and classical chromatic numbers is known [5, 6, 3]. Furthermore, the quantum colorings of these graphs use only rank one projectors. In [8], Cameron et. al. show that when restricted to using rank one projectors, larger than exponential separations cannot be achieved, and thus the above separation is best possible in this case. However, in [16] it is shown that rank one projectors are not always sufficient to attain $\chi_q(X)$, and therefore it may be possible to have larger separations in general. In particular, it is not known if there exists a sequence of graphs X_n for which $\lim_{n \rightarrow \infty} \chi(X_n) = \infty$ but $\lim_{n \rightarrow \infty} \chi_q(X_n) < \infty$. Another fundamental question regarding quantum colorings which is still open, is whether determining if a graph has a quantum n -coloring is decidable. Hopefully, the study of quantum homomorphisms in general can help answer these and other questions regarding quantum colorings.

Our main result concerning quantum chromatic number is the following lower bound which is a direct consequence of Theorem 6.7.1:

Corollary 6.10.1. *For any graph X ,*

$$\bar{\vartheta}^+(X) \leq \chi_q(X).$$

Proof. Let $c = \chi_q(X)$. Then $X \xrightarrow{q} K_c$ and by Theorem 6.7.1, $\bar{\vartheta}^+(X) \leq \bar{\vartheta}^+(K_c) = c$. \square

The author originally proved the above lower bound, as well as Theorem 6.7.1, for $\bar{\vartheta}$ only. At the time, this was a new and best lower bound on χ_q ¹. However, J. Briët, H. Buhrman, M. Laurent, and T. Piovesan later showed that $\bar{\vartheta}^+$ is a lower bound on another parameter which also lower bounds χ_q , implying the above corollary. The author then extended Theorem 6.7.1 to $\bar{\vartheta}^-$ and $\bar{\vartheta}^+$ which gives the above proof of the corollary.

The above lower bound is of great practical use, since $\bar{\vartheta}^+$ can be written as a semidefinite program and can be computed efficiently [22]. This is particularly nice since it is not even known if there exists an algorithm which computes $\chi_q(X)$ in finite time.

We now use this lower bound to compute the exact value of the quantum chromatic number for a well-known class of graphs. We will actually use $\bar{\vartheta}$ instead of $\bar{\vartheta}^+$, but this suffices in this case. For $n \in \mathbb{N}$, let Ω_n be the graph whose vertices are the ± 1 vectors of

¹This had also been proved (but not published) by M. Laurent, G. Scarpa, and A. Varvitsiotis.

length n with orthogonal vectors being adjacent. We only consider the case when $4|n$, since otherwise Ω_n is either empty or bipartite. In such a case, a result of Frankl and Rödl [15] implies that $\chi(\Omega_n)$ is exponential in n . On the other hand, it is known that $\chi_q(\Omega_n) \leq n$ for all $n \in \mathbb{N}$ [3]. However, it remained unknown whether this inequality is tight for $4|n$. To show this, we first compute $\bar{\vartheta}(\Omega_n)$ for $4|n$.

Lemma 6.10.2. *If n is divisible by 4, then $\bar{\vartheta}(\Omega_n) = n$.*

*Proof.*² Let $n = 4m$. In [34], it was shown that if X is regular (all vertices have equal degree) and edge transitive, then

$$\vartheta(X) = |V(X)| \frac{\lambda_{\min}}{\lambda_{\min} - \lambda_{\max}}$$

where λ_{\max} and λ_{\min} are the largest and smallest eigenvalues of the adjacency matrix of X , respectively. Since Ω_n is regular and edge transitive, we can use this equality to compute $\vartheta(\Omega_n)$. The maximum eigenvalue of a regular graph is just the degree of the graph, and so for Ω_n this is equal to $\binom{4m}{2m}$. The minimum eigenvalue of Ω_n for $4|n$ was computed in [18] and is equal to $-\binom{4m}{2m}/(4m-1)$. Therefore,

$$\vartheta(\Omega_n) = 2^n \frac{-\frac{\binom{4m}{2m}}{4m-1}}{-\frac{\binom{4m}{2m}}{4m-1} - \binom{4m}{2m}} = 2^n \frac{1}{1 + (4m-1)} = \frac{2^n}{n}.$$

Another result from [34] states that for a vertex transitive graph X ,

$$\vartheta(X)\bar{\vartheta}(X) = |V(X)|.$$

Since Ω_n is vertex transitive as well, we obtain that

$$\bar{\vartheta}(\Omega_n) = n$$

as desired. □

Combining the above with Corollary 6.10.1 yields the following:

Corollary 6.10.3. *If n is divisible by 4, then $\chi_q(\Omega_n) = n$.* □

Curiously, $\{\Omega_{4n}\}_{n \in \mathbb{N}}$ is an infinite family of graphs for which the quantum chromatic number is known exactly, while the chromatic number remains unknown (for $n > 2$). The above two results are joint work with Laura Mančinska.

²This proof is a minor modification of a proof from unpublished notes of Simone Severini.

6.11 Quantum Independence Number

In the previous section we saw examples of graphs X such that $X \xrightarrow{q} K_n$ but $X \not\rightarrow K_n$ for some $n \in \mathbb{N}$. In this section, we will see examples of graphs Z such that $K_m \xrightarrow{q} Z$ but $K_m \not\rightarrow Z$ for some $m \in \mathbb{N}$, i.e. $\omega(Z) < \omega_q(Z)$. We will actually phrase our results in terms of independence numbers as opposed to clique numbers for reasons which will become clear below. Our main result will be that for any two graphs X and Y such that $X \xrightarrow{q} Y$ but $X \not\rightarrow Y$, we have that $\alpha(X \times Y) < \alpha_q(X \times Y) = |V(X)|$. In order to show this, we first show that the projective packing number upper bounds the quantum independence number.

Lemma 6.11.1. *If X is a graph, then X has a projective packing of value $\alpha_q(X)$ and therefore*

$$\alpha_q(X) \leq \tilde{\alpha}(X).$$

Proof. Suppose that $m = \alpha_q(X)$, and thus $K_m \xrightarrow{q} \overline{X}$. Note that $\overline{K_m}$ has a projective packing of value m given by assigning the identity matrix to every vertex. Therefore by Theorem 6.8.4, X has a projective packing of value m . \square

Applying Lemma 6.8.2 immediately gives us the following corollary:

Corollary 6.11.2. *For any graph X ,*

$$\alpha_q(X) \leq \vartheta^-(X) \text{ and } \omega_q(X) \leq \bar{\vartheta}^-(X).$$

\square

We note that the above corollary also follows easily from the monotonicity of $\bar{\vartheta}^-$ with respect to quantum homomorphisms. As with Corollary 6.8.3, the above corollary is weaker than its preceding lemma, but perhaps of more practical use.

The next result is the solution to Exercise 7 in Chapter 2 of [26], and is likely the original motivation for the definition of the homomorphic product.

Lemma 6.11.3. *For graphs X and Y , we have that $X \rightarrow Y$ if and only if $\alpha(X \times Y) = |V(X)|$.*

Proof. Let $m = |V(X)|$ and $n = |V(Y)|$ and note that

$$\alpha(X \times Y) = \omega(\overline{X \times Y}) \leq \chi(\overline{X \times Y}) \leq m$$

by Property (6.7). So $\alpha(X \times Y) = m$ if and only if $X \times Y$ has an independent set of size m .

Suppose that $\varphi : X \rightarrow Y$ is a homomorphism. Let $S_\varphi = \{(x, \varphi(x)) : x \in V(X)\}$. Note that $|S_\varphi| = m$. For $x, x' \in V(X)$, either $x \not\sim x'$ or $x \sim x'$. In the former case, it is easy to see that $(x, \varphi(x)) \not\sim (x', \varphi(x'))$ in $X \times Y$. In the latter case, the vertices $\varphi(x)$ and $\varphi(x')$ are adjacent in Y and thus $(x, \varphi(x)) \not\sim (x', \varphi(x'))$ in $X \times Y$. Therefore, S_φ is an independent set of size m in $X \times Y$.

Conversely, let S be an independent set of size m in $X \times Y$. Since the set $\{(x, y) : y \in V(Y)\}$ induces a clique for all $x \in V(X)$, and there are m such sets partitioning $V(X \times Y)$, there must be exactly one vertex of S whose first coordinate is x for every $x \in V(X)$. Define $\phi : V(X) \rightarrow V(Y)$ as follows:

$$\phi(x) = y \text{ for the unique } y \text{ such that } (x, y) \in S.$$

It is straightforward to check that ϕ is a homomorphism. □

Somewhat surprisingly, we are able to prove a quantum analog of the above lemma. The “only if” direction of the lemma below is joint work with Laura Mančinska, as is the corollary following it.

Lemma 6.11.4. *For graphs X and Y , we have that $X \xrightarrow{q} Y$ if and only if $\alpha_q(X \times Y) = |V(X)|$.*

Proof. Let $m = |V(X)|$. Note that

$$\alpha_q(X \times Y) \leq \tilde{\alpha}(X \times Y) \leq m$$

as discussed above. So $\alpha_q(X \times Y) = m$ if and only if $K_m \xrightarrow{q} \overline{X \times Y}$. Suppose that $X \xrightarrow{q} Y$ and that E_{xy} for $x \in V(X)$, $y \in V(Y)$ are projectors which give a quantum homomorphism from X to Y . We must use these to construct projectors which give a quantum homomorphism from K_m to $\overline{X \times Y}$. Since $m = |V(X)|$, we may assume that K_m and X have the same vertex set. For all $x_1, x_2 \in V(X)$ and $y \in V(Y)$, define $F_{x_1(x_2, y)}$ as follows:

$$F_{x_1(x_2, y)} = \begin{cases} E_{x_1 y} & \text{if } x_1 = x_2 \\ 0 & \text{o.w.} \end{cases}$$

It is easy to see that the $F_{x_1(x_2, y)}$ are projectors, and that

$$\sum_{(x, y) \in V(X \times Y)} F_{x_1(x, y)} = I$$

for all $x_1 \in V(K_m) = V(X)$. So we only need to show that $F_{x_1(x_2,y)}F_{x'_1(x'_2,y')} = 0$ whenever $x_1 \sim x'_1$ in K_m and $(x_2, y) \not\sim (x'_2, y')$ in $\overline{X \times Y}$. This trivially holds whenever $x_1 \neq x_2$ or $x'_1 \neq x'_2$, so we may assume that $x_1 = x_2 = x$ and $x'_1 = x'_2 = x'$ for some $x, x' \in V(X)$. In this case, $F_{x_1(x_2,y)}F_{x'_1(x'_2,y')} = E_{xy}E_{x'y'}$. If $x \sim x'$ in K_m , then $x \neq x'$. Given this, if $(x, y) \not\sim (x', y')$ in $\overline{X \times Y}$, then $(x, y) \sim (x', y')$ in $X \times Y$ and thus $x \sim x'$ in X and $y \not\sim y'$ in Y . However, this implies that $E_{xy}E_{x'y'} = 0$ as desired. Therefore, $K_m \xrightarrow{q} \overline{X \times Y}$ and thus $\alpha_q(X \times Y) = m$.

Conversely, suppose that $\alpha_q(X \times Y) = |V(X)|$. By Lemma 6.11.1, there exists a projective packing of $X \times Y$ of value $|V(X)|$ and therefore $X \xrightarrow{q} Y$. \square

As a consequence of the above two lemmas, we have the following:

Corollary 6.11.5. *If $X \xrightarrow{q} Y$ but $X \not\rightarrow Y$, then*

$$\alpha(X \times Y) < \alpha_q(X \times Y) = |V(X)|.$$

Proof. If $X \xrightarrow{q} Y$, then $\alpha_q(X \times Y) = |V(X)|$ by Lemma 6.11.4. On the other hand, since $X \not\rightarrow Y$ we have that $\alpha(X \times Y) < |V(X)|$ by Lemma 6.11.3. \square

With this corollary we can convert the separation between quantum and classical chromatic number given in the previous section into a separation between quantum and classical independence numbers. However, the above corollary does not tell us anything about the size of this separation. Since the separation for chromatic numbers was large, we would like the separation for independence numbers to be large as well. Though we will not do so here, we note that a large separation can be obtained by noting that $X \times K_n = X \square K_n$ and applying a result of Vizing's which states that for any graphs X and Y ,

$$\alpha(X \square Y) \leq \min\{\alpha(X)|V(Y)|, \alpha(Y)|V(X)|\}.$$

The full proof uses a result of Frankl and Rödl [15] which bounds the size of independent sets of Ω_{4n} , and can be found in [42].

6.12 No Quantum No-Homomorphism Lemma

In this section we give a simple proof of the quantum clique-coclique bound, and briefly discuss why a quantum version of the no-homomorphism lemma does not hold.

Lemma 6.12.1. *If X is a vertex transitive graph, then*

$$\alpha_q(X)\omega_q(X) \leq |V(X)|.$$

Proof. In [34], it is shown that $\vartheta(X)\bar{\vartheta}(X) = |V(X)|$ for any vertex transitive graph X . Applying Corollary 6.11.2 we have

$$\alpha_q(X)\omega_q(X) \leq \vartheta(X)\bar{\vartheta}(X) = |V(X)|.$$

□

There is actually a nice direct proof of the above which uses the (K_m, X) - and (K_n, \bar{X}) -homomorphism games for $m = \omega_q(X)$ and $n = \alpha_q(X)$. One can combine winning strategies for these games to construct a winning strategy for the $(K_{mn}, K_{|V(X)|})$ -homomorphism game thus proving the result. The technique for combining the two strategies is similar to the proof of Lemma 2.7.2 given in Section 2.8.

Recall from Section 2.7 that the clique-coclique bound is a special case of the no-homomorphism lemma. Since a quantum version of the clique-coclique bound holds, it is natural to ask whether a quantum version of the no-homomorphism lemma holds as well. There are perhaps several ways to formulate a quantum analog of the no-homomorphism lemma, but here we will consider the following: If Y is a vertex transitive graph and $X \xrightarrow{q} Y$, then

$$\frac{\alpha_q(X)}{|V(X)|} \geq \frac{\alpha_q(Y)}{|V(Y)|}.$$

The quantum clique-coclique bound is the special case of the above where X is a complete graph. If the above holds, then it must hold when Y is a complete graph, in which case we obtain the following: If $X \xrightarrow{q} K_n$, then

$$\frac{\alpha_q(X)}{|V(X)|} \geq \frac{1}{n}.$$

Letting $n = \chi_q(X)$, one can rearrange the above to obtain

$$\chi_q(X) \geq \frac{|V(X)|}{\alpha_q(X)},$$

which can be viewed as a quantum analog of the well-known “trivial” lower bound on chromatic number. It turns out that even this special case of the possible quantum no-homomorphism lemma does not hold. The counterexample is a graph we have already seen, the graph Ω_n .

Since $\chi_q(\Omega_n) = n$ for $4|n$, if the above quantum version of the no-homomorphism lemma were true, then we would have that $\alpha_q(\Omega_n) \geq 2^n/n$ whenever $4|n$. However, for $4|n$

$$\alpha_q(\Omega_n) \leq \vartheta(\Omega_n) = 2^n/n.$$

Since $\alpha_q(\Omega_n)$ must be an integer, if n is a multiple of four and not a power of two, then $\alpha_q(\Omega_n) < 2^n/n$ contradicting the above quantum analog of the no-homomorphism lemma.

The fact that we cannot even say that $\chi_q(X) \geq |V(X)|/\alpha_q(X)$ is a little surprising given the other ways in which quantum homomorphisms behave similarly to homomorphisms. However, in Section 6.13 we will introduce a parameter which is defined similarly to χ_q but does satisfy the above inequality with α_q replaced with $\tilde{\alpha}$.

6.13 Projective Rank

Here we will introduce a new graph parameter, projective rank, which is closely connected to both $\tilde{\alpha}$ and χ_q . The definitions of projective representations and projective rank given below were originally given in [42] and are a joint work with Laura Mančinska.

Recall from Lemma 6.2.3 that $X \xrightarrow{q} Y$ if and only if there exists $d \in \mathbb{N}$ and projectors $E_{xy} \in \mathbb{C}^{d \times d}$ for $x \in V(X)$, $y \in V(Y)$ which give a quantum homomorphism from X to Y and all have the same rank. Suppose that such projectors have rank r . Recall that for each $x \in V(X)$ we have

$$\sum_{y \in V(Y)} E_{xy} = I.$$

Since the terms in the above sum are pairwise orthogonal, this is equivalent to

$$\sum_{y \in V(Y)} \text{rk}(E_{xy}) = d.$$

Since $\text{rk}(E_{xy}) = r$ for all $x \in V(X)$ and $y \in V(Y)$, this is equivalent to $d = r|V(Y)|$. Therefore, we can replace the summing to identity requirement with the requirement that $d = r|V(Y)|$. Since the orthogonality requirements for the E_{xy} correspond to the adjacencies in $X \times Y$, as discussed in Section 6.8, we have the following observation: There exists a quantum homomorphism from a graph X to vertex transitive graph Y if and only if there exists an assignment of rank r projectors of $\mathbb{C}^{d \times d}$ to the vertices of $X \times Y$ such that adjacent vertices receive orthogonal projectors and $d = r|V(Y)|$. This observation motivates the following definition.

Definition. A d/r -projective representation (or simply a d/r -representation) of a graph X is an assignment of rank r projectors in $\mathbb{C}^{d \times d}$ to the vertices of X such that adjacent vertices are assigned orthogonal projectors. We say that the *value* of a d/r -representation is the rational number $\frac{d}{r}$.

Note that a 3/1-representation is not a 6/2-representation, though they do have the same value, namely 3. In the language of projective representations, we can rephrase the above observation as follows: A graph X has a quantum homomorphism to a vertex transitive graph Y if and only if $X \times Y$ has a d/r -representation for some $d, r \in \mathbb{N}$ of value $|V(Y)|$. As the graph $X \times Y$ always contains a clique of size $|V(Y)|$, it is easy to see that this is the smallest possible value of any projective representation of $X \times Y$. This prompts the definition of projective rank:

Definition. The *projective rank* of a graph X , denoted $\xi_f(X)$, is given by

$$\xi_f(X) = \inf \left\{ \frac{d}{r} : X \text{ has a } d/r\text{-representation} \right\}.$$

The reader may have noticed that a projective representation is simply a projective packing with the extra condition that all projectors must have the same rank. This is correct, though the value of a d/r -representation of a graph X is d/r while the same assignment of projectors viewed as a projective packing would have value $r|V(X)|/d$. The reason we define such seemingly similar parameters is that one of them behaves as a type of independence number while the other behaves like a chromatic number. In fact, the following lemma shows that these two parameters share the same relationship that independence number and fractional chromatic number share.

Lemma 6.13.1. *For any graph X ,*

$$\xi_f(X) \geq \frac{|V(X)|}{\tilde{\alpha}(X)}.$$

Furthermore, if X is vertex transitive then equality holds.

Proof. As noted above, a d/r -representation is also a projective packing of value

$$\frac{r|V(X)|}{d} = \frac{|V(X)|}{d/r}.$$

This implies that $\tilde{\alpha}(X) \geq |V(X)|/\xi_f(X)$ which gives the above inequality.

To prove that equality holds when X is vertex transitive, it suffices to show that if X has a projective packing of value γ , then it has a projective packing of value γ in which all of the projectors have the same rank. This we proceed to do.

Suppose that $x \mapsto E_x \in \mathbb{C}^{d \times d}$ for $x \in V(X)$ is a projective packing of value γ . Consider the rows and columns of $|\text{Aut}(X)| \times |\text{Aut}(X)|$ matrices as being indexed by the elements

of $\text{Aut}(X)$. For each $\sigma \in \text{Aut}(X)$, let M_σ be the matrix whose $\sigma\sigma$ -entry is 1 and all other entries are 0. Note that $M_\sigma M_{\sigma'} = \delta_{\sigma\sigma'} M_\sigma$, where $\delta_{\sigma\sigma'}$ is the Kronecker delta function. For all $x \in V(X)$, let

$$F_x = \sum_{\sigma \in \text{Aut}(X)} E_{\sigma(x)} \otimes M_\sigma.$$

We claim that this is a projective packing of X of value γ such the F_x all have the same rank. First note that the terms in the sum defining F_x are pairwise orthogonal due to the property of the M_σ 's noted above. Therefore, the F_x are projectors. For $x \sim x' \in V(X)$,

$$\begin{aligned} F_x F_{x'} &= \sum_{\sigma, \sigma' \in \text{Aut}(X)} E_{\sigma(x)} E_{\sigma'(x')} \otimes M_\sigma M_{\sigma'} \\ &= \sum_{\sigma \in \text{Aut}(X)} E_{\sigma(x)} E_{\sigma(x')} \otimes M_\sigma \\ &= 0 \end{aligned}$$

since σ preserves adjacency. Now consider the rank of F_x :

$$\begin{aligned} \text{rk}(F_x) &= \text{rk} \left(\sum_{\sigma \in \text{Aut}(X)} E_{\sigma(x)} \otimes M_\sigma \right) \\ &= \sum_{\sigma \in \text{Aut}(X)} \text{rk}(E_{\sigma(x)}) \text{rk}(M_\sigma) \\ &= \sum_{\sigma \in \text{Aut}(X)} \text{rk}(E_{\sigma(x)}). \end{aligned}$$

Since X is vertex transitive, this does not depend on x , and so all of the projectors have the same rank. If we let $R = \gamma d$, then R is equal to the sum of the ranks of the E_x . Considering the sum of the ranks of the F_x , we see that

$$\begin{aligned} \sum_{x \in V(X)} \text{rk}(F_x) &= \sum_{x \in V(X)} \sum_{\sigma \in \text{Aut}(X)} \text{rk}(E_{\sigma(x)}) \\ &= \sum_{\sigma \in \text{Aut}(X)} \sum_{x \in V(X)} \text{rk}(E_{\sigma(x)}) \\ &= \sum_{\sigma \in \text{Aut}(X)} R \\ &= R |\text{Aut}(X)|. \end{aligned}$$

As the F_x are $d|\text{Aut}(X)| \times d|\text{Aut}(X)|$ matrices, they give a projective packing of value

$$\frac{R|\text{Aut}(X)|}{d|\text{Aut}(X)|} = \frac{R}{d} = \gamma.$$

□

As a corollary to the above and Lemma 6.11.1, we have the following:

Corollary 6.13.2. *If X is a vertex transitive graph, then*

$$\xi_f(X) \leq \frac{|V(X)|}{\alpha_q(X)}.$$

□

Projective rank also relates well to some previously known parameters. An orthogonal representation of a graph is an assignment of nonzero vectors to its vertices such that adjacent vertices receive orthogonal vectors. The minimum dimension in which there exists an orthogonal representation of a graph X is known as its *orthogonal rank*, and is denoted by $\xi(X)$. This parameter is investigated in various papers including [8], in which they relate it to the “rank-1 quantum chromatic number”. Clearly, a projective representation using rank 1 projectors is equivalent to an orthogonal representation, and one can think of projective rank as a fractional version of orthogonal rank. Due to this, we have that

$$\xi_f(X) \leq \xi(X)$$

for all graphs X .

One can also think of projective rank as a subspace version of fractional chromatic number. A homomorphism from X to $K_{d,r}$ can be transformed into a d/r -projective representation by simply mapping an r -subset, S , of $[d]$ to the projection onto $\text{span}\{e_i \in \mathbb{C}^d : i \in S\}$, where e_i is the i^{th} standard basis vector. Since the fractional chromatic number of a graph X is the minimum value of d/r such that $X \rightarrow K_{d,r}$, we see that

$$\xi_f(X) \leq \chi_f(X)$$

for all graphs X .

Slightly less trivially, we have the following Lemma which is joint work with Laura Mančinska:

Lemma 6.13.3. *For any graph X ,*

$$\xi_f(X) \leq \chi_q(X).$$

Proof. Suppose that X has a quantum n -coloring, i.e. a quantum homomorphism to K_n . Since K_n is vertex transitive, this implies that there exist rank r projectors, $E_{x_i} \in \mathbb{C}^{nr \times nr}$, which give a quantum homomorphism from X to K_n . From the orthogonality requirements on the E_{x_i} , we see that $E_{x_1}E_{x'_1} = 0$ whenever $x \sim x'$. Therefore the projectors E_{x_1} for $x \in V(X)$ give a projective representation of X of value $\frac{nr}{r} = n$. Therefore $\xi_f(X) \leq n$. \square

The above lemma is also implied by the fact that projective rank is quantum homomorphism monotone, which we will prove in the next section.

6.14 Relation to $\bar{\vartheta}^+$ and Monotonicity of Projective Rank

In the previous section we saw that $\xi_f(X) \leq \xi(X)$ and $\xi_f(X) \leq \chi_f(X)$ for all graphs X . It was shown in [34] that these inequalities also hold when ξ_f is replaced by $\bar{\vartheta}$. So it is natural to consider whether ξ_f and $\bar{\vartheta}$ compare in some nice way. We will show that they do, and in fact that $\bar{\vartheta}^+(X) \leq \xi_f(X)$ for all graphs X . To do so, we first note that, as with projective packings, we do not change the value of projective rank by restricting to real matrices, since we can use the same transformation given in Lemma 6.2.2.

Lemma 6.14.1. *For any graph X ,*

$$\bar{\vartheta}^+(X) \leq \xi_f(X).$$

Proof. Suppose that $x \mapsto E_x$ for $x \in V(X)$ is a d/r -representation of X using real matrices. For each $x \in V(X)$, let

$$v_x = \frac{1}{\sqrt{r}} \text{vec}(E_x) - \frac{\sqrt{r}}{d} \text{vec}(I),$$

where I is the $d \times d$ identity. Note that these are not unit vectors. We will show that, after normalization, these vectors give a rigid vector $\frac{d}{r}$ -coloring. We will not normalize until the end, since it keeps the computations a little cleaner.

Recall from the definition of rigid vector colorings, we must show that the inner product of two vectors assigned to adjacent vertices is $-\frac{1}{d/r-1}$, and that any two vectors assigned

to the vertices of X have inner product at least this value. For arbitrary $x, x' \in V(X)$,

$$\begin{aligned}
v_x^T v_{x'} &= \left(\frac{1}{\sqrt{r}} \text{vec}(E_x) - \frac{\sqrt{r}}{d} \text{vec}(I) \right)^T \left(\frac{1}{\sqrt{r}} \text{vec}(E_{x'}) - \frac{\sqrt{r}}{d} \text{vec}(I) \right) \\
&= \frac{1}{r} \text{Tr}(E_x E_{x'}) - \frac{1}{d} \text{Tr}(E_x) - \frac{1}{d} \text{Tr}(E_{x'}) + \frac{r}{d^2} \text{Tr}(I) \\
&= \frac{1}{r} \text{Tr}(E_x E_{x'}) - \frac{r}{d} - \frac{r}{d} + \frac{r}{d} \\
&= \frac{1}{r} \text{Tr}(E_x E_{x'}) - \frac{r}{d}.
\end{aligned}$$

If $x \sim x'$, then $\text{Tr}(E_x E_{x'}) = 0$ and so $v_x^T v_{x'} = -r/d$ in this case. Since the E_x 's are positive semidefinite, $v_x^T v_{x'}$ is at least $-r/d$ for any $x, x' \in V(X)$. If $x = x'$, then

$$\text{Tr}(E_x E_{x'}) = \text{Tr}(E_x^2) = \text{Tr}(E_x) = r,$$

and so $\|v_x\|^2 = 1 - r/d = (d - r)/d$. Therefore, after we normalize, the inner product of any two vectors assigned to adjacent vertices of X will be

$$-\frac{r}{d} \cdot \frac{d}{d - r} = -\frac{r}{d - r} = -\frac{1}{d/r - 1}.$$

Furthermore, the inner product of any two vectors assigned to the vertices of X is at least this value. \square

Since $X \xrightarrow{q} Y$ for vertex transitive Y if and only if $X \times Y$ has a projective representation of value $|V(Y)|$, we obtain a corollary to the above lemma similar to Corollary 6.8.3 to Lemma 6.8.2.

Corollary 6.14.2. *If $X \xrightarrow{q} Y$, then*

$$\bar{\vartheta}^+(X \times Y) = |V(Y)|.$$

\square

The reader may have noticed that the projective rank of a graph can be naturally defined in terms of homomorphisms. Indeed, if $\Omega(r, d)$ is the infinite graph whose vertices are the rank r projectors in $\mathbb{C}^{d \times d}$ such that orthogonal projectors are adjacent, then X has a d/r -representation if and only if $X \rightarrow \Omega(r, d)$. Thinking of things in this way makes it easy to see that ξ_f is homomorphism monotone, however we can actually show that it is in fact quantum homomorphism monotone.

Theorem 6.14.3. *If $X \xrightarrow{q} Y$, and Y has a projective representation of value $\gamma \in \mathbb{Q}$, then X has a projective representation of value γ . Therefore, $\xi_f(X) \leq \xi_f(Y)$.*

Proof. Suppose that $X \xrightarrow{q} Y$, and that $E_{xy} \in \mathbb{C}^{d' \times d'}$ for $x \in V(X)$, $y \in V(Y)$ are the projectors which give the quantum homomorphism. Furthermore, let $y \mapsto F_y$ for $y \in V(Y)$ be a d/r -representation of Y . Define P_x for $x \in V(X)$ as follows:

$$P_x = \sum_{y \in V(Y)} E_{xy} \otimes F_y.$$

We will show that this is a projective representation of value $\frac{d}{r}$ for X . First we must check that the P_x are projectors. Since each term in the above sum is a projector, it suffices to show that the terms are pairwise orthogonal. For distinct $y, y' \in V(Y)$, $E_{xy}E_{xy'} = 0$, and therefore the P_x are projectors.

Next we must check that $P_x P_{x'} = 0$ for $x \sim x'$. We have that

$$P_x P_{x'} = \left(\sum_{y \in V(Y)} E_{xy} \otimes F_y \right) \left(\sum_{y' \in V(Y)} E_{x'y'} \otimes F_{y'} \right) = \sum_{y, y' \in V(Y)} E_{xy} E_{x'y'} \otimes F_y F_{y'}.$$

However, if $x \sim x'$, then $E_{xy}E_{x'y'} = 0$ whenever $y \not\sim y'$. Furthermore, if $y \sim y'$, then $F_y F_{y'} = 0$. Therefore, all of the terms in the final sum above are 0 when $x \sim x'$, and thus $P_x P_{x'} = 0$ in this case.

Now all that is left to do is check that this representation has value $\frac{d}{r}$. The dimension of this representation is simply dd' since the E_{xy} are in dimension d' and the F_y are in dimension d . We noted above that the terms in the sum defining P_x are pairwise orthogonal, therefore

$$\begin{aligned} \text{rk}(P_x) &= \text{rk} \left(\sum_{y \in V(Y)} E_{xy} \otimes F_y \right) = \sum_{y \in V(Y)} \text{rk}(E_{xy} \otimes F_y) \\ &= \sum_{y \in V(Y)} \text{rk}(E_{xy}) \text{rk}(F_y) = r \sum_{y \in V(Y)} \text{rk}(E_{xy}) = rd'. \end{aligned}$$

Therefore, the representation has value $\frac{dd'}{rd'} = \frac{d}{r}$ and thus $\xi_f(X) \leq \frac{d}{r}$. □

The contrapositive of the above theorem can be used to show that a quantum homomorphism does not exist between certain graphs, and in the next section we will see an example of this for which the monotonicity of $\bar{\vartheta}^-$, $\bar{\vartheta}$, and $\bar{\vartheta}^+$ is not able to accomplish the same.

6.15 Projective Rank of Some Special Graphs

We have already noted that the projective rank of the complete graph K_n is n , and this is easy to see. In this section we will show that we can give the exact value of the projective rank for all Kneser graphs and odd cycles as well.

Recall that $\xi_f(X) \leq \chi_f(X)$ for all graphs X . For $n \geq 2r$, the Kneser graph $K_{n:r}$ has $\chi_f(K_{n:r}) = n/r$. Therefore we immediately see that $\xi_f(K_{n:r}) \leq n/r$. On the other hand, it was shown in [34], that $\bar{\vartheta}(K_{n:r}) = n/r$ whenever $n \geq 2r$. Applying Lemma 6.14.1, we have the following:

Lemma 6.15.1. *If $n \geq 2r$, then*

$$\xi_f(K_{n:r}) = \frac{n}{r}.$$

Note that if $n < 2r$, then $K_{n:r}$ is empty and thus $\xi_f(K_{n:r}) = 1$ in this case.

For the odd cycles we are not so lucky. It was shown in [34] that

$$\bar{\vartheta}(C_{2k+1}) = \frac{1 + \cos\left(\frac{\pi}{2k+1}\right)}{\cos\left(\frac{\pi}{2k+1}\right)}$$

for all $k \in \mathbb{N}$. A strict vector coloring of this value can be achieved by assigning the vertices of C_{2k+1} to the vertices of a regular $(2k+1)$ -gon of radius 1 centered at the origin such that adjacent vertices of C_{2k+1} are assigned to vertices of the $(2k+1)$ -gon that are as far apart as possible. This is easily seen to be a rigid vector coloring as well, and thus $\bar{\vartheta}^+(C_{2k+1}) = \bar{\vartheta}(C_{2k+1})$ for all $k \in \mathbb{N}$. We will see below that $\xi_f(C_{2k+1}) = \frac{2k+1}{k}$ for all $k \in \mathbb{N}$, and since

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

for $k \geq 2$, we can not use $\bar{\vartheta}^+$ to tightly bound ξ_f from below, as we did with the Kneser graphs.

On the other hand, the fractional chromatic number does serve as a tight upper bound on the projective rank of C_{2k+1} . Indeed, it is easily seen that the vertices

$$\{i, i+1, \dots, i+k-1\} \in V(K_{2k+1:k})$$

for $i \in [2k+1]$ induce an odd cycle of length $2k+1$, and thus $C_{2k+1} \rightarrow K_{2k+1:k}$ implying that $\chi_f(C_{2k+1}) \leq \frac{2k+1}{k}$.

In order to prove the necessary lower bound on $\xi_f(C_{2k+1})$, we will need the following well-known lemma, which we state without proof.

Lemma 6.15.2. *If U and W are subspaces of a vector space V , then*

$$\dim(U) + \dim(W) = \dim(U \oplus W) + \dim(U \cap W).$$

□

We are now able to give the projective rank of all odd cycles. This is joint work with Laura Mančinska.

Theorem 6.15.3. *For $k \in \mathbb{N}$,*

$$\xi_f(C_{2k+1}) = \frac{2k+1}{k}.$$

Proof. Let $[2k+1]$ be the vertex set of C_{2k+1} such that $i \sim i+1$ modulo $2k+1$ for all $i \in [2k+1]$. Suppose that there exists a d/r -representation of C_{2k+1} . If we replace the projectors assigned to each vertex of C_{2k+1} with the subspaces of \mathbb{C}^d they project onto, then we will have an assignment of r -dimensional subspaces of \mathbb{C}^d to the vertices of C_{2k+1} such that adjacent vertices receive orthogonal subspaces. Let V_i be the subspace assigned to vertex $i \in [2k+1]$. We will prove that

$$\dim(V_1 \cap V_{2j+1}) \geq (2j+1)r - jd,$$

for all $j \in \{0, 1, \dots, k\}$ by induction. Clearly the statement is true for $j = 0$. Suppose that it holds for $j = \ell - 1$ for some $\ell \leq k$. The subspaces $(V_1 \cap V_{2\ell-1})$ and $V_{2\ell+1}$ are contained in the orthogonal complement of $V_{2\ell}$ which has dimension $d - r$. By Lemma 6.15.2, we have

$$d - r \geq \dim(V_{2\ell+1} \oplus (V_1 \cap V_{2\ell-1})) = \dim(V_{2\ell+1}) + \dim(V_1 \cap V_{2\ell-1}) - \dim(V_{2\ell+1} \cap (V_1 \cap V_{2\ell-1})).$$

Ignoring the middle expression and rearranging, we get

$$\begin{aligned} \dim(V_{2\ell+1} \cap (V_1 \cap V_{2\ell-1})) &\geq (r - d) + \dim(V_{2\ell+1}) + \dim(V_1 \cap V_{2\ell-1}) \\ &\geq 2r - d + (2\ell - 1)r - (\ell - 1)d \\ &= (2\ell + 1)r - \ell d \end{aligned}$$

by induction. Since $\dim(V_1 \cap V_{2\ell+1}) \geq \dim(V_{2\ell+1} \cap (V_1 \cap V_{2\ell-1}))$, we have proven the claim. Letting $j = k$ we see that

$$\dim(V_1 \cap V_{2k+1}) \geq (2k+1)r - kd.$$

However, since $1 \sim 2k+1$ we must have that $\dim(V_1 \cap V_{2k+1}) = 0$. Therefore,

$$(2k+1)r - kd \leq 0,$$

and rearranging gives

$$\frac{d}{r} \geq \frac{2k+1}{k}.$$

□

We mentioned above that $C_{2k+1} \rightarrow K_{2k+1:k}$ and thus $\chi_f(C_{2k+1}) \leq \frac{2k+1}{k}$. This in fact holds with equality which can be seen by noting that C_{2k+1} has an independent set of size k , or by the above theorem. This in particular implies that the odd girth of $K_{2k+1:k}$ is $2k+1$, since $C_{2k'+1} \rightarrow K_{2k+1:k}$ implies that

$$\frac{2k'+1}{k'} \leq \frac{2k+1}{k} \Rightarrow k' \geq k$$

by the homomorphism monotonicity of χ_f . Moreover, it has been shown [40] that the odd girth of $K_{n:r}$ is $2\lceil \frac{r}{n-2r} \rceil + 1$ for $n \geq 2r$. It is not hard to see that $\lceil \frac{r}{n-2r} \rceil$ is the smallest value of k such that

$$\frac{2k+1}{k} \leq \frac{n}{r}. \quad (6.8)$$

In other words, $C_{2k+1} \rightarrow K_{n:r}$ if and only if (6.8) holds. This implies that $C_{2k+1} \xrightarrow{q} K_{n:r}$ if (6.8) holds. On the other hand, Lemma 6.15.1 and Theorem 6.15.3 along with the monotonicity of ξ_f with respect to quantum homomorphisms implies the converse of this statement. Therefore $C_{2k+1} \xrightarrow{q} K_{n:r}$ if and only if (6.8) holds and we have the following:

Theorem 6.15.4. *For $n \geq 2r$,*

$$\text{og}_q(K_{n:r}) = \text{og}(K_{n:r}) = 2 \left\lceil \frac{r}{n-2r} \right\rceil + 1.$$

□

It is important to note that the above theorem cannot be proved using $\bar{\vartheta}^-$, $\bar{\vartheta}$, or $\bar{\vartheta}^+$. We mentioned above that the values of both $\bar{\vartheta}(C_{2k-1})$, and $\bar{\vartheta}^+(C_{2k-1})$ are

$$\frac{1 + \cos\left(\frac{\pi}{2k-1}\right)}{\cos\left(\frac{\pi}{2k-1}\right)}.$$

It can be shown [18] that $\bar{\vartheta}^-(X) = \bar{\vartheta}(X)$ for any arc transitive graph X , and thus the above is also the value of $\bar{\vartheta}^-(C_{2k-1})$. However,

$$\bar{\vartheta}^-(K_{2k+1:k}) = \bar{\vartheta}(K_{2k+1:k}) = \bar{\vartheta}^+(K_{2k+1:k}) = \frac{2k+1}{k},$$

and it is not hard to show that

$$\frac{1 + \cos\left(\frac{\pi}{2k-1}\right)}{\cos\left(\frac{\pi}{2k-1}\right)} \leq \frac{2k+1}{k}$$

for $k \geq 3$. Therefore $\bar{\vartheta}(C_{2k-1}) \leq \bar{\vartheta}(K_{2k+1:k})$ (and similarly for $\bar{\vartheta}^-$ and $\bar{\vartheta}^+$), and so we cannot conclude that C_{2k-1} does not have a quantum homomorphism to $K_{2k+1:k}$ using this approach.

6.16 Lattice Properties of Quantum Homomorphisms

In this section and the next, we will shift focus from graph parameters to the more order theoretic properties of quantum homomorphisms. We will not formally define the quantum homomorphism order until the next section, but here we will lay the ground work for showing that, like the homomorphism order, it is a lattice. To do this we will simply need to show that the categorical product and disjoint union of graphs satisfies the same order theoretic properties for quantum homomorphisms as they do for homomorphisms. More specifically, we have the following two lemmas.

Lemma 6.16.1. *A graph Z satisfies $Z \xrightarrow{q} X$ and $Z \xrightarrow{q} Y$ if and only if $Z \xrightarrow{q} X \times Y$.*

Proof. Suppose that $Z \xrightarrow{q} X \times Y$. Since $X \times Y \rightarrow X$ and $X \times Y \rightarrow Y$, we have that $Z \xrightarrow{q} X$ and $Z \xrightarrow{q} Y$. Now suppose that projectors E_{zx} for $z \in V(Z)$, $x \in V(X)$ and F_{zy} for $z \in V(Z)$, $y \in V(Y)$ give quantum homomorphisms from Z to X and from Z to Y respectively. For $z \in V(Z)$ and $(x, y) \in V(X \times Y)$, let

$$P_{z(x,y)} = E_{zx} \otimes F_{zy}.$$

We have that

$$\sum_{(x,y) \in V(X \times Y)} P_{z(x,y)} = \sum_{x \in V(X)} E_{zx} \otimes \sum_{y \in V(Y)} F_{zy} = I \otimes I = I.$$

If $z \sim z'$ and $(x, y) \not\sim (x', y')$, then without loss of generality we have that $x \not\sim x'$ and thus $E_{zx}E_{zx'} = 0$. Therefore,

$$P_{z(x,y)}P_{z'(x',y')} = (E_{zx}E_{zx'}) \otimes (F_{zy}F_{zy'}) = 0$$

in this case. Therefore $Z \xrightarrow{q} X \times Y$. □

The lemma will allow us to show that any two elements of the quantum homomorphism order have a meet. The next lemma gives the corresponding result for joins.

Lemma 6.16.2. *A graph Z satisfies $X \xrightarrow{q} Z$ and $Y \xrightarrow{q} Z$ if and only if $X \cup Y \xrightarrow{q} Z$.*

Proof. If $X \cup Y \xrightarrow{q} Z$, then since $X \rightarrow X \cup Y$ and $Y \rightarrow X \cup Y$, we have that $X \xrightarrow{q} Z$ and $Y \xrightarrow{q} Z$. Conversely, suppose that $E_{xz} \in \mathbb{C}^{d \times d}$ for $x \in V(X)$, $z \in V(Z)$ and $F_{yz} \in \mathbb{C}^{d' \times d'}$ for $y \in V(Y)$, $z \in V(Z)$ give quantum homomorphisms from X to Z and from Y to Z respectively. Let I_d and $I_{d'}$ be the $d \times d$ and $d' \times d'$ identity matrices respectively. For $w \in V(X \cup Y)$ and $z \in V(Z)$, let

$$P_{wz} = \begin{cases} E_{wz} \otimes I_{d'} & \text{if } w \in V(X) \\ F_{wz} \otimes I_d & \text{if } w \in V(Y) \end{cases}$$

For $w \in V(X)$, we have that

$$\sum_{z \in V(Z)} P_{wz} = \sum_{z \in V(Z)} E_{wz} \otimes I_{d'} = I_d \otimes I_{d'} = I_{dd'},$$

and similarly for when $w \in V(Y)$. If $w \sim w'$ and $z \not\sim z'$, then without loss of generality $w, w' \in V(X)$. Therefore,

$$P_{wz} P_{w'z'} = (E_{wz} \otimes I_{d'}) (E_{w'z'} \otimes I_d) = (E_{wz} E_{w'z'}) \otimes I_d = 0.$$

Thus we have shown that $X \cup Y \xrightarrow{q} Z$. □

Together, these two lemmas allow us to show that the quantum homomorphism order is a lattice with the same meet and join operations as the homomorphism order.

6.17 The Quantum Homomorphism Order and Measurement Graphs

As with the relation “ \rightarrow ”, we can use the relation “ \xrightarrow{q} ” to construct a partial order. We say that graphs X and Y are *quantum homomorphically equivalent*, or simply *q-equivalent* for short, if $X \xrightarrow{q} Y$ and $Y \xrightarrow{q} X$. We will denote this by $X \equiv_q Y$. Again, similarly to with homomorphic equivalence, the relation “ \equiv_q ” is an equivalence relation which gives rise to *quantum homomorphic equivalence classes*, or simply *q-equivalence classes*. We will use the notation $\mathcal{H}_q(X)$ to refer to the q-equivalence class to which X belongs. The *quantum homomorphism order*, denoted \mathcal{G}_q , is a partial order on q-equivalence classes such that $\mathcal{H}_q(X) \leq_q \mathcal{H}_q(Y)$ if $X \xrightarrow{q} Y$. As with the homomorphism order, we may abuse notation

and simply write $X \leq_q Y$. By Lemmas 6.16.1 and 6.16.2 in the previous section, two q -equivalence classes, $\mathcal{H}_q(X)$ and $\mathcal{H}_q(Y)$, have $\mathcal{H}_q(X \times Y)$ as their meet and $\mathcal{H}_q(X \cup Y)$ as their join. Therefore \mathcal{G}_q is a lattice, and it has the same meet and join operations as \mathcal{G} .

Perhaps the first question to ask is what the class $\mathcal{H}_q(X)$ “looks like” in terms of homomorphic equivalence classes. Since $X \rightarrow Y$ implies $X \xrightarrow{q} Y$, the q -equivalence classes are unions of homomorphic equivalence classes. Furthermore, if $\mathcal{H}(X) \leq \mathcal{H}(Y)$, then $\mathcal{H}_q(X) \leq_q \mathcal{H}_q(Y)$. This means that if Ψ is the map which takes $\mathcal{H}(X)$ to $\mathcal{H}_q(X)$, then Ψ is a homomorphism of partial orders from \mathcal{G} to \mathcal{G}_q . Therefore, \mathcal{G}_q is a homomorphic image of \mathcal{G} . Moreover, since

$$\Psi(\mathcal{H}(X) \wedge \mathcal{H}(Y)) = \Psi(\mathcal{H}(X \times Y)) = \mathcal{H}_q(X \times Y) = \mathcal{H}_q(X) \wedge \mathcal{H}_q(Y)$$

and similarly $\Psi(\mathcal{H}(X) \vee \mathcal{H}(Y)) = \mathcal{H}_q(X) \vee \mathcal{H}_q(Y)$, we have that Ψ is a lattice homomorphism. It is not clear if any of these observations imply anything of importance, but we find them quite interesting.

Since q -equivalence classes are unions of homomorphic equivalence classes, it is natural to ask “how many” homomorphic equivalence classes form a given q -equivalence class. In particular, is it possible that $\mathcal{H}_q(X) = \mathcal{H}(X)$? This is true for $X \equiv K_2$ or $X \equiv K_1$, since $X \xrightarrow{q} K_2$ if and only if $X \rightarrow K_2$, and similarly for K_1 . Are these the only such examples of this? We believe so, and we have the following conjecture:

Conjecture 6.17.1. *If $K_2 < X$, then $\mathcal{H}_q(X) \neq \mathcal{H}(X)$.*

One approach to proving the above conjecture is to consider the measurement graphs. In Section 6.3, we showed that $X \xrightarrow{q} Y$ if and only if $X \rightarrow M(Y, d)$ for some $d \in \mathbb{N}$. If there exists a graph X such that $X \not\rightarrow Y$ but $X \rightarrow M(Y, d)$ for some $d \in \mathbb{N}$, then $X \cup Y \xrightarrow{q} Y$ and $Y \xrightarrow{q} X \cup Y$. Therefore, $X \cup Y \equiv_q Y$, but $X \cup Y \not\rightarrow Y$ and thus $X \cup Y \neq Y$. This implies that $\mathcal{H}_q(Y) \neq \mathcal{H}(Y)$. Furthermore, for any graph X satisfying this property, the image of X in $M(Y, d)$ also satisfies this property, and therefore it suffices to find a finite subgraph of $M(Y, d)$ which does not have a homomorphism to Y to obtain $\mathcal{H}_q(Y) \neq \mathcal{H}(Y)$. The existence of a finite subgraph of $M(Y, d)$ which does not admit a homomorphism to Y is equivalent to $M(Y, d) \not\rightarrow Y$ (this can be proved by a trivial modification of Gottschalk’s proof of the De Bruijn-Erdős theorem [21]). Therefore, this approach for showing $\mathcal{H}_q(Y) \neq \mathcal{H}(Y)$ can succeed if and only if $\mathfrak{M}(Y) \not\rightarrow Y$. So we have the following strengthening of the above conjecture:

Conjecture 6.17.2. *If $K_2 < X$, then $\mathfrak{M}(X) \not\rightarrow X$.*

On the other hand, this is not necessary for $\mathcal{H}_q(Y) \neq \mathcal{H}(Y)$. If there exists a graph X such that $X \rightarrow Y$, $Y \not\rightarrow X$, and $Y \rightarrow M(X, d)$ for some $d \in \mathbb{N}$, then $X \equiv_q Y$ and thus $\mathcal{H}_q(Y) \neq \mathcal{H}(Y)$ even if $\mathfrak{M}(Y) \rightarrow Y$. Conversely, suppose that $\mathfrak{M}(Y) \rightarrow Y$ and $\mathcal{H}_q(Y) \neq \mathcal{H}(Y)$. Then there exists a graph X such that $X \equiv_q Y$ but $X \not\equiv Y$. This implies that $X \rightarrow \mathfrak{M}(Y)$ by Theorem 6.3.2 and thus $X \rightarrow Y$. Since $X \not\equiv Y$, we also have that $Y \not\rightarrow X$. So we see that, $X < Y$, and furthermore this is true of any element of $\mathcal{H}_q(Y)$ not homomorphically equivalent to Y . Therefore, if $\mathfrak{M}(Y) \rightarrow Y$ then Y is *the* maximal (with respect to \leq) element of $\mathcal{H}_q(Y)$ up to homomorphic equivalence.

If $\mathcal{H}_q(Y) \neq \mathcal{H}(Y)$ for some graph Y , then we can ask how much “bigger” $\mathcal{H}_q(Y)$ is in comparison to $\mathcal{H}(Y)$. In this case, there exists some graph $X \in \mathcal{H}_q(Y)$ such that $X \not\equiv Y$. Therefore, either $X \not\rightarrow Y$ or $Y \not\rightarrow X$. Without loss of generality we can assume that $Y \not\rightarrow X$. Therefore, $X \cup Y \not\rightarrow X$ and $X \rightarrow X \cup Y$. This implies that $X < Y$. Since \mathcal{G} is dense, this shows that for any graph Y , either $\mathcal{H}_q(Y) = \mathcal{H}(Y)$, or $\mathcal{H}_q(Y)$ is the union of an infinite number of homomorphic equivalence classes.

For any graph Y , we know that $\mathcal{H}_q(Y)$ must contain every finite graph X such that $Y \rightarrow X \rightarrow \mathfrak{M}(Y)$. But it may contain others as noted above. We have in fact seen an example of a graph Y for which $\mathcal{H}_q(Y)$ does not contain only graphs X such that $Y \rightarrow X \rightarrow \mathfrak{M}(Y)$. The graph Ω_n has a quantum homomorphism to K_n for all n , and if n is a power of two, then $K_n \rightarrow \Omega_n$ since an n -clique in Ω_n corresponds to a Hadamard matrix of order n , which always exists for n a power of two. Therefore, if n is a power of two, then $\Omega_n \equiv_q K_n$. Furthermore, if n is large enough, then $\Omega_n \not\rightarrow K_n$ and thus K_n does not lie between Ω_n and $\mathfrak{M}(\Omega_n)$ in the homomorphism order.

Since $X \rightarrow \mathfrak{M}(Y)$ for any graph $X \in \mathcal{H}_q(Y)$, the graph $\mathfrak{M}(Y)$ serves as a sort of “upper bound” for the graphs in $\mathcal{H}_q(Y)$, but of course $\mathfrak{M}(Y)$ is not finite and therefore not contained in $\mathcal{H}_q(Y)$. We noted above that if $\mathfrak{M}(Y) \rightarrow Y$ then Y is the maximal element of $\mathcal{H}_q(Y)$. Conversely, suppose that Y is such that $X \rightarrow Y$ for all $X \in \mathcal{H}_q(Y)$. If Z is a finite subgraph of $\mathfrak{M}(Y)$, then $Z \xrightarrow{q} Y$ and therefore $Z \cup Y \equiv_q Y$. This implies that $Z \cup Y \rightarrow Y$ and thus $Z \rightarrow Y$. Therefore, every finite subgraph of $\mathfrak{M}(Y)$ admits a homomorphism to Y , and this implies that $\mathfrak{M}(Y) \rightarrow Y$ by the above mentioned modification of the De Bruijn-Erdős theorem. So a graph Y is the maximal element of $\mathcal{H}_q(Y)$ with respect to \leq if and only if $\mathfrak{M}(Y) \rightarrow Y$.

Having considered maximal elements of q -equivalence classes, it is natural to consider minimal elements as well. Suppose X is a graph such that $X \rightarrow Y$ for all $Y \in \mathcal{H}_q(X)$. It is easy to see that any other graph of $\mathcal{H}_q(X)$ with this property must be homomorphically equivalent to X , thus any q -equivalence class contains at most one such graph up to homomorphic equivalence. Furthermore, if X is such a graph, then every graph $Y \in \mathcal{H}_q(X)$

satisfies $X \rightarrow Y \rightarrow \mathfrak{M}(X)$, and so $\mathcal{H}_q(X)$ simply consists of all finite graphs Y such that $X \rightarrow Y \rightarrow \mathfrak{M}(X)$. The converse also trivially holds.

If X is a minimal element of $\mathcal{H}_q(X)$ with respect to \leq , then $Y \not\rightarrow X$ for any graph $Y \in \mathcal{H}_q(X)$ that is not homomorphically equivalent to X . It turns out that this condition is also sufficient for X to be a minimal element of $\mathcal{H}_q(X)$. To see this suppose that no element of $\mathcal{H}_q(X)$ which is not homomorphically equivalent to X admits a homomorphism to X . For contradiction suppose that there exists a graph $Z \in \mathcal{H}_q(X)$ such that $X \not\rightarrow Z$. Since both X and Z are contained in $\mathcal{H}_q(X)$, so is the graph $X \times Z$. However, $X \times Z \rightarrow X$ by Lemma 6.16.1, but $X \not\rightarrow X \times Z$ since $X \not\rightarrow Z$. Therefore $X \times Z$ is an element of $\mathcal{H}_q(X)$ which admits a homomorphism to X but is not homomorphically equivalent to X , a contradiction.

To finish this section we give three straightforward but interesting lemmas which will solidify some of the intuition of the above discussion.

Lemma 6.17.3. *For $d \in \mathbb{N}$,*

$$M(X, d) \xrightarrow{q} X,$$

and therefore $M(X, d) \equiv_q X$.

Proof. This actually immediately follows from Theorem 6.3.1, since $M(X, d) \rightarrow M(X, d)$. However, we will also give an explicit quantum homomorphism. Recall that a vertex of $M(X, d)$ is of the form $\mathbf{E} = (E_x)_{x \in V(X)}$ such that the E_x 's are projectors in $\mathbb{C}^{d \times d}$ which sum to identity. We claim that the projectors

$$P_{\mathbf{E}_x} = E_x$$

give a quantum homomorphism from $M(X, d)$ to X . Clearly,

$$\sum_{x \in V(X)} P_{\mathbf{E}_x} = \sum_{x \in V(X)} E_x = I.$$

Now suppose that $\mathbf{E}^1 = (E_x^1)_{x \in V(X)}$ and $\mathbf{E}^2 = (E_x^2)_{x \in V(X)}$ are adjacent vertices of $M(X, d)$. By the definition of $M(X, d)$, we have that $E_x^1 E_{x'}^2 = 0$ whenever $x \not\sim x'$. Therefore, if $x \not\sim x'$, then

$$P_{\mathbf{E}^1 x} P_{\mathbf{E}^2 x'} = E_x^1 E_{x'}^2 = 0.$$

Thus the $P_{\mathbf{E}_x}$ do indeed give a quantum homomorphism from $M(X, d)$ to X . \square

Note that the above does not imply that $\mathfrak{M}(X) \xrightarrow{q} X$. This is because the definition of quantum homomorphism requires that all projectors have the same dimension. However,

it does imply that the disjoint union of a finite number of $M(X, d)$ does have a quantum homomorphism to X , since we can take tensor products with appropriately sized identity matrices to obtain projectors of all the same dimension. It also implies the following:

Lemma 6.17.4. *If $X \xrightarrow{q} Y$, then $M(X, d) \xrightarrow{q} Y$ for all $d \in \mathbb{N}$.*

Proof. This follows from the above lemma and the transitivity of quantum homomorphisms. \square

The last lemma we give here shows that the existence of a quantum homomorphism between two finite graph is equivalent to the existence of a homomorphism between two measurement graphs.

Lemma 6.17.5. *For X and Y being finite graphs, $X \xrightarrow{q} Y$ if and only if $\mathfrak{M}(X) \rightarrow \mathfrak{M}(Y)$.*

Proof. Since X is a subgraph of $\mathfrak{M}(X)$, if $\mathfrak{M}(X) \rightarrow \mathfrak{M}(Y)$ then $X \rightarrow \mathfrak{M}(Y)$ and thus $X \xrightarrow{q} Y$ by Theorem 6.3.2. To prove the converse we will show that if $X \xrightarrow{q} Y$, then for any $d \in \mathbb{N}$ there exists $d' \in \mathbb{N}$ such that $M(X, d) \rightarrow M(Y, d')$. Since $X \xrightarrow{q} Y$, we have that $M(X, d) \xrightarrow{q} Y$ by the above lemma. Therefore, by Theorem 6.3.1, there exists $d' \in \mathbb{N}$ such that $M(X, d) \rightarrow M(Y, d')$. So $M(X, d) \rightarrow \mathfrak{M}(Y)$ for all $d \in \mathbb{N}$ and thus $\mathfrak{M}(X) \rightarrow \mathfrak{M}(Y)$. \square

Note that the above proof does not work for infinite graphs since Theorem 6.3.2 does not hold for infinite graphs in general.

One interesting consequence of the above lemma is that the quantum homomorphism order of finite graphs is isomorphic to the suborder of the homomorphism order of infinite graphs induced by the measurement graphs. It would be nice to be able to say the same thing for the quantum homomorphism order of infinite graphs, but we believe a slightly different definition of quantum homomorphism may be needed. One possibility is to attempt to come up with a definition which uses infinite dimensional projectors, but it is unclear whether such a definition would correspond to a quantum strategy (using infinite dimensional measurements) for the (X, Y) -homomorphism game.

6.18 Discussion of Results

One of the goals of this chapter was to compare homomorphisms to quantum homomorphisms. Towards that, we have seen several similarities as well as some differences between these two notions. We saw that both the relations \rightarrow and \xrightarrow{q} were transitive and gave rise

to partial orders which are also lattices. Furthermore, we saw that there exists a lattice homomorphism from the homomorphism order to the quantum homomorphism order. We also showed that the existence of a quantum homomorphism from X to Y is equivalent to the existence of a homomorphism from $\mathfrak{M}(X)$ to $\mathfrak{M}(Y)$. We introduced the homomorphism game which can be used as a tool for studying quantum homomorphisms, and we saw that $X \rightarrow Y$ and $X \xrightarrow{q} Y$ are equivalent to the existence of classical and quantum strategies to the (X, Y) -homomorphism game respectively. Perhaps most importantly we have seen that certain theorems which hold for homomorphisms have natural quantum analogs which hold for quantum homomorphisms. In particular, we have seen that $X \xrightarrow{q} Y$ if and only if $\alpha_q(X \times Y) = |V(X)|$ and that a quantum version of the clique-coclique bound holds. We have seen that one can use quantum homomorphisms to naturally define quantum analogs of many graph parameters. These quantum parameters are easily seen to be quantum homomorphism monotone for the same reason their classical counterparts are homomorphism monotone. Somewhat surprisingly, we found $\bar{\vartheta}^-$, $\bar{\vartheta}$, and $\bar{\vartheta}^+$ are quantum homomorphism monotone, even though they are defined in terms of homomorphisms. We also introduced the parameters $\tilde{\omega}$ and ξ_f which we showed to be quantum homomorphism monotone.

One notable difference between homomorphisms and quantum homomorphisms is that the obvious adaptation of the no-homomorphism lemma to quantum homomorphisms does not hold. However, by Lemma 6.13.1 and Theorem 6.14.3, if Y is vertex transitive and $X \xrightarrow{q} Y$, then

$$\frac{\tilde{\alpha}(X)}{|V(X)|} \geq \frac{1}{\xi_f(X)} \geq \frac{1}{\xi_f(Y)} = \frac{\tilde{\alpha}(Y)}{|V(Y)|}.$$

So if we like, we can take this to be our quantum version of the no-homomorphism lemma.

The graphs Ω_{4n} which had large separations between chromatic and quantum chromatic numbers serve as an example of how greatly the behavior of quantum homomorphisms can deviate from that of homomorphisms. From a quantum computing/information point of view these examples, which have a so-called “quantum advantage”, demonstrate the power that entanglement can have.

We discussed very many graph parameters in this chapter, and so as a summary of our results we have the following inequality in which we write $f \leq g$ if f and g are graph parameters such that $f(X) \leq g(X)$ for all X :

$$\omega \leq \omega_q \leq \tilde{\omega} \leq \bar{\vartheta}^- \leq \bar{\vartheta} \leq \bar{\vartheta}^+ \leq \xi_f \leq \chi_q \leq \chi.$$

Furthermore, we have shown that all of the parameters listed above other than ω and χ are quantum homomorphism monotone. Also, for each of the above inequalities, there exists a graph for which it is strict. We give examples for each below.

$\omega(\mathbf{X}) < \omega_q(\mathbf{X})$: In Section 6.11 we noted that Corollary 6.11.5 can be used to show that $\alpha(\Omega_n \square K_n) < \alpha_q(\Omega_n \square K_n)$ for certain values of n . Taking complements we obtain $\omega(\overline{\Omega_n \square K_n}) < \omega_q(\overline{\Omega_n \square K_n})$.

$\omega_q(\mathbf{X}) < \tilde{\omega}(\mathbf{X})$: Since ω_q must be an integer, any graph X such that $\tilde{\omega}(X)$ is not an integer must satisfy $\omega_q(X) < \tilde{\omega}(X)$. Recall that for $4|n$, we saw that $\xi_f(\Omega_n) = n$. Therefore,

$$\tilde{\omega}(\overline{\Omega_n}) = \tilde{\alpha}(\Omega_n) = \frac{|V(\Omega_n)|}{\xi_f(\Omega_n)} = \frac{2^n}{n}$$

since Ω_n is vertex transitive. So when $4|n$ and n is not a power of two, we have that $\omega_q(\Omega_n) < \tilde{\omega}(\Omega_n)$.

$\tilde{\omega}(\mathbf{X}) < \bar{\vartheta}^-(\mathbf{X})$: In Section 6.15 we mentioned that $\bar{\vartheta}^-(C_5) = \sqrt{5}$. However, Theorem 6.15.3 shows that $\xi_f(C_5) = 5/2$ and thus

$$\tilde{\omega}(C_5) = \tilde{\alpha}(\overline{C_5}) = \tilde{\alpha}(C_5) = \frac{|V(C_5)|}{\xi_f(C_5)} = \frac{5}{5/2} = 2 < \sqrt{5}.$$

Therefore $\tilde{\omega}(C_5) < \bar{\vartheta}^-(C_5)$.

$\bar{\vartheta}^-(\mathbf{X}) < \bar{\vartheta}(\mathbf{X})$: In [47], Schrijver gave an example of a graph \mathfrak{S} such that $\vartheta^-(\mathfrak{S}) < \vartheta(\mathfrak{S})$. Taking complements yields $\bar{\vartheta}^-(\overline{\mathfrak{S}}) < \bar{\vartheta}(\overline{\mathfrak{S}})$. (The graph \mathfrak{S} consists of the 01-strings of length six such that strings at Hamming distance at most three are adjacent.)

$\bar{\vartheta}(\mathbf{X}) < \bar{\vartheta}^+(\mathbf{X})$: In [48], Szegedy showed that $\bar{\vartheta}^-(Y)\bar{\vartheta}^+(\overline{Y}) = |V(Y)|$ for all vertex transitive graphs Y . Also, Lovász showed in [34] that $\bar{\vartheta}(Y)\bar{\vartheta}(\overline{Y}) = |V(Y)|$ for all vertex transitive graphs Y . The graph \mathfrak{S} given by Schrijver mentioned above is vertex transitive and therefore

$$\bar{\vartheta}^+(\mathfrak{S}) = \frac{|V(\mathfrak{S})|}{\bar{\vartheta}^-(\overline{\mathfrak{S}})} > \frac{|V(\mathfrak{S})|}{\bar{\vartheta}(\overline{\mathfrak{S}})} = \bar{\vartheta}(\mathfrak{S}).$$

$\bar{\vartheta}^+(\mathbf{X}) < \xi_f(\mathbf{X})$: We mentioned in Section 6.15 that $\bar{\vartheta}^+(C_5) = \sqrt{5}$, and Theorem 6.15.3 shows that $\xi_f(C_5) = 5/2$. Therefore $\bar{\vartheta}^+(C_5) < \xi_f(C_5)$.

$\xi_f(\mathbf{X}) < \chi_q(\mathbf{X})$: Since $\chi_q(X)$ must be an integer, $\xi_f(C_5) < \chi_q(C_5)$.

$\chi_q(\mathbf{X}) < \chi(\mathbf{X})$: We have already seen that $\chi_q(\Omega_n) < \chi(\Omega_n)$ for certain values of n .

We also briefly discussed the parameters χ_f and ξ . We showed that $\xi_f \leq \xi, \chi_f$, and it is well known that $\xi, \chi_f \leq \chi$. But how do ξ and χ_f compare to each other?

$\xi(\mathbf{X}) < \chi_f(\mathbf{X})$: A result of Frankl and Rödl [15] gives an upper bound on $\alpha(\Omega_{4n})$ which shows that $\chi_f(\Omega_{4n})$ is exponential in n . However, $\xi(\Omega_{4n}) = 4n$ by definition of Ω_{4n} .

$\chi_f(\mathbf{X}) < \xi(\mathbf{X})$: Since $\xi_f \leq \xi \leq \chi$ and ξ is integer valued, we have that $\xi(C_5) = 3 > 5/2 = \chi_f(C_5)$.

Therefore the parameters ξ and χ_f are incomparable. This also implies that neither of them is always equal to ξ_f nor always equal to χ . The two examples above also show that χ_f and χ_q are incomparable. The only remaining unknown relationship is that of χ_q and ξ . We do not currently have an example of a graph for which we know that these two parameters differ. However we believe that they are incomparable.

6.19 Open Questions

The most practical open question is on the decidability of the existence of quantum homomorphisms. Given graphs X and Y , is there an algorithm which decides if $X \xrightarrow{q} Y$ in finite time? It was noted in [8] that if we fix the dimension in which we are looking for projectors that give a quantum homomorphism from X to Y , then the question becomes decidable since it reduces to determining if a particular set of quadratic equations in a finite number of variables has a solution. So one approach to resolving the decidability of this question would be to bound the dimension required for a quantum homomorphism between two graphs.

We have seen separations between chromatic and quantum chromatic number, as well as the classical and quantum versions of independence and clique numbers. Can we find separations between other parameters and their quantum analogs? For instance, is there a graph X for which $\text{og}(X) > \text{og}_q(X)$? Such separations are of interest to the quantum information community, because they exhibit the advantage quantum entanglement can sometimes provide. Moreover, they provide a quantitative measure of this advantage.

More generally, can we find more examples of graphs X and Y such that $X \xrightarrow{q} Y$ but $X \not\rightarrow Y$? All of the examples we have seen have at least one of X and Y being complete. Can we find an example in which neither X nor Y is a complete graph? This question can be answered in the affirmative, but possibly in an uninteresting way. Recall that the graph Ω_n satisfies $\chi_q(\Omega_n) < \chi_f(\Omega_n)$ for certain values of n . This implies that $\Omega_n \xrightarrow{q} K_{nr:r}$ but $\Omega_n \not\rightarrow K_{nr:r}$ for these values of n and any value of r . However, $\Omega_n \xrightarrow{q} K_n \rightarrow K_{nr:r}$ and $\Omega_n \not\rightarrow K_n$, and so this example does not seem fundamentally different than when one of the graphs is complete. So we have the following revised question: Can we find an example of graphs X and Y such that $X \xrightarrow{q} Y$, $X \not\rightarrow Y$, and there does not exist $n \in \mathbb{N}$ such that $X \xrightarrow{q} K_n \xrightarrow{q} Y$ and either $X \not\rightarrow K_n$ or $K_n \not\rightarrow Y$? We point out that an example of a graph

Y such that $\text{og}(Y) > \text{og}_q(Y) \geq 5$ would give us such an example in addition to providing a separation for odd girth.

It would be nice to know that the supremum and infimum in the definitions of $\tilde{\alpha}$ and ξ_f can be replaced by maximum and minimum respectively. In particular this would allow us to say that $X \xrightarrow{q} Y$ if and only if $\tilde{\alpha}(X \times Y) = |V(X)|$, which is more aesthetically pleasing than the current theorem. However, it is not clear how one would prove that the supremum or infimum is always obtained. We do not even know that these parameters are always rationally valued. One possible approach would be to first attempt to find another proof that χ_f is always attained which uses the definition of χ_f involving homomorphisms to Kneser graphs, and then to attempt to adapt this proof to ξ_f . The current proof that χ_f is always attained relies on the fact that χ_f can be written as a linear program, and therefore always attains its optimal value.

In Section 6.15, we saw that the quantum odd girth of a Kneser graph is equal to its odd girth. Is the quantum chromatic number equal to chromatic number for Kneser graphs? This question has implications for the order \mathcal{G}_q , since the fact that there exists Kneser graphs with arbitrary large odd girth and chromatic number allows one to construct infinite antichains (sets of pairwise incomparable graphs) in \mathcal{G} . Therefore, if $\chi_q(K_{n:r}) = \chi(K_{n:r})$, then \mathcal{G}_q also contains infinite antichains. Infinite antichains are also part of at least one proof of the density of \mathcal{G} , so answering this question could possibly help to prove that \mathcal{G}_q is dense.

What other properties of \mathcal{G} does \mathcal{G}_q share? We have seen that \mathcal{G}_q is a lattice, but this is really just scratching the surface of the types of questions we can ask about this order. We mentioned the question of density above, but another interesting question is whether or not \mathcal{G}_q is universal. An affirmative answer to this would serve as evidence for \mathcal{G}_q being an interesting and robust order worthy of study. On the other hand, a negative answer would be a stark illustration of the effect entanglement has on homomorphisms, which should be of interest to those in the field of quantum information. A graph with a separation between quantum and classical versions of some parameter offers concrete quantitative measure of the advantage of entanglement, but a difference in the structure of \mathcal{G}_q from \mathcal{G} allows one to see the qualitative effect of entanglement on a global scale.

We saw some examples of theorems previously known for homomorphisms which turned out to be true for quantum homomorphisms as well. For instance, a connected graph X has a quantum homomorphism to a graph Y if and only if it has a quantum homomorphism to one of the components of Y . Are there other theorems concerning homomorphisms that can be adapted into true theorems for quantum homomorphisms? Furthermore, what is it that makes one theorem adaptable to quantum homomorphisms but another one not?

This is not a very precise question, but we think it is an interesting one.

The parameters $\bar{\vartheta}^-$, $\bar{\vartheta}$, $\bar{\vartheta}^+$, and ξ_f can all be defined naturally in terms of homomorphisms. However, they are not just homomorphism monotone, which is expected given these definitions, they are also quantum homomorphism monotone. This is in contrast with other homomorphism monotone parameters such as χ or ω . What essential quality of these parameters is the cause of this?

Can we find a formulation of $\tilde{\omega}$ in terms of homomorphisms? More specifically, is there a class of graphs $\{X_s : s \in S\}$, and function $f : S \rightarrow \mathbb{R}$ such that

$$\tilde{\omega}(X) = \sup\{f(s) : X_s \rightarrow X\}?$$

Clique number can be defined in this way, and since we view $\tilde{\omega}$ as a type of clique number, it would be nice to be able to give a similar definition for this parameter.

Lastly, we are interested in whether or not one can define a “quantum core” of a graph in a reasonable way. Cores are in some sense the canonical representatives of homomorphic equivalence classes, and so this question is asking whether or not there are canonical representatives of q-equivalence classes. Obviously, if $\mathcal{H}_q(Y) = \mathcal{H}(Y)$, then it makes sense to define the quantum core of Y as being simply the core of Y , but what about when $\mathcal{H}_q(Y) \neq \mathcal{H}(Y)$? There are several equivalent definitions of the core of a graph, and so we can borrow inspiration from some of them. For instance, we could consider defining the quantum core of Y to be the vertex minimal graph contained in $\mathcal{H}_q(Y)$. But unlike in the case of cores, we do not know if this is guaranteed to be unique. Further consideration of this question tends to lead to more questions, such as what is a proper quantum endomorphism, or what is a quantum isomorphism?

References

- [1] Peyman Afshani, Mahsa Ghandehari, Mahya Ghandehari, Hamed Hatami, Ruzbeh Tusserkani, and Xuding Zhu. Circular chromatic index of graphs of maximum degree 3. *Journal of Graph Theory*, 49(4):325–335, 2005. URL: <http://dx.doi.org/10.1002/jgt.20086>, doi:10.1002/jgt.20086.
- [2] Michael O. Albertson and Karen L. Collins. Homomorphisms of 3-chromatic graphs. *Discrete Math.*, 54(2):127–132, 1985. doi:10.1016/0012-365X(85)90073-1.
- [3] David Avis, Jun Hasegawa, Yosuke Kikuchi, and Yuuya Sasaki. A quantum protocol to win the graph colouring game on all Hadamard graphs. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.*, E89-A(5):1378–1381, 2006. arXiv:quant-ph/0509047v4, doi:10.1093/ietfec/e89-a.5.1378.
- [4] S. Beigi and P. W. Shor. On the complexity of computing zero-error and Holevo capacity of quantum channels. *ArXiv e-prints*, September 2007. arXiv:0709.2090.
- [5] Gilles Brassard, Richard Cleve, and Alain Tapp. Cost of exactly simulating quantum entanglement with classical communication. *Phys. Rev. Lett.*, 83:1874–1877, 1999. arXiv:quant-ph/9901035, doi:10.1103/PhysRevLett.83.1874.
- [6] Harry Buhrman, Richard Cleve, and Avi Wigderson. Quantum vs. classical communication and computation. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, STOC '98, pages 63–68. ACM, 1998. arXiv:quant-ph/9802040, doi:10.1145/276698.276713.
- [7] Peter J. Cameron and Priscila A. Kazanidis. Cores of symmetric graphs. *J. Aust. Math. Soc.*, 85(2):145–154, 2008. URL: <http://dx.doi.org/10.1017/S1446788708000815>, doi:10.1017/S1446788708000815.

- [8] Peter J. Cameron, Ashley Montanaro, Michael W. Newman, Simone Severini, and Andreas Winter. On the quantum chromatic number of a graph. *Electr. J. Comb.*, 14(1), 2007. [arXiv:quant-ph/0608016](https://arxiv.org/abs/quant-ph/0608016).
- [9] Richard Cleve, Peter Hoyer, Ben Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *19th IEEE Annual Conference on Computational Complexity*, pages 236–249, 2004. [arXiv:quant-ph/0404076v2](https://arxiv.org/abs/quant-ph/0404076v2), [doi:10.1109/CCC.2004.1313847](https://doi.org/10.1109/CCC.2004.1313847).
- [10] Bruno Codenotti, Ivan Gerace, and Sebastiano Vigna. Hardness results and spectral techniques for combinatorial problems on circulant graphs. *Linear Algebra Appl.*, 285:3–123, 1998.
- [11] P. Delsarte. *An Algebraic Approach to the Association Schemes of Coding Theory*. Philips journal of research / Supplement. N.V. Philips’ Gloeilampenfabrieken, 1973. URL: <http://books.google.ca/books?id=zna0SgAACAAJ>.
- [12] Runyao Duan, Simone Severini, and Andreas Winter. Zero-error communication via quantum channels and a quantum Lovász ϑ function. In *Proc. Int. Symp. Inf. Theory*, pages 64–68, 2011. [arXiv:1002.2514](https://arxiv.org/abs/1002.2514), [doi:10.1109/ISIT.2011.6034211](https://doi.org/10.1109/ISIT.2011.6034211).
- [13] P. Erdős. Graph theory and probability. *Canad. J. Math.*, 11:34–38, 1959.
- [14] R. Fraïssé. *Theory of relations*, volume 118 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, 1986.
- [15] Peter Frankl and Vojtěch Rödl. Forbidden intersections. *Trans. Amer. Math. Soc.*, 300(1):259–286, 1987. [doi:10.2307/2000598](https://doi.org/10.2307/2000598).
- [16] Junya Fukawa, Hiroshi Imai, and François Le Gall. Quantum coloring games via symmetric SAT games. In *Asian Conference on Quantum Information Science (AQIS’11)*, 2011.
- [17] Viktor Galliard and Stefan Wolf. Pseudo-telepathy, entanglement, and graph colorings. In *IEEE International Symposium on Information Theory*, page 101, 2002. [doi:10.1109/ISIT.2002.1023373](https://doi.org/10.1109/ISIT.2002.1023373).
- [18] Chris Godsil. Interesting graphs and their colourings. Unpublished notes, 2003.
- [19] Chris Godsil and Gordon Royle. *Algebraic Graph Theory*, volume 207 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2001. [doi:10.1007/978-1-4613-0163-9](https://doi.org/10.1007/978-1-4613-0163-9).

- [20] Chris Godsil and Gordon F. Royle. Cores of geometric graphs. *Ann. Comb.*, 15(2):267–276, 2011. URL: <http://dx.doi.org/10.1007/s00026-011-0094-5>, doi:10.1007/s00026-011-0094-5.
- [21] W. H. Gottschalk. Choice functions and Tychonoff’s theorem. *Proc. Amer. Math. Soc.*, 2:172, 1951.
- [22] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981. URL: <http://dx.doi.org.proxy.lib.uwaterloo.ca/10.1007/BF02579273>, doi:10.1007/BF02579273.
- [23] Geňa Hahn and Claude Tardif. Graph homomorphisms: structure and symmetry. In *Graph symmetry*, volume 497 of *NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci.*, pages 107–166. Kluwer Acad. Publ., 1997.
- [24] Stephen Travis Hedetniemi. *Homomorphisms of Graphs and Automata*. ProQuest LLC, Ann Arbor, MI, 1966. Thesis (Ph.D.)–University of Michigan. URL: http://gateway.proquest.com.proxy.lib.uwaterloo.ca/openurl?url_ver=Z39.88-2004&rft_val_fmt=info:ofi/fmt:kev:mtx:dissertation&res_dat=xri:pqdiss&rft_dat=xri:pqdiss:6701754.
- [25] Pavol Hell and Jaroslav Nešetřil. The core of a graph. *Discrete Math.*, 109(1-3):117–126, 1992. Algebraic graph theory (Leibnitz, 1989). URL: [http://dx.doi.org.proxy.lib.uwaterloo.ca/10.1016/0012-365X\(92\)90282-K](http://dx.doi.org.proxy.lib.uwaterloo.ca/10.1016/0012-365X(92)90282-K), doi:10.1016/0012-365X(92)90282-K.
- [26] Pavol Hell and Jaroslav Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.
- [27] Alan J. Hoffman. Some recent results on spectral properties of graphs. In *Beiträge zur Graphentheorie (Kolloquium, Manebach, 1967)*, pages 75–80. Teubner, Leipzig, 1968.
- [28] Ian Holyer. The NP-completeness of edge-coloring. *SIAM J. Comput.*, 10(4):718–720, 1981. URL: <http://dx.doi.org/10.1137/0210055>, doi:10.1137/0210055.
- [29] Jan Hubička and Jaroslav Nešetřil. Universal partial order represented by means of oriented trees and other simple graphs. *European J. Combin.*, 26(5):765–778, 2005. URL: <http://dx.doi.org/10.1016/j.ejc.2004.01.008>, doi:10.1016/j.ejc.2004.01.008.

- [30] John B. Johnston. Universal infinite partially ordered sets. *Proc. Amer. Math. Soc.*, 7:507–514, 1956.
- [31] Bjarni Jónsson. Universal relational systems. *Math. Scand.*, 4:193–208, 1956.
- [32] David Karger, Rajeev Motwani, and Madhu Sudan. Approximate graph coloring by semidefinite programming. *J. ACM*, 45(2):246–265, 1998. [arXiv:cs/9812008](https://arxiv.org/abs/cs/9812008), [doi:10.1145/274787.274791](https://doi.org/10.1145/274787.274791).
- [33] Benoit Larose, François Laviolette, and Claude Tardif. On normal Cayley graphs and hom-idempotent graphs. *European Journal of Combinatorics*, 19(7):867 – 881, 1998. URL: <http://www.sciencedirect.com/science/article/pii/S0195669898902347>, [doi:10.1006/eujc.1998.0234](https://doi.org/10.1006/eujc.1998.0234).
- [34] László Lovász. On the Shannon capacity of a graph. *IEEE Trans. Inform. Theory*, 25(1):1–7, 1979. [doi:10.1109/TIT.1979.1055985](https://doi.org/10.1109/TIT.1979.1055985).
- [35] R. Lukoka and J. Mazk. Cubic graphs with given circular chromatic index. *SIAM Journal on Discrete Mathematics*, 24(3):1091–1103, 2010. URL: <http://epubs.siam.org/doi/abs/10.1137/090752316>, [arXiv:http://epubs.siam.org/doi/pdf/10.1137/090752316](http://epubs.siam.org/doi/pdf/10.1137/090752316), [doi:10.1137/090752316](https://doi.org/10.1137/090752316).
- [36] Laura Mančinska, Giannicola Scarpa, and Simone Severini. New separations in zero-error channel capacity through projective Kochen-Specker sets and quantum coloring. *IEEE Transactions on Information Theory*, 59(6):4025–4032, 2013. [doi:10.1109/TIT.2013.2248031](https://doi.org/10.1109/TIT.2013.2248031).
- [37] Robert J. McEliece, Eugene R. Rodemich, and Howard C. Rumsey, Jr. The Lovász bound and some generalizations. *J. Combin. Inform. System Sci.*, 3(3):134–152, 1978.
- [38] Jaroslav Nešetřil. Structure of graph homomorphisms. I. *Graph Theory Notes N. Y.*, 35:7–12, 1998. New York Graph Theory Day, 35 (1998).
- [39] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, 2000.
- [40] Svatopluk Poljak and Zsolt Tuza. Maximum bipartite subgraphs of Kneser graphs. *Graphs Combin.*, 3(2):191–199, 1987. URL: <http://dx.doi.org.proxy.lib.uwaterloo.ca/10.1007/BF01788540>, [doi:10.1007/BF01788540](https://doi.org/10.1007/BF01788540).

- [41] Aleš Pultr and Věra Trnková. *Combinatorial, Algebraic and Topological Representations of Groups, Semigroups and Categories*, volume 22 of *North-Holland Mathematical Library*. North-Holland Publishing Co., Amsterdam, 1980.
- [42] David Roberson and Laura Mančinska. Graph homomorphisms for quantum players. 2012. URL: <http://arxiv.org/abs/1212.1724>, [arXiv:1212.1724](https://arxiv.org/abs/1212.1724).
- [43] David E. Roberson. Cores of vertex transitive graphs. *Electr. J. Comb.*, 20(2), 2013. URL: <http://www.combinatorics.org/ojs/index.php/eljc/article/view/v20i2p45>.
- [44] Gert Sabidussi. Vertex-transitive graphs. *Monatsh. Math.*, 68:426–438, 1964.
- [45] Giannicola Scarpa and Simone Severini. Kochen-Specker sets and the rank-1 quantum chromatic number. *IEEE Trans. Inf. Theory*, 58(4):2524–2529, 2012. [arXiv:1106.0712](https://arxiv.org/abs/1106.0712), [doi:10.1109/TIT.2011.2178018](https://doi.org/10.1109/TIT.2011.2178018).
- [46] Edward R. Scheinerman and Daniel H. Ullman. *Fractional Graph Theory*. Dover Publications Inc., Mineola, NY, 2011. A rational approach to the theory of graphs, With a foreword by Claude Berge, Reprint of the 1997 original.
- [47] Alexander Schrijver. A comparison of the Delsarte and Lovász bounds. *IEEE Trans. Inform. Theory*, 25(4):425–429, 1979. URL: <http://dx.doi.org.proxy.lib.uwaterloo.ca/10.1109/TIT.1979.1056072>, [doi:10.1109/TIT.1979.1056072](https://doi.org/10.1109/TIT.1979.1056072).
- [48] Mario Szegedy. A note on the theta number of Lovász and the generalized Delsarte bound. In *FOCS*, pages 36–39, 1994.
- [49] R. Šámal. Cycle-continuous mappings – order structure. *ArXiv e-prints*, December 2012. [arXiv:1212.6909](https://arxiv.org/abs/1212.6909).
- [50] Emo Welzl. Symmetric graphs and interpretations. *Journal of Combinatorial Theory, Series B*, 37(3):235 – 244, 1984. URL: <http://www.sciencedirect.com/science/article/pii/009589568490056X>, [doi:10.1016/0095-8956\(84\)90056-X](https://doi.org/10.1016/0095-8956(84)90056-X).
- [51] Hassler Whitney. Congruent Graphs and the Connectivity of Graphs. *Amer. J. Math.*, 54(1):150–168, 1932. URL: <http://dx.doi.org.proxy.lib.uwaterloo.ca/10.2307/2371086>, [doi:10.2307/2371086](https://doi.org/10.2307/2371086).
- [52] Xuding Zhu. Circular colouring of graphs. Workshop on Graph Homomorphisms at the Fields Institute, 2011. URL: <http://www.fields.utoronto.ca/audio/11-12/homomorphisms/zhu/>.

- [53] Xuding Zhu. The fractional version of Hedetniemi's conjecture is true. *European J. Combin.*, 32(7):1168–1175, 2011. URL: <http://dx.doi.org.proxy.lib.uwaterloo.ca/10.1016/j.ejc.2011.03.004>, doi:10.1016/j.ejc.2011.03.004.

Index

- (X, Y) -homomorphism game, 105
- $M(Y, d)$, 103
- $X * Y$, 18
- $X + Y$, 21
- $X \boxtimes Y$, 18
- $X \square Y$, 18
- $X \cup Y$, 20
- $X \times Y$, 118
- $X \times Y$, 18
- $X(G, C)$, 25
- $X[Y]$, 18
- $\mathfrak{M}(Y)$, 104
- α , 9
- χ , 8
- χ_f , 10
- \equiv , 12
- $\mathcal{H}(X)$, 12
- Im , 6
- ω , 9
- \xrightarrow{q} , 98
- Tr_A , 89
- ξ , 135
- ξ_f , 133
- d/r -representation, 132
- n -clique, 9
- n -colorable, 8
- n -coloring, 8
- \mathcal{G} , 13
- arc, 22
- arc transitive, 22
- automorphism, 7
- Cartesian product, 18
- categorical product, 18
- Cayley graph, 25
- chromatic number, 8
- class I, 49
- class II, 49
- clique number, 9
- color class, 8
- coloring, 8
- connection set, 25
- core, 13
- core of X , 14
- dense, 33
- dense from above, 57
- density matrix, 84
- disjoint union of graphs, 20
- disjunctive product, 18
- edge transitive, 22
- edge union, 21
- endomorphism, 7
- entangled state, 89
- equitable partition, 39
- exponential graph, 34
- fibre, 6
- fractional chromatic number, 10
- gap, 53

homomorphic equivalence class, 12
 homomorphic product, 118
 homomorphically equivalent, 12
 homomorphism, 6
 homomorphism game, 105
 homomorphism monotone, 11
 homomorphism order of graphs, 13
 homomorphism order of line graphs, 51
 hull, 44

 image, 6
 independence number, 9
 independence ratio, 23
 independent set, 9
 isomorphism, 7

 join, 32

 Kneser graph, 11

 lattice, 32
 left stochastic matrix, 81
 line graph, 47

 matching, 49
 maximally entangled state, 89
 measurement graph, 103, 104
 meet, 31
 multiple, 19

 non-edge transitive, 22
 normal Cayley graph, 42

 odd girth, 9
 orthogonal projectors, 97
 orthogonal rank, 135

 partial order, 12
 partial trace, 89
 partially ordered set, 12

 perfect matching, 58
 poset, 12
 post-measurement state, 83
 projective measurement, 86
 projective packing, 119
 projective packing number, 119
 projective rank, 133
 projective representation, 132
 projector, 86
 proper endomorphism, 8
 pure states, 82

 q-equivalence class, 143
 q-equivalent, 143
 quantum homomorphic equivalence class, 143
 quantum homomorphically equivalent, 143
 quantum homomorphism, 98
 quantum homomorphism order, 143
 quotient matrix, 39

 reduced state, 89
 retract, 8
 retraction, 8

 Schmidt decomposition, 98
 stabilizer, 25
 stochastic vectors, 81
 strong product, 18
 suborder, 51
 superposition, 82

 total order, 53

 unitary matrix, 82
 universal, 35

 vector coloring, 115
 vector coloring, rigid, 115
 vector coloring, strict, 115
 vertex transitive, 21

weakly d -maximal, [65](#)