DIPLOMARBEIT

The Erdős-Ko-Rado Theorem in Graph Theory

ausgeführt am

Institut für
Diskrete Mathematik und Geometrie
TU Wien

unter der Anleitung von

Univ.Prof. Dipl.-Ing. Dr.techn. Michael Drmota

durch

Lucas Unterberger B.Sc.
Matrikelnummer: 01325438
Studienkennzahl: UE 066 394

Wien, am 11. Mai 2022

Michael Drmota

Lucas Unterberger

Abstract

The Erdős-Ko-Rado theorem is a well known theorem in the field of combinatorics. It was discovered in the 1960s [8] and from the 1970s to the 1980s many generalizations have been proven [6]. In essence it provides an upper bound for the size of certain intersecting subfamilies of the power set of n elements. Since sets are a very general construct, that are used by all mathematical fields, it is no surprise that the Erdős-Ko-Rado theorem has applications outside of the field it was discovered in. This thesis aims to investigate its consequences in graph theory.

In the first chapter we introduce the modern version of the Erdős-Ko-Rado theorem and provide a short proof, due to Gyula O. H. Katona [18]. We also refer to Deza and Frankl [6], who gathered further generalizations. Afterwards we recap basic definitions and theorems for graph theory. Since a big part of this work uses algebraic graph theory, we define algebraic constructs like groups and homomorphisms and their applications in graph theory.

The goal of the second chapter is to define Kneser graphs and provide an additional proof for the Erdős-Ko-Rado theorem. This is done by considering independent sets on Kneser graphs and subsequently prove that the independence number of Kneser graphs is bounded by the same bound as the Erdős-Ko-Rado theorem. The proof however is far more involved and makes use of multiple algebraic graph theoretic concepts.

Chapter 3 is about spectral graph theory. This mathematical field studies the spectrum of graphs and can be described as the union of graph theory and linear algebra. In particular we want to define eigenvalues for graphs and also characterize the spectrum of Kneser graphs. There is still a lot of discovery to be made in this field, however some graph invariants can already be bound by eigenvalues.

In chapter 4 we explore a different application of the Erdős-Ko-Rado theorem, by defining EKR graphs. These types of graphs obey a bound that can be thought of as a generalization of the Erdős-Ko-Rado theorem for more complicated structures. We prove the EKR property for some subfamilies of trees and show closure under lexicographic products with complete graphs. We end on a conjecture by Holroyd and Talbot [16] that ensures EKR properties for a big class of graphs, which still has to be proven or disproven. We also point towards further research that managed to prove the conjecture for certain graph classes.

Kurzfassung

Das Erdős-Ko-Rado Theorem ist ein bekannter Satz in der Kombinatorik. Es wurde erstmals in den 1960ern entdeckt [8] und in den 1970ern bis 1980ern wurden einige Verallgemeinerungen bewiesen [6]. Im wesentlichen garantiert es eine obere Schranke für die Mächtigkeit gewisser schneidender Teilfamilien der Potenzmenge von n Elementen. Da Mengen ein sehr allgemeines Konstrukt sind, das in praktisch allen mathematischen Teilgebieten verwendet wird, ist es keine Überraschung, dass das Erdős-Ko-Rado Theorem viele Anwendungen auch außerhalb der Mengentheorie hat. Das Ziel dieser Arbeit ist die Auswirkungen und Anwendungen auf die Graphentheorie zu untersuchen.

Im ersten Kapitel führen wir die moderne Version des Erdős-Ko-Rado Theorems ein und beweisen es mittels eines kurzen Beweises, von Gyula O. H. Katona [18]. Wir berufen uns außerdem auf Deza and Frankl [6], die weitere Verallgemeinerungen gesammelt haben. Danach wiederholen wir grundlegende Definitionen und Sätze der Graphentheorie. Da ein großer Teil dieser Arbeit Algebraische Graphentheorie benutzt, definieren wir außerdem algebraische Konstrukte wie Gruppen und Homomorphismen und deren Anwendungen in der Graphentheorie.

Das Ziel des zweiten Kapitels ist Kneser Graphen zu definieren und einen zweiten Beweis des Erdős-Ko-Rado Theorems auszuarbeiten. Dafür betrachten wir unabhängige Mengen in Kneser Graphen und beweisen, dass die Unabhängigkeitszahl von Kneser Graphen durch dieselbe obere Schranke beschränkt ist, die im Erdős-Ko-Rado Theorem vorkommt. Der Beweis ist allerdings signifikant aufwendiger und benutzt viele Konzepte aus der Algebraischen Graphentheorie.

Kapitel 3 handelt von Spektraler Graphentheorie. Dieses Gebiet erforscht das Spektrum von Graphen und kann als die Vereinigung von Graphentheorie und linearer Algebra gesehen werden. Wir definieren Eigenwerte von Graphen und charakterisieren das Spektrum von Kneser Graphen. In diesem Gebiet gibt es noch viele Forschungsmöglichkeiten, jedoch können einige Graphen Invarianten schon durch Eigenwerte beschränkt werden.

In Kapitel 4 betrachten wir eine andere Anwendung des Erdős-Ko-Rado Theorems, indem wir EKR Graphen definieren. Diese Graphen erfüllen eine Ungleichung, die als eine Art Verallgemeinerung des Erdős-Ko-Rado Theorems für komplexere Strukturen gesehen werden kann. Wir beweisen die EKR Eigenschaft für einige Teilfamilien von Bäumen und zeigen Abgeschlossenheit bezüglich des lexikographischen Produkts mit vollständigen Graphen. Wir enden mit einer Vermutung von Holroyd und Talbot [16], die die EKR Eigenschaft für eine große Klasse an Graphen zeigen würde. Beweis oder Widerlegung stehen noch aus. Schließlich verweisen wir auf laufende Forschung, die diese Vermutung schon für einige Klassen von Graphen beweisen konnte.

Acknowledgement

I want to thank my supervisor, Professor Dr. Michael Drmota, for introducing me to this interesting mathematical field and providing guidance. Above all, I want to thank him for his excellent communication and constant support throughout this journey.

I want to thank my family for believing in me and especially my brother, who always had time to ponder about mathematics with me.

Finally I want to thank my wonderful girlfriend, who supported me not only in my study, but also in everything else.

Eidesstattliche Erklärung

Ich erkläre an Eides statt, dass ich die vorliegende Diplomarbeit selbstständig und ohne fremde Hilfe verfasst, andere als die angegebenen Quellen und Hilfsmittel nicht benutzt bzw. die wörtlich oder sinngemäß entnommenen Stellen als solche kenntlich gemacht habe.

Wien, am 11. Mai 2022

Lucas Unterberger

Contents

1 Preliminaries

This chapter is supposed to provide the foundations and necessary definitions that are used in this thesis. We start with the Erdős-Ko-Rado theorem, followed by graph theory and algebraic concepts like groups and homomorphisms.

The Erdős-Ko-Rado theorem provides a bound for intersecting families. Our goal is to observe its applications in graph theory, since this field is closely related to combinatorial structures.

We start by stating some definitions and the modern version of the Erdős-Ko-Rado theorem. It is assumed that the reader is familiar with basic set theory, mathematical notation and proof concepts like induction.

1.1 The Erdős-Ko-Rado Theorem

This theorem was originally stated slightly different [8] and over time the lower bound for n has been improved. We will prove the theorem of Erdős, Ko and Rado, for intersecting families, using Gyula O. H. Katona's substantially shorter proof [18]. For l-intersecting families and other variations we will refer to [6].

We call a set $S \subseteq \{1, \dots, n\} =: \Omega$ with $|S| = k$ a k-set. A family of sets $\mathcal{S}$, where every two members have non empty intersection is called an *intersecting* family. Such a family $\mathcal{S}$ is called *t-intersecting* if for all $S_i, S_j \in \mathcal{S}$, with $S_i \neq S_j$, holds $|S_i \cap S_j| \geq t$. We denote the family of all k-subsets of Ω with $\binom{\Omega}{k}$. Further we define $\mathbb{N}^+ := \mathbb{N} \setminus \{0\}$.

Theorem 1.1 (Erdős-Ko-Rado)**.** *Let $n, k, l \in \mathbb{N}$ with $n > 0, k \leq \frac{n}{2}$ and let $\mathcal{F} \subseteq \binom{\Omega}{k}$. If $\mathcal{F}$ is intersecting, then*

$$|\mathcal{F}| \leq \binom{n-1}{k-1}. \tag{1}$$

Proof. The proof provided by Katona uses double counting. For $i \in \Omega$ let $B_i \subseteq \Omega$ be the set of numbers which are congruent to

$$(i-1)k+1, (i-1)k+2, \dots, ik \mod n.$$

These sets do not need to be pairwise different. We start by proving that for any intersecting subset $B_{i_1}, \dots, B_{i_d}$ of the set described above, $d \leq k$ holds. We may assume that $i_1 = 1$ by symmetry. For every $b \in B_1$ holds that b is congruent to one of $1, 2, \dots, k \mod n$, thus for $k \leq n/2$ we see that $B_1 = \{1, \dots, k\}$. Now for $j > 1$ we have $(j-1)k + 1 > k$. Thus $B_1 \cap B_j \neq \emptyset$ if and only if either

$$jk = q_1 n + r_1 \tag{2}$$

where $0 \leq q_1 < k$ and $1 \leq r_1 \leq k$ or

$$(j-1)k + 1 = q_2 n + r_2 \tag{3}$$

where $0 \leq q_2 < k$ and $1 \leq r_2 \leq k$ holds. Equation (2) represents the intersection of the last element of B_j with B_1 and equation (3) represents the intersection of the first element of B_j with B_1.
There is at most one pair (j, r_1) that satisfies equation (2) for fixed q_1 and at most one pair (j, r_2) satisfying equation (3) for fixed q_2. If equation (2) holds for some j with $r_1 = k$, then $(j-1)k + 1 = q_1 n + 1$ and equation (3) can only be satisfied with the same j. Lets suppose that (2) holds for some j_0, with $r_1 < k$. Then (3) holds with $q_2 = q_1$ only if $j = j_0 + 1$, since $j_0 k + 1 = q_1 n + r_1 + 1$. But $B_{j_0} \cap B_{j_0+1} = \emptyset$, since $2k \leq n$. Therefore for every q there is at most one j with $B_1 \cap B_j \neq \emptyset$, such that either (2) or (3) holds. Further if $q = 0$, then $j = 1$, which proves our first statement.
Let $F_1, \ldots, F_{n!}$ be the sequences obtained from $F_1 = (B_1, \ldots, B_n)$, by permuting the elements of Ω. We denote the elements of $\mathcal{F}$ with A_i for $i \in \{1, \ldots, |\mathcal{F}|\}$. Now we count the pairs (F_i, A_j), where A_j is contained in the sequence F_i, in two different ways. Our first statement guarantees that a fixed F_i can contain at most k different intersecting A_j. Therefore the amount of pairs has to be less than or equal to $n!k$.
On the other hand there are $k!(n-k)!$ ways to permute a fixed B_r into an A_j, since $|B_r| = k$, meaning that there are k possible ways to map an element of B_r to A_j, while the other $(n-k)$ elements must not be mapped onto A_j. Thus there are $nk!(n-k)!$ many F_i, that contain a fixed A_j, since $|F_i| = n$. Further there are at exactly $|\mathcal{F}|$ many A_j. Therefore we can conclude that

$$n!k \geq |\mathcal{F}| nk!(n-k)!,$$

which proves the stated bound of this theorem, since the above equation holds if and only if

$$|\mathcal{F}| \leq \frac{n!k}{nk!(n-k)!} = \frac{(n-1)!}{(k-1)!(n-k)!} = \binom{n-1}{k-1}.$$

□

A natural question would be to ask whether this bound is tight. Lets assume that $\mathcal{F} = \{F \in \binom{\Omega}{k} : x_0 \subseteq F\}$, meaning that every set in $\mathcal{F}$ has the element x_0 in common. Then any $F \in \mathcal{F}$ can consist of $k-1$ other elements that can be chosen out of $n-1$ elements. By construction, $\mathcal{F}$ is intersecting and we have

$$|\mathcal{F}| = \binom{n-1}{k-1}.$$

On the other hand if $k < \frac{n}{2}$, then the example above, which is called a trivial family, is the only family that fulfills equality of the Erdős-Ko-Rado bound. This was shown by Hilton and Miller [15], however Peter Frankl and Zoltán Füredi constructed a shorter and easier to read proof [11].

In [6] M. Deza and P. Frankl summarize improvements and generalizations of the EKR-theorem, which improves their bound for n. For l-intersecting families we get the following theorem.

Theorem 1.2 (Erdős-Ko-Rado)**.** *Let $n, k, l \in \mathbb{N}$ with $0 < l < k < n$ and let $\mathcal{F} \subseteq \binom{\Omega}{k}$. If $\mathcal{F}$ is l-intersecting and n is sufficiently larger than k and l, then*

$$|\mathcal{F}| \leq \binom{n-l}{k-l}. \tag{4}$$

If $k < \frac{n}{2}$, then $\mathcal{F} = \{F \in \binom{\Omega}{k} : x_0 \subseteq F\}$ for some $x_0 \in \Omega$ if and only if

$$|\mathcal{F}| = \binom{n-l}{k-l}. \tag{5}$$

1.2 Graphs

This section aims to introduce most definitions and corollaries needed for the field of algebraic graph theory. This includes basic graph theory, some group theory and homomorphisms on graph structures. We will introduce some well known graph classes like trees and planar graphs and their properties.

There are many books about graph theory, however we will focus on Godsil and Royle's Algebraic Graph Theory [12], due to their wide cover of this topic. This introductory chapter is based on chapter 1, 2 and 6 of [12].

A graph X consists of two sets $X = (V(X), E(X))$, where $V(X)$ is any nonempty set, called the *vertex* set and $E(X) \subseteq V(X) \times V(X)$ is called the *edge* set. We write $n = |V(X)|$ for the cardinality of the vertex set. If the context is clear we drop the dependencies and simply write $X = (V, E)$. Due to their simple structure, graphs are widely applicable. They are used in computer science, chemistry, communication networks and many more fields [1].
As a side note most authors denote a graph by $G = (V, E)$ [22] [3], but this thesis is in large parts based on Godsil and Royle's Algebraic Graph Theory [12]. Group theory plays a vital role in this subject and groups are usually denoted by G across mathematics. That is why we will use the slightly uncommon notation defined above in chapters where groups are considered.

1.2.1 Simple Graphs

We say that an edge $e \in E$ *connects* the vertices it is made of $e = (x, y)$ and we also write $e = xy$ for this edge. This edge is *incident* to both vertices x and y. An edge $e = (x, x)$ connecting a vertex to itself is called a *loop*. Two vertices are *adjacent*, or *neighboring*, if there exists an edge that connects them and we will denote this by $x \sim y$. A graph is called *undirected*, if $e = xy \in E$ implies $yx \in E$, otherwise it is called *directed*. An edge in an undirected graph is also called an *arc*, to further distinguish these two classes. A graph is called *finite* if V is finite.

Throughout this thesis, if not further mentioned, we will assume that X is a finite *simple graph*, meaning that X is undirected, finite and loop-free. Sometimes it is

useful to consider a multigraph, where E is a multiset, which allows multiple distinct edges connecting the same two vertices.

Two graphs $X_1 = (V_1, E_1)$ and $X_2 = (V_2, E_2)$ are *equal* if $V_1 = V_2$ and $E_1 = E_2$. However this kind of equality is quite restrictive, since graphs tend to be represented very differently even if they have the same structure. Therefore we define two graphs X_1 and X_2 to be isomorphic, if there exists a bijective function $f : V_1 \to V_2$ such that $\forall x, y \in V_1$ holds $x \sim y$ if and only if $f(x) \sim f(y)$ in V_2. We consider two isomorphic graphs to be equal, since they only differ in their vertex labels.

Figure 1: Two isomorphic graphs

As shown in Figure 1 two isomorphic graphs can also have different drawings. A *drawing* of a graph is usually an embedding of the graph onto the 2D plane, where vertices are represented by distinct dots or circles and edges by (not necessarily straight) lines. Depending on the context, the vertices or edges can be labeled by their respective elements. Directed graphs are drawn with arrows representing edges, with $e = (x, y)$ depicted by an arrow going from x to y. Note that in general you can draw graphs on any surface, but we will only consider drawings on the plane.

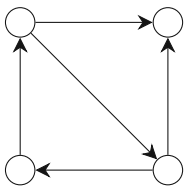

Figure 2: An example of a directed graph

1.2.2 Degrees

Let $X = (V, E)$ be a simple graph and $x \in V$, then the *neighborhood* of x is defined as the set of all neighboring vertices. It is denoted by $N(x)$. The *degree* or *valency* of a vertex x is the number of neighboring vertices $deg(x) = |N(x)|$. We also write $d(v)$ for the degree of vertex v. Note that for directed graphs there is a distinction between the number of incoming edges $d(x)_+$ and the number of outgoing edges $d(x)_-$. These

definitions allow the formulation of what many regard the simplest graph theoretical theorem.

Lemma 1.1 (Handshaking Lemma)**.** *Let $X = (V, E)$ be a graph. Then*

$$\sum_{v \in V} d(v) = 2|E|. \tag{6}$$

Proof. The degree of a vertex counts the number of all incident edges for that vertex. Every edge connects exactly 2 vertices and contributes to two degrees at a time. Summation over all degrees therefore counts all edges twice. □

A graph with at least one vertex, but no edges is called an *empty* graph E_n and a graph with every possible edge $E = V \times V$ is called a *complete* graph K_n, where $n = |V|$. A graph is called *regular* if $d(x_i) = d(x_j)$ for all $x_i, x_j \in V$.

1.2.3 Subgraphs

Let $X = (V(X), E(X))$ be a graph then $Y = (V(Y), E(Y))$ is a *subgraph* of X if Y is a graph, $V(Y) \subseteq V(X)$ and $E(Y) \subseteq E(X)$. We also write $Y \subseteq X$ and call X the *supergraph* of Y. Subgraphs are an important concept for graphs and yield further definitions. A subgraph $Y \subseteq X$ is called an *induced* subgraph if for every pair of vertices $x_1, x_2 \in V(Y)$ holds that $x_1x_2 \in E(X)$ implies $x_1x_2 \in E(Y)$. Induced subgraphs basically contain all possible edges of the supergraph. Therefore if $Y \subseteq X$ is an induced subgraph and $V(Y) = V(X)$ then $Y = X$.

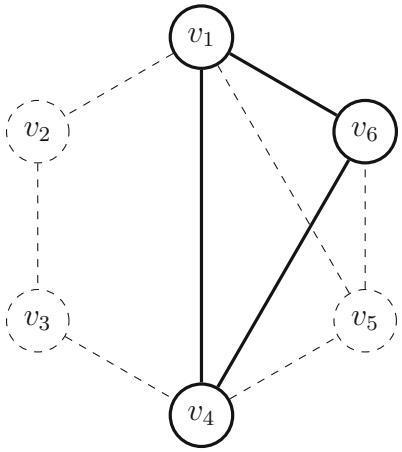

Figure 3: The subgraph induced by the vertices v_1, v_4 and v_6

Two important classes of subgraphs are cliques and independent sets. A *clique* is a subgraph that is complete and an *independant set* is an induced subgraph that is empty. In general a subset of vertices $I \subseteq V$ is called *independant* if $x, y \in I$ implies $x \not\sim y$. We write $\omega(X)$ for the size of the largest clique in a graph X. The size of the largest independent set in X is called the *independence number* $\alpha(X)$. Especially $\alpha(X)$ will remain a recurring figure throughout this thesis.

1.2.4 Connectivity

We will now define connectivity for graphs. Let $X = (V, E)$ be a graph, a *path* $x = x_0, x_1, x_2, \ldots, x_r = y$ of length r from x to y is a sequence of $r + 1$ distinct vertices, such that for all $i \in \{0, \ldots, r-1\}$ holds that $x_i x_{i+1}$ is an edge. This path connects the vertices x and y. It is also possible to define such a path as r distinct edges, such that an end-vertex of an edge, other than y, has to be the start-vertex of the next edge. The smallest number of edges in a path from vertices x to y is called the *distance* from x to y. A graph X is *connected* if there exists a path between any two vertices of X. Otherwise X is a disconnected graph.

Induced subgraphs that are maximal, subject to being connected, are called *connected components* of X. Every disconnected graph therefore decomposes into connected induced subgraphs and can be reconstructed by their disjoint union.

A *cycle* is a connected graph, where every vertex has exactly two neighbors. We write C_n for the cycle with n vertices. The smallest cycle is the complete graph K_3. Usually a cycle of a graph refers to a subgraph, where every vertex has degree two.

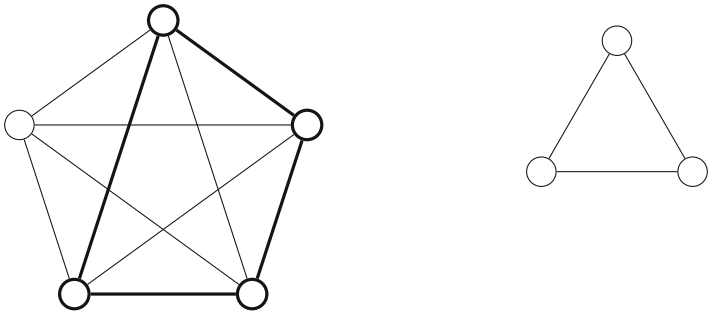

Figure 4: K_5 with C_4 as a subgraph and the cycle K_3

1.2.5 Trees

A *forest* is a graph without cycles and a *tree* is a connected forest. Thus all forests can be viewed as a union of trees. Trees are a very important graph class and they have many well-studied properties. It is easily provable that all paths in a tree are unique, because they have no cycles, by definition.

Lemma 1.2. *Let $T = (V, E)$ be a tree with $|V| = n$ then $|E| = n - 1$.*

Proof. We prove this by induction over n. Assume that the property holds for n and we have a tree with $n + 1$ vertices. Now take any edge e and remove it. Then we get two disconnected trees each with less than or equal to n vertices, since trees do not contain cycles. Suppose they have i and j vertices respectively and therefore $i - 1$ and $j - 1$ edges by the induction hypothesis. If we add those edges and the single edge we deleted we get $i - 1 + j - 1 + 1 = i + j - 1$ edges. Since $i + j = n + 1$, we have shown that the tree with $n + 1$ vertices has n edges. □

We call vertices of a tree with degree one *leafs* and vertices with degree two *internal* or *inner* vertices. A *rooted tree* is a tree with a designated vertex called the *root*. Often rooted trees are drawn with the root at the top layer, its neighbors in the next layer below and so on, until you end with the leafs. In a rooted tree the *parent* of a vertex is the unique neighbor that is closer to the root vertex. All other neighbors are the *children* of that vertex. The root has no parent and the leafs have no children.

A *k-ary tree* is a rooted tree, where each vertex has at most k children. An example can be seen in Figure 5.

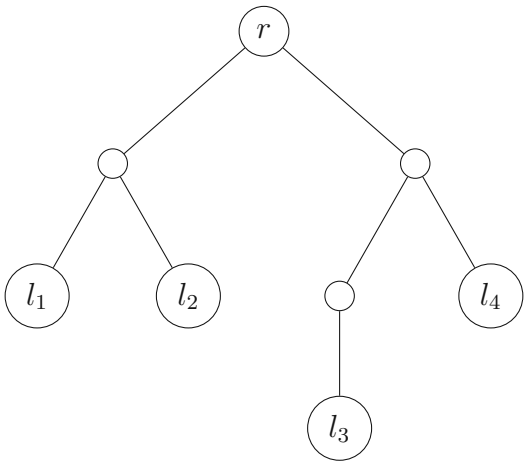

Figure 5: A *binary* tree with root r and leafs l_i

1.2.6 Planar Graphs

As we have already seen in Figure 1, two isomorphic graphs can always be drawn in different ways. We want to define a *planar* graph as a graph that can be drawn without crossing any of its edges. This formulation leaves some ambiguities however. Therefore we define a function that, given a graph, maps each vertex onto a distinct point in the plane and each edge onto a continuous non self-intersecting curve joining its endpoints. We call this function a *planar embedding* if the curves corresponding to incident edges meet only at their common vertex and all other edges do not meet. A graph is then called *planar* if and only if it has a planar embedding. A *plane graph* is a planar graph together with a fixed embedding.

Given a plane graph, its edges divide the plane into regions, called $faces$. There is exactly one unbounded face, called the *external face* and all other faces are bounded. Some examples of planar graphs we have already seen include all complete graphs K_i with $i < 5$, trees and cycles. We will quickly proof this for forests.

Lemma 1.3. *Let $X = (V, E)$ be a forest, then X is planar.*

Proof. First we consider that every forest can be written as a union of distinct trees. Further we can assume that each tree can be drawn arbitrarily far away from the other

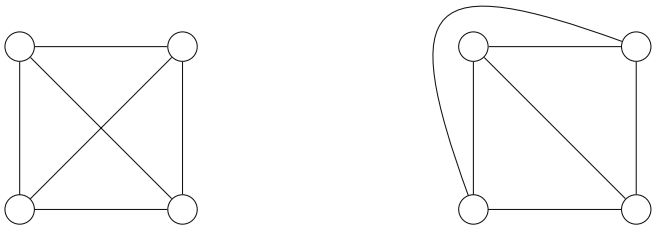

Figure 6: A planar and a non planar drawing of the K_4

trees, meaning we only need to prove that a single tree is planar. We choose any vertex of this tree and declare it to be a root. Now we start to draw the root on the plane and draw all of its children in the level below with sufficiently large distance. Let the distance between the first and the last child be one unit. For the next level, draw the children of the children of the root, with distance $\frac{1}{n}$ from first to last. Since there are at most $n-1$ children, there are no intersections of vertices or edges. If we repeat this process for the other levels, we obtain a planar drawing of our tree. □

There is a remarkable connection between vertices, edges and faces for planar graphs, found by Euler, which has since been proven in a number of different ways [7].

Theorem 1.3 (Euler)**.** *Let $X = (V, E)$ be a connected, planar graph and F the set of faces given by a planar embedding. Then*

$$|V| - |E| + |F| = 2. \tag{7}$$

Proof. We will proof this theorem by reducing X to a tree. We start by observing that a tree satisfies the above condition, since $n - (n-1) + 1 = 2$.
Let X be a planar graph. For every cycle in X we remove one edge of this cycle and iterate this process until X is a tree. In every step we lower the number of edges by one and lower the number of faces by one, canceling out the effect on the equation. Since a tree fulfills the equation, the planar graph had to fulfill it as well.

□

1.2.7 Matchings

A *matching* M in a graph X is a set of edges such that two edges do not share the same vertex. The size of a matching is the number of elements it contains. A vertex incident to an edge of M is *covered* by M. A matching that covers all vertices of X is called a $perfect\ matching$ or a $1\text{-}factor$. If a graph contains a perfect matching it has to have an even number of vertices. A *maximum matching* is a matching containing the largest number of edges.

1.3 Groups

Groups are a fundamental concept in mathematics and algebraic graph theory explores the consequences of applying group theory to graphs. We will introduce the basic def-

initions and use them later on to gain insights into transitive graphs.

A *group* is a tuple $(G, \star)$ of a set G coupled with a binary operation $\star : G \times G \to G$ that fulfills the following axioms:

$$\begin{aligned} &[Associativity]: && \forall a, b, c \in G : (a \star b) \star c = a \star (b \star c). \\ &[Identity\ Element]: && \forall a \in G\ \exists e \in G : e \star a = a. \\ &[Inverse\ Element]: && \forall a \in G\ \exists a^{-1} \in G : a^{-1} \star a = e. \end{aligned}$$

Typical notation for the *identity element* is e and the *inverse* of an element a is denoted by a^{-1}. We also write $f \star g$, or simply fg instead of $\star(f, g)$, similar to how we denote addition or multiplication. It is easy to show that the identity and inverse elements are unique and that they are left- and right-identity or left- and right-inverse elements respectively [14].

Let G be a group, then H is a *subgroup* of G if $H \subseteq G$ and H is a group with the same operation. The left cosets of H in G are defined as $gH := \{gh : h \in H\}$ for all $g \in G$. The right cosets are similarly defined as $Hg := \{hg : h \in H\}$ for all $g \in G$.

Lemma 1.4. *Any two left (or right) cosets are either disjoint or equal.*

Proof. Let g_1H and g_2H be two left cosets. Lets assume that $g_1H \cap g_2H \neq \emptyset$. Thus there exist $h_1, h_2 \in H$ such that $g_1h_1 = g_2h_2$. After multiplying with h_2^{-1} we get

$$g_1h_1h_2^{-1} = g_2h_2h_2^{-1} = g_2e = g_2.$$

Since H is a group $h_1h_2^{-1} \in H$ and we can find every element of g_2H in $g_1h_1h_2^{-1}H = g_1H$. The last equality holds, since if $h \in H$ then $h_2h_1^{-1}h$ is also in H. To prove the other inclusion, simply switch the roles of g_1 and g_2. This works similarly for right cosets. □

The *order* of a group G is defined as its cardinality, while the *order* of an element $g \in G$ is the least a such that $g^a = e$, where g^a denotes the multiplication of g with itself a times.

1.3.1 Permutation Groups

Let V be a set. Bijective functions that map V to V are called *permutations*. The *symmetric group* of a set V is the set of all permutations, with function composition $\circ$ for its group operation. The symmetric group is indeed a group, since the identity function is a permutation and inverse elements of bijective functions exist and are unique. We denote the symmetric group by $Sym(V)$, or $Sym(n)$, if $|V| = n$. A *permutation group* on V is a subgroup of $Sym(V)$. The image of an element $v \in V$ under a permutation $g \in Sym(V)$ is denoted by v^g. Let $S \subseteq V$, then we define the image of S under permutation g as $S^g = \{s^g \in S : s \in S\}$.

An isomorphism of a graph to itself is called an *automorphism*. The set of all automorphisms of X is a group called the *automorphism group*, denoted by $Aut(X)$. The

name is justified, because the identity function corresponds to the identity element of the group, for $f, g \in Aut(X)$ the composition $f \circ g$ is in $Aut(X)$ and since f is bijective f^{-1} exists and also maps to X.

The automorphism group of a graph X is a subgroup of $Sym(V(X))$. Thus $Aut(X)$ is a permutation group and can be interpreted as the group of label permutations of its vertices that satisfy all edge requirements. Automorphisms have the pleasant property, that they preserve distance, meaning that an automorphism f fulfills that the distance from x to y is the same as the distance from $f(x)$ to $f(y)$. This can be seen as the image of an automorphism has to map a path to a path with the same number of vertices, thus preserving distance.

A *permutation representation* of a group G is a homomorphism from G into $Sym(V)$ for some set V. It is also referred to as an *action* of G on V. We say that G acts on V. A representation is called $faithful$ if its kernel is the identity group $\{e\}$.

If a group G acts on a set V, then it naturally induces other actions. Let $S \subseteq V$ be a subset of a set V, then for any $g \in G$ the translate S^g is also a subset of V. With S^g we refer to the image of the permutation associated to g, by our group action. This however yields an action of G on the power set 2^V, since each $g \in G$ determines a permutation of the subsets of V. Also $|S^g| = |S|$ holds, because permutations are bijections. Therefore for any fixed k, the action of G on V also induces an action on the k-subsets of V. In a similar manner we also get an action of G on the ordered k-tuples of elements of V.

Let G be a permutation group on the set V, a subset $S \subseteq V$ is called *G-invariant* if for all $g \in G$ and $s \in S$ holds, that $s^g \in S$. If S is invariant under G, then each $g \in G$ permutes the elements of S. We denote the *restriction* of a permutation g to S by $g \restriction S$. The mapping $g \mapsto g \restriction S$ is a homomorphism from G into $Sym(V)$ and the image of G under this homomorphism is a permutation group on S. We denote this image as $G \restriction S$.

A permutation group G on V is *transitive* if for every $x, y \in V$ there exists $g \in G$, such that $x^g = y$. A G-invariant subset S of V is called an *orbit* of G if $G \restriction S$ is transitive on S. Let $x \in V$, then it is easy to check, that $x^G := \{x^g : g \in G\}$ is an orbit of G. If $y \in x^G$, then $y^G = x^G$ and if $y \notin x^G$, then $y^G \cap x^G = \emptyset$, so each element lies in a unique orbit of G. The orbits of G partition V and any G-invariant subset of V is a union of orbits of G.

Let G be a permutation group on V. We define the *stabilizer* G_x of x as the set of all permutations $g \in G$ such that $x^g = x$. Now G_x is a subgroup of G, because the identity is in G_x and for every $g, h \in G_x$ it holds that $x^{gh} = x$. If $x_1, \ldots, x_r$ are distinct elements of V, then we define the *pointwise stabilizer* of $\{x_1, \ldots, x_n\}$ as the

group

$$G_{x_1,\dots,x_r} := \bigcap_{i=1}^{r} G_{x_i}.$$

If S is a subset of V, then the stabilizer G_S of S is the set of all permutations g such that $S^g = S$. This is sometimes called the *setwise stabilizer*, since we do not require the elements to be fixed points of the permutations. If $S = \{x_1, \dots, x_n\}$, then $G_{x_1,\dots,x_r}$ is a subgroup of G_S.

Lemma 1.5. *Let G be a permutation group acting on V and let S be an orbit of G. If $x, y \in S$ then the set $\{g \in G : x^g = y\}$ is a right coset of G_x. On the other hand all elements in a right coset of G_x map x to the same point in S.*

Proof. Since G acts transitively on its orbit S, there exists a $g \in G$ such that $x^g = y$. Suppose that there exists $h \in G$ with $x^h = y$. Then $x^g = x^h$ and therefore $x = x^{hg^{-1}}$. Thus hg^{-1} is in G_x, which implies that $h \in G_x g$. We have shown that all permutations, that map x to y have to be in the same right coset. For the other implication consider elements hg of $G_x g$, for some $h \in G_x$. Since h is in G_x the image x^{hg} equals $(x^h)^g = x^g$. Thus every element of $G_x g$ maps x to x^g. □

1.4 Homomorphisms

Let X and Y be graphs. A *homomorphism* from X to Y is a mapping $f : V(X) \to V(Y)$, such that two adjacent vertices $x \sim y$ with $x, y \in V(X)$ are mapped to adjacent vertices $f(x) \sim f(y)$ in Y. Note that if $x \sim y$ then $f(x) \neq f(y)$, since $f(x) \not\sim f(x)$ holds for simple graphs. If X and Y are directed graphs, then a map $f : V(X) \to V(Y)$ is a homomorphism if for every arc (x, y) in X, the image $(f(x), f(y))$ is an arc in Y. Homomorphisms preserve structure and throughout mathematics it proved very fruitful to explore them and their interactions with various objects.

A graph is called *bipartite* if its vertex set V can be partitioned into two parts $V = V_1 \cup V_2$ with $V_1 \cap V_2 = \emptyset$, such that every edge $e = xy$ has one end in V_1 and the other in V_2. The *complete bipartite* graph $K_{m,n}$ is a bipartite graph such that $|V_1| = m$, $|V_2| = n$ and all vertices of V_1 are adjacent to all vertices of V_2. For any bipartite graph there exists a homomorphism to K_2, since you can always map V_1 onto one vertex and V_2 onto the other.

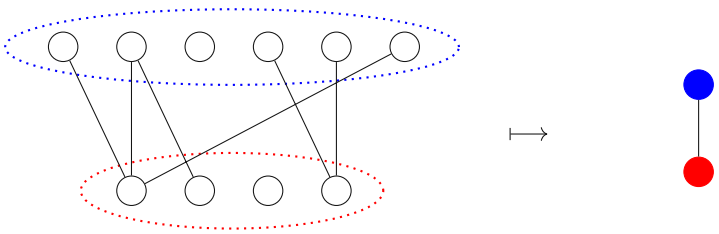

Figure 7: A bipartite graph with a homomorphism to K_2

1.4.1 Graph Colorings

Another important example of graph homomorphisms are graph colorings. A *proper coloring* of a graph X is a mapping from $V(X)$ into some finite set of colors, such that two adjacent vertices have different colors. If X can be properly colored with a set of k colors, then we say that X can be properly k-colored. The *chromatic number* of X, denoted by $\chi(X)$, is the least value for k such that X can be properly k-colored. As seen in Figure 7 all bipartite graphs have a chromatic number of two. Vertices with the same image are said to be in the same *color class*. Every color class is an independent set.

Lemma 1.6. *The chromatic number $\chi(X)$ of a graph X is the least integer r such that there is a homomorphism from X to K_r.*

Proof. Suppose X can be properly colored with r colors. Consider a mapping $f : X \to K_r$ that maps each color class onto a different vertex of K_r. This mapping is a homomorphism, since the images of adjacent vertices are necessarily adjacent. Lets now assume that there is a homomorphism from X to K_r. If $y \in V(Y)$, we define the preimage $f^{-1}(y)$ of y by

$$f^{-1}(y) := \{x \in V(X) : f(x) = y\}.$$

Because we have no loops and f is a homomorphism, $f^{-1}(y)$ is an independent set. Therefore we can assign every preimage a color and we receive r distinct color classes. Thus we get a proper r-coloring of X, which implies that $\chi(X) \leq r$. □

A *retraction* is a homomorphism f from a graph X to a subgraph $Y \subseteq X$ such that the restriction $f \restriction Y$ to $V(Y)$ is the identity map. If such a retraction from X to Y exists, then we say that Y is a *retract* of X. If X has a clique of size $k = \chi(X)$, then any k-coloring of X determines a retraction onto the clique.

When we consider graphs with loops, homomorphisms can behave very differently. Since the loop vertex x is adjacent to itself all neighbors of x can also be mapped to it. In this case many of the above derivations are no longer applicable. For example every graph can be mapped to the trivial graph with one vertex x and a loop xx, as seen in Figure 8. We will explicitly mention if we allow loops and possible repercussions that come with this generalization.
If there exists a surjective homomorphism from a graph X to a graph Y and a surjective homomorphism from Y to X, then it is easy to see that X is isomorphic to Y. Let f be a homomorphism from X to Y, then the preimages $f^{-1}(y)$ of each vertex $y \in V(Y)$ are called the *fibres* of Y. The fibres of f determine a partition π of $V(X)$, called the *kernel* of f. As we already have seen if Y has no loops, then the kernel is a partition into independent sets. For a given graph X together with a partition π of $V(X)$ we define a graph X/π such that the cells of the partition are the vertices and there is an edge between two cells, if there exists an edge in X with endpoints in those cells. If there is an edge with endpoints in the same cell, then we have a loop in X/π. There is a natural homomorphism from X to X/π with kernel π.

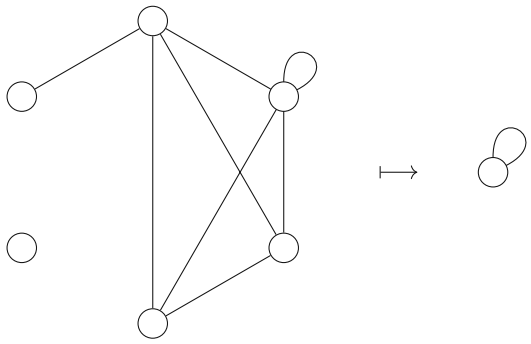

Figure 8: A homomorphism that maps every graph to the trivial graph

1.4.2 Cores

If X, Y and Z are graphs and there exist homomorphisms f from X to Y and g from Y to Z, then the composition $g \circ f$ is a homomorphism from X to Z. This motivates us to define a relation "$\to$" on the class of all graphs by $X \to Y$ if there exists a homomorphism from X to Y. This relation is transitive and reflexive, because the identity map is a homomorphism, hence $X \to X$. It is not antisymmetric however. Consider the example of the complete bipartite graph $K_{3,3}$ and K_2, then $K_{3,3} \to K_2$ and $K_2 \to K_{3,3}$ holds, but $K_{3,3} \neq K_2$ and even $K_{3,3} \ncong K_2$. In order to obtain a partial order we therefore have to define an equivalency on this structure. We say two graphs X and Y are *homomorphically equivalent* if $X \to Y$ and $Y \to X$. If we consider the collection of these equivalency classes then "$\to$" is a partial ordering. We will now take a closer look at properties of these equivalency classes.

A graph X is a *core* if any homomorphism from X to itself is a bijection, implying that any endomorphism is also an automorphism. The simplest examples of cores are the complete graphs. A subgraph Y of X is a core of X if Y is a core and there is a homomorphism from X to Y. We will show that every graph X has a unique core, up to isomorphism, which we will denote by $X^\bullet$. If Y is a core of X and f is a homomorphism from X to Y, then $f \restriction Y$ must be an automorphism of Y. Let $g : Y \to Y$ be the inverse of this restriction, thus $g \circ f \restriction Y$ is the identity on Y. This implies the existence of a retraction from X to Y, since $g \circ f : X \to Y$ is a homomorphism and restricted to Y the identity map. Therefore every core of X is a retract. Any complete graph must be its own core, because there exists no homomorphism to a proper subgraph.

Lemma 1.7. *Let X and Y be cores. Then X and Y are homomorphically equivalent if and only if they are isomorphic.*

Proof. Suppose X and Y are homomorphically equivalent, with homomorphisms $f : X \to Y$ and $g : Y \to X$. Then $f \circ g$ is a homomorphism from Y to itself and therefore an automorphism, because Y is a core. Because $f \circ g$ and $g \circ f$ are surjective, f and g have to be surjective, which implies that X and Y are isomorphic. Conversely if X and Y are isomorphic then they are also homomorphically equivalent. □

Lemma 1.8. *Every graph X has a core, which is an induced subgraph and is unique up to isomorphism.*

Proof. Since X is finite and the identity is a homomorphism on X, the family of subgraphs of X to which X has a homomorphism is finite and nonempty and therefore has a minimal element with respect to inclusion. This minimal element is a core. It is also an induced subgraph, because every core is a retract. Suppose that Y_1 and Y_2 are cores of X, with the homomorphisms $f_i : X \to Y_i$. Then $f_1 \restriction Y_2$ is a homomorphism from Y_2 to Y_1 and $f_2 \restriction Y_1$ is a homomorphism from Y_1 to Y_2, because they are cores and a restriction of a homomorphism is a homomorphism. Therefore the above lemma implies that Y_1 and Y_2 are isomorphic. □

The following Lemma explains when two graphs are homomorphically equivalent.

Lemma 1.9. *Two graphs X and Y are homomorphically equivalent if and only if their cores are isomorphic.*

Proof. Assume that X and Y are homomorphically equivalent and $f : X \to Y$ is a homomorphism. Then there is a sequence of homomorphisms

$$X^\bullet \to X \to Y \to Y^\bullet,$$

which yields a homomorphism from $X^\bullet \to Y^\bullet$. Therefore $X^\bullet$ and $Y^\bullet$ are homomorphically equivalent. Conversely if $X^\bullet$ and $Y^\bullet$ are homomorphically equivalent, then we have a sequence of homomorphisms

$$X \to X^\bullet \to Y^\bullet \to Y,$$

which yields a homomorphism form X to Y.
Hence two graphs are homomorphically equivalent if and only if their cores are. By the previous lemma, two cores are homomorphically equivalent if and only if they are isomorphic, which proves this lemma. □

Now we can prove the claims we stated for "$\to$".

Corollary 1.1. *The relation "$\to$" is a partial order on the set of isomorphism classes of cores.*

Proof. As stated before "$\to$" is a transitive and reflexive relation on the set of isomorphism classes of graphs. Therefore it is also reflexive and transitive on isomorphism classes of cores. By Lemma 1.5. homomorphically equivalent cores are isomorphic, thus "$\to$" is antisymmetric, reflexive and transitive. □

1.5 Products

If X and Y are graphs, we define their *product* $X \times Y$ as a graph with vertex set $V(X) \times V(Y)$ and $(x, y) \sim (x', y')$ if and only if $x \sim x'$ and $y \sim y'$. The map that maps (x, y) to (y, x) is an isomorphism from $X \times Y$ to $Y \times X$, implying commutativity of the product. Similarly for associativity, there is an isomorphism from $(X \times Y) \times Z$

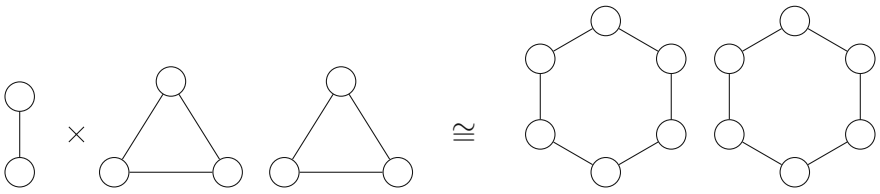

Figure 9: $K_2 \times 2K_3$ is isomorphic to $2C_6$

to $X \times (Y \times Z)$. This justifies the name of product. However if $X \times Y_1 \cong X \times Y_2$ it does not follow that $Y_1 \cong Y_2$. Consider for example $K_2 \times 2K_3 \cong K_2 \times C_6 \cong 2C_6$: Both products are isomorphic to $2C_6$, but unlike "normal" products, their factors can be non isomorphic.

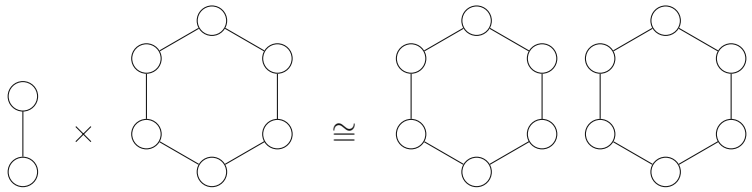

Figure 10: $K_2 \times C_6$ is also isomorphic to $2C_6$

Further noteworthy properties include, that for any graph X the product $X \times K_1$ is always the empty graph, since we do not allow loops, and the product of two connected graphs is connected if and only if at least one of the factors is not bipartite.
For any $x \in V(X)$ the set $\{(x, y)|y \in V(Y)\}$ is always an independent set, therefore we can define a map

$$p_X : (x, y) \mapsto x, \tag{8}$$

which is a homomorphism from $X \times Y$ to X. It is called the *projection* from $X \times Y$ to X. Analogously for Y we can define the projection p_Y from $X \times Y$ to Y.

Theorem 1.4. *Let X, Y and Z be graphs. If $f : Z \to X$ and $g : Z \to Y$, then there exists a unique homomorphism $\phi : Z \to X \times Y$ such that $f = p_X \circ \phi$ and $g = p_Y \circ \phi$.*

Proof. Assume there are homomorphisms $f : Z \to X$ and $g : Z \to Y$. We define

$$\phi : z \mapsto (f(z), g(z)),$$

which is a homomorphism from Z to $X \times Y$. Now ϕ fulfills the requirements $p_X \circ \phi = f$ and $p_Y \circ \phi = g$. But ϕ is also uniquely determined by f and g, which completes our proof. □

There are many other types of graph products, apart from the (natural) product of two graphs, that can be defined. In chapter 4 we consider the lexicographic product and explore its properties and utilities.

2 Kneser Graphs

In this chapter we describe Kneser graphs and provide an alternative proof of the Erdős-Ko-Rado theorem. We also cover a wide range of other graph classes with interesting properties. This chapter is based on chapter 3,4 and 7 of [12].

In the first half we start by exploring multiple graph classes and show that Kneser graphs are regular and transitive, among other properties. We briefly discuss the Petersen graph, which is probably the most famous Kneser graph, due to its many appearances as a counter example for certain conjectures. Then we define fractional colorings and show that Kneser graphs serve a similar purpose in that setting as complete graphs did for regular colorings. We also show that fractional colorings can be related to the field of linear programming.

The goal of the second half of this chapter is to prove that the independence number of Kneser graphs has the same bound as stated in the Erdős-Ko-Rado theorem. We calculate the fractional chromatic number of Kneser graphs and show that the fractional chromatic number is related to the independence number for vertex transitive graphs.

2.1 $J(v, k, i)$ Graphs

We will start by introducing a family of graphs $J(v, k, i)$, that contains many important graph classes. They are defined by a combinatoric construction and are therefore closely related to the mathematical field of combinatorics.

Definition 2.1. *Let $v, k, i \in \mathbb{N}$, with $i \leq k \leq v$ and let Ω be a set with $|\Omega| = v$. We define $J(v, k, i)$ to be the graph with vertex set $V = \{S \subset \Omega : |S| = k\}$, where two subsets are adjacent if their intersection has size i.*

Vertices of this graph are subsets with size k, thus $J(v, k, i)$ has exactly $\binom{v}{k}$ vertices. It is also a regular graph with degree $\binom{k}{i}\binom{v-k}{k-i}$, since for every i elements in a vertex, there are $k - i$ out of $v - k$ other candidates for possible vertices. Further $J(v, 1, 0)$ is the complete graph K_v and $J(v, 1, i)$ with $i > 1$ is the empty graph E_v.

Lemma 2.1. *If $i \leq k \leq v$, then $J(v, k, i) \cong J(v, v - k, v - 2k + i)$.*

Proof. Consider a function that maps a k-subset of Ω to its complement with regards to Ω. This function is an isomorphism between $J(v, k, i)$ and $J(v, v - k, v - 2k + i)$, since $\binom{v}{k} = \binom{v}{v-k}$ and thus sets of size k get mapped to sets of size $v - k$. The images of two adjacent vertices have to be adjacent again, because they intersect in at least $v-2k$ elements. They also share i elements and therefore intersect in exactly $v-2k+i$ elements. □

Due to this isomorphism wlog we can assume that $2k \leq v$, otherwise we choose the isomorphic graph.

Definition 2.2. *Let $2k \leq v$, then we define the Johnson graphs as $J(v, k, k-1)$ and the Kneser graphs as $J(v, k, 0)$.*

For the rest of this chapter Kneser graphs will be our focus.

Let g be a permutation on Ω, then g also permutes the subsets of Ω and in particular the subsets of size k. If $S \subseteq \Omega$ and $T \subseteq \Omega$, then

$$|S \cap T| = |S^g \cap T^g|.$$

Thus g is an automorphism on $J(v, k, i)$. This leads to the following lemma.

Lemma 2.2. *Let $v \geq k \geq i$, then $Aut(J(v, k, i))$ contains a subgroup isomorphic to $Sym(v)$.*

Proof. Since all permutations $g \in Sym(\Omega)$ are also automorphisms on $J(v, k, i)$ the lemma holds. □

A question one could ask now is whether $Aut(J(v, k, i)) \cong Sym(\Omega)$. In general this is not the case, because $Aut(J(v, k, i))$ is a permutation group acting on a set of size $\binom{v}{k}$, meaning that they do not have to be equal.

2.2 Transitive graphs

A graph X is *vertex transitive*, or just *transitive* if its automorphism group acts transitively on $V(X)$. As a reminder that means that for all $x, y \in V(X)$ there exists an automorphism $g \in Aut(X)$ such that $x^g = y$. This is a very strong property that guarantees a lot of algebraic structure within a graph.

Theorem 2.1. *The $J(v, k, i)$ graphs are vertex transitive.*

Proof. Let Ω be the underlying set, then all permutations of $Sym(v)$ also permute the subsets of Ω. There are permutations that map k-subsets to any other k-subset and thus are valid automorphisms to prove transitivity. □

Another family of transitive graphs are the *k-cubes*. A *k-cube* Q_k is a graph with vertex set equal to $\{0, 1\}^k$ and two binary k-tuples are adjacent if and only if they differ in exactly one coordinate position.

Lemma 2.3. *The k-cube Q_k is vertex transitive.*

Proof. Let v be a fixed k-tuple. The map $\rho_v : x \mapsto x + v$, with addition mod 2 in every coordinate, is a permutation of the vertices of Q_k. This mapping is also an automorphism, because two k-tuples x and y differ in exactly one coordinate position if and only if $x + v$ and $y + v$ differ in exactly one coordinate position. For every $v \in \{0, 1\}^k$ we obtain an automorphism this way and they form a subgroup H of the automorphism group $Aut(Q_k)$. Now H acts transitively on V, because for any two $x, y \in V$ the automorphism $\rho_{y-x} \in H$ maps x to y, since

$$x^{\rho_{y-x}} = x + y - x = y.$$

□

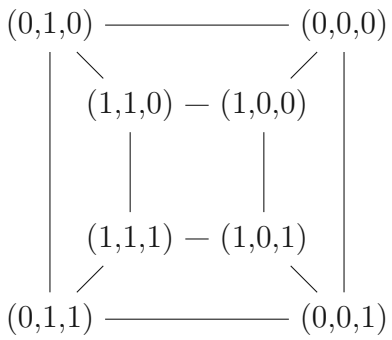

Figure 11: The 3-cube Q_3

Note that the group H above is not equal to the automorphism group $Aut(Q_k)$. As we have already seen H has size 2^k. Any permutation of the k coordinate positions is an automorphism of Q_k. The set of these permutations K is a subgroup of $Aut(Q_k)$, that is isomorphic to $Sym(k)$. Since H and K are both subgroups of $Aut(Q_k)$, their complex product HK is also a subgroup. The size of HK is given by

$$|HK| = \frac{|H||K|}{|H \cap K|}.$$

Now $H \cap K$ is the identity group. To see this consider any non trivial element of K and any non trivial element of H. The image of $(0, \ldots, 0)$ under any permutation of its positions is itself, while no non trivial addition can map $(0, \ldots, 0)$ to $(0, \ldots, 0)$. Thus $|H \cap K| = 1$ and we get $|HK| = 2^k k!$. Since HK is a subgroup of $Aut(Q_k)$, we can conclude, that $|Aut(Q_k)| \geq 2^k k!$.

Another example of a family of transitive graphs are the *circulants*. With $\mathbb{Z}_n$ we denote the additive group of integers modulo n. Let C be a subset of $\mathbb{Z}_n \setminus \{0\}$, then we define a directed graph $X = X(\mathbb{Z}_n, C)$ with vertex set $\mathbb{Z}_n$. Two vertices i, j are an arc $(i, j) \in E$ if and only if $j - i \in C$. The graph $X(\mathbb{Z}_n, C)$ is called a *circulant* of order n and C is called its *connection set*. If C is closed under additive inverses, then (i, j) is an arc if and only if (j, i) is an arc, in essence providing an undirected graph. That is because with $c \in C$ we get $-c \in C$ and therefore $(i, j) \in C$ if and only if $j - i \in C \Leftrightarrow i - j \in C$ which implies $(j, i) \in C$.

The permutation that maps i to $i + 1$ is an automorphism of X. To see this suppose that (i, j) is an arc in X. Then $j - i \in C$, but this implies that $(j + 1) - (i + 1) \in C$. Thus $(i + 1, j + 1)$ is also an arc in X, proving the homomorphism property of the permutation. If C is closed under additive inverses, then the permutation that maps i to $-i$ is also an automorphism. Therefore, if X is undirected, its automorphism group has to have order at least $2n$. Every cycle C_n is isomorphic to a circulant of order n, with connection set $\{1, -1\}$. We map every vertex of C_n to $\{0, \ldots, n - 1\}$, such that i is adjacent to j if and only if $j - i \equiv \pm 1 \mod n$. This is indeed an isomorphism and we immediately get that $C = \{1, -1\}$. The complete and empty graphs are also

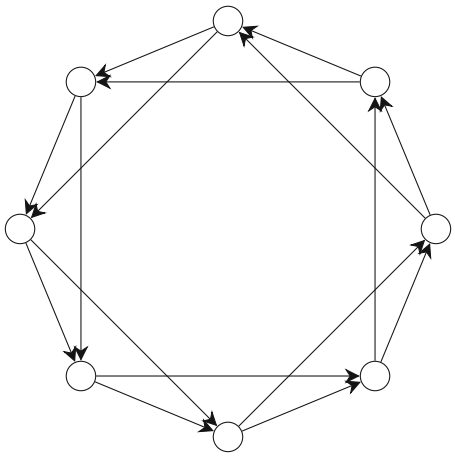

Figure 12: The circulant $X(\mathbb{Z}_8, \{1, 2\})$

circulants, with $C = \mathbb{Z}_n$ and $C = \emptyset$ respectively.

Since the above mentioned permutation, that maps i to $i + 1$, is an automorphism we can map any vertex x to y by applying this automorphism a certain number of times.

2.2.1 Cayley Graphs

Most of the graph examples we have seen so far in this chapter are also members of a more general family of vertex transitive graphs. Let G be a group and let C be a subset of G, that is inverse-closed and does not contain the identity. We define the *Cayley graph* $X(G, C)$ to be the graph with vertex set G and edge set

$$E(X(G, C)) = \{gh : hg^{-1} \in C\}.$$

If C is an arbitrary subset of G, then we can define a directed graph $X(G, C)$ with vertex set G and arc set $\{(g, h) : hg^{-1} \in C\}$. If C is inverse-closed and does not contain the identity, then this graph reduces to a Cayley graph.

Theorem 2.2. *The Cayley graph $X(G, C)$ is vertex transitive.*

Proof. For each $g \in G$ the mapping

$$\rho_g : x \mapsto xg$$

is a permutation of the elements of G, since G is a group. This permutation is also an automorphism on $X(G, C)$, because

$$\rho_g(y)\rho_g(x)^{-1} = (yg)(xg)^{-1} = ygg^{-1}x^{-1} = yx^{-1}.$$

Thus $xg \sim yg$ if and only if $x \sim y$. The set of permutations ρ_g for all $g \in G$ forms a subgroup of the automorphism group $Aut(X(G, C))$. To show that the automorphism

group acts transitively on G, let $g, h \in G$ be arbitrary vertices. The automorphism $\rho_{g^{-1}h}$ maps g to h, since $\rho_{g^{-1}h}(g) = gg^{-1}h = h$. Thus for every pair of arbitrary vertices we have found an automorphism, that maps one to the other. □

In the proof above we have shown that we only need elements of G to ensure transitivity, rather than the whole automorphism group.

Corollary 2.1. *Let $Y = X(G, C)$ be a Cayley graph. Then a subgroup of $Aut(Y)$ isomorphic to G is transitive on Y.*

The k-cube Q_k is a Cayley graph for the abelian group $(\mathbb{Z}_2)^k$ and a circulant on n vertices is a Cayley graph for the cyclic group of order n.

2.2.2 Cores of Vertex Transitive Graphs

We will briefly gather a few short facts about cores of vertex transitive graphs.

Theorem 2.3. *Let X be a vertex-transitive graph, then its core $X^\bullet$ is also vertex-transitive.*

Proof. Suppose $x \neq y$ are two distinct vertices of $X^\bullet$. Since X is vertex-transitive, there exists an automorphism that maps x to y. The retraction from X to $X^\bullet$ is a homomorphism, thus composition with this automorphism yields a homomorphism $f : X \to X^\bullet$. Therefore the restriction $f \restriction X^\bullet$ is an automorphism of $X^\bullet$, that maps x to y. □

Theorem 2.4. *Let X be a vertex-transitive graph, then $|V(X^\bullet)|$ divides $|V(X)|$.*

Proof. We show that for any homomorphism f from X to $X^\bullet$ the fibres of f, which are the preimages $f^{-1}(v)$, have the same size. Thus showing, that f maps the same amount of vertices to every image proving the theorem.
Lets assume that $f : X \to X$ is a homomorphism and $f(X) =: Y$ is a core of X. For every automorphism $g \in Aut(X)$, the translate Y^g gets mapped onto Y by f. Thus Y^g has one vertex in each fibre of f. Suppose that $v \in V(X)$ and let F be the fibre of f that contains v. Because X is vertex transitive the number of automorphisms that contain v does not depend on the actual choice of v. Lets call this number N. Now every image Y^g meets F and therefore we get

$$|Aut(X)| = |F|N.$$

Since N does not depend on F this shows that every fibre of f has the same size. □

2.2.3 Edge Transitive Graphs

A graph $X = (V, E)$ is *edge transitive* if its automorphism group acts transitively on E. Meaning that for all $e_1, e_2 \in E$ there exists an automorphism $g \in Aut(X)$ such that $e_1^g = e_2$. The image of an edge xy under an automorphism f is to be understood as $f(x)f(y) \in E$, because the images of adjacent vertices have to be adjacent again.
Edge transitive graphs are not necessarily vertex transitive. Consider the complete

bipartite graph $K_{m,n}$ where $m \neq n$. This graph is not vertex transitive, because a vertex with degree m can not be mapped to a vertex with degree n by an automorphism. However it is edge transitive, since we can permute any vertex within its degree class. There are also graphs that are vertex transitive but not edge transitive. We have already been introduced to a class of graphs that is both vertex transitive and edge transitive.

Lemma 2.4. *The $J(v, k, i)$ graphs are edge transitive.*

Proof. As mentioned previously all automorphisms of the set Ω also define automorphisms on the set of edges of $J(v, k, i)$ graphs. Let $xy \in E$ and $ab \in E$ be two distinct edges of the graph. Because they are edges we know that $|x \cap y| = |a \cap b| = i$. Now we construct a permutation on Ω by mapping all vertices in the intersection of one edge to all vertices in the intersection of the other edge. The other $k - i$ elements of x get mapped to the $k - i$ elements of a and similarly for y and b. All other elements can be fixed points. This is indeed a permutation on Ω and thus an automorphism on $J(v, k, i)$, which also maps xy to ab. □

We can also prove that all edge transitive graphs that are not vertex transitive are bipartite.

Lemma 2.5. *Let X be an edge transitive graph with no isolated vertices. If X is not vertex transitive, then $Aut(X)$ has exactly two orbits, and these two orbits are a bipartition of X.*

Proof. Suppose X is edge transitive but not vertex transitive and let $xy \in E(X)$. If $w \in V(X)$, then w lies on an edge and there is an element of $Aut(X)$ that maps this edge onto xy. Thus any vertex of X lies in either the orbit of x under $Aut(X)$, or the orbit of y. Therefore $Aut(X)$ has to have exactly two orbits. An edge that joins two vertices in the same orbit cannot be mapped by an automorphism to an edge that contains a vertex from the other orbit. Since X is edge transitive and every vertex lies in an edge, there is no edge joining two vertices in the same orbit. Thus X is bipartite and the orbits are the bipartition for it. □

2.2.4 Arc Transitive Graphs

A graph is *arc transitive* if its automorphism group acts transitively on the set of arcs. As a reminder an arc is an ordered pair of adjacent vertices. This is in fact a stronger property than being vertex transitive or edge transitive, meaning that every arc transitive graph is also vertex and edge transitive. To see this suppose that X is an arc transitive graph with no isolated vertices. Then X has to be edge transitive, since any undirected graph can be interpreted as a directed graph with arcs in both directions. Let $x, y \in V(X)$, we have to show that there exists an automorphism that maps x to y. Since X has no isolated vertices x and y are both part of possibly different edges. Due to the X being arc transitive, there is an automorphism that maps the edge with endpoint x to the other edge with endpoint y and thus x to y.

Lemma 2.6. *The $J(v, k, i)$ graphs are arc transitive.*

Proof. Consider the vertex $\{1, \ldots, k\}$. The stabilizer of this vertex contains $Sym(k) \times Sym(v-k)$, since we can permute these k elements and the other $v-k$ elements of Ω. Any two k-sets meeting the vertex $\{1, \ldots, k\}$ in an i-set can be mapped to each other by this group. We simply permute the $k-i$ elements from one vertex to the other. □

2.3 Kneser Graphs

The $J(v, k, 0)$ graphs are also known as *Kneser* graphs. We denote them by $K_{v:k}$ and they are the graphs consisting of k-subsets of a set Ω with $|\Omega| = v$, where two vertices are adjacent if they are disjoint. As we have already seen $K_{v:1}$ are the complete graphs K_v.

The goal of this chapter is to state the Erdős-Ko-Rado theorem for Kneser graphs. We introduce fractional colorings and use them to prove the theorem. But first we take a closer look at a certain Kneser graph that has many interesting properties.

2.3.1 Petersen Graph

One of the most famous Kneser graphs is $J(5, 2, 0)$, also known as the *Petersen graph*. In the previous chapter we proved that all Kneser graphs are arc transitive and therefore edge transitive, vertex transitive and regular. Due to being a $J(5, 2, 0)$ graph,

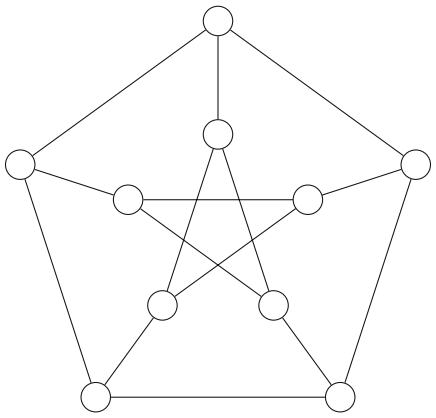

Figure 13: The Petersen graph $J(5, 2, 0)$

we know that $Sym(5)$ acts on it and therefore we know that its automorphism group has order at least $5! = 120$. In fact it can be shown that the automorphism group of the Petersen graph has order exactly 120 [12]. Further the Petersen graph is *distance transitive*, a *Moore graph* and a *strongly regular graph*, but not a Cayley graph [12].

2.3.2 Fractional Colorings

Fractional Colorings are a very important concept, that will help us prove the main theorem of this chapter. Given by their name, they are closely related to proper graph colorings.

Let X be a graph. With $\mathcal{I}(X)$ we denote the set of all independent sets of X and with $\mathcal{I}(X, i)$ the independent sets that contain the vertex i. A *fractional coloring* of a graph X is a non negative function $f : \mathcal{I}(X) \to \mathbb{R}$ such that

$$\sum_{S \in \mathcal{I}(X,x)} f(S) \geq 1$$

holds for any vertex x. We define the *weight* of a fractional coloring as the sum of all of its values and the *fractional chromatic number* $\chi^*(X)$ as the minimum possible weight of a fractional coloring. A fractional coloring is called *regular* if for each vertex x holds

$$\sum_{S \in \mathcal{I}(X,x)} f(S) = 1.$$

We show now that proper colorings are also fractional colorings. If X has a proper k-coloring, then its color classes $V_1, \ldots, V_k$ are a partition of $V(X)$. We can define a function f such that $f(V_i) = 1$ and $f(S) = 0$ for all other independent sets S. Then f is a fractional coloring of X with weight k and thus

$$\chi^*(X) \leq \chi(X)$$

holds for any graph X.
Suppose now that we have a 01-valued fractional coloring of X with weight k. Then the support of f consists of k independent sets $V_1, \ldots, V_k$ whose union is $V(X)$. We can construct a proper k-coloring by coloring the vertex x with the smallest i such that $x \in V_i$, to avoid coloring twice. Therefore the chromatic number $\chi(X)$ is the minimum weight of a 01-valued fractional coloring.

Figure 14 shows an example of a fractional coloring f of C_5, with vertex set $V = \{v_1, \ldots, v_5\}$. For this example we defined the support of f to be all maximal independent sets of C_5. These are exactly the sets $\{v_1, v_3\}$, $\{v_1, v_4\}$, $\{v_2, v_4\}$, $\{v_2, v_5\}$ and $\{v_3, v_5\}$. The value of f for these is set to $1/2$. This satisfies the definition of a fractional coloring and can be thought of as choosing a distinct color for every maximal independent vertex set. Since all vertices are in exactly two of these sets, this can be visualized by each vertex having two colors. The weight of this fractional coloring is $5/2$, which is lower than the chromatic number $\chi(C_5) = 3$.

Another example to consider is the Kneser graph $K_{v:r}$. The r-sets that contain the element l form an independent set of size $\binom{v-1}{r-1}$, because l has to be fixed for every r-set. Since the vertices of $K_{v:r}$ themselves consist of r elements, each vertex has to lie in exactly r of those independent sets mentioned above, one for each element. We

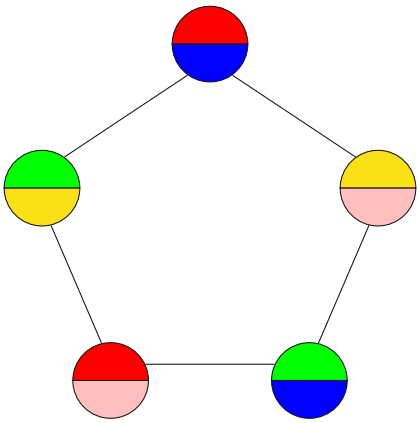

Figure 14: A fractional coloring with weight $\frac{5}{2}$

define a function f with $f(x) = 1/r$ for x in these independent sets and zero elsewhere. This is a fractional coloring of $K_{v:r}$ with weight

$$\sum_{S \in \mathcal{I}(X,x)} f(S) = \sum_{S \in \mathcal{I}(X,x)} \frac{1}{r} = \frac{v}{r}.$$

Thus we get $\chi^*(K_{v:r}) \leq v/r$.

2.3.3 Fractional Cliques

Fractional cliques are a generalization of cliques and we foreshadow that they are dual to fractional colorings in a way. They will prove to be very useful for us.

A *fractional clique* of a graph X is a non-negative real-valued function on V such that the sum of its values on vertices of any independent set is at most one. The *weight* of a fractional clique is the sum of its values. The *fractional clique number* ω^* of X is the maximum possible weight of a fractional clique. Given a clique of size k in X, its characteristic function is a 01-valued fractional clique of weight k and therefore it holds that

$$\omega(X) \leq \omega^*(X).$$

Let $\alpha(X)$ be the maximum size of an independent set in X, then we define a function $g := \alpha(X)^{-1}\mathbf{1}$, where $\mathbf{1} : V \to \mathbb{R}$ is the function that maps every vertex to $1 \in \mathbb{R}$. Now g is a fractional clique, because the sum of values of vertices on any independent set has to be less than or equal to one.

Lemma 2.7. *Let X be a graph, then*

$$\omega^*(X) \geq \frac{|V(X)|}{\alpha(X)}.$$

Proof. Since we have already shown, that g defined as above is always a fractional clique, it follows that the fractional clique number is at least

$$\sum_{x \in V(X)} g(v) = \frac{1}{\alpha(X)} \sum_{x \in V(X)} 1 = \frac{|V(X)|}{\alpha(X)}.$$

□

For vertex transitive graphs, we can determine the fractional clique number exactly.

Lemma 2.8. *Let X be a vertex transitive graph, then*

$$\omega^*(X) = \frac{|V(X)|}{\alpha(X)}.$$

Further $\alpha(X)^{-1}\mathbf{1}$ is a fractional clique with this weight.

Proof. Suppose g is a non zero fractional clique of X. Let $\gamma \in Aut(X)$ be an automorphism on X, then we define a function

$$g^\gamma(x) := g(x^\gamma).$$

Now g^γ is a fractional clique, since independent sets get mapped to independent sets by automorphisms. Further g^γ has the same weight as g. We can now define a function $\hat{g}$, that is also a fractional clique with the same weight as g by

$$\hat{g} := \frac{1}{|Aut(X)|} \sum_{\gamma \in Aut(X)} g^\gamma.$$

Let $x, y \in V$ be a vertices. We will show that $\hat{g}(x) = \hat{g}(y)$. Since X is vertex transitive, we know that there exists an $\omega \in Aut(X)$ such that $x^\omega = y$. Because $Aut(X)$ is a group it follows that

$$g^\gamma(x) = g(x^\gamma) = g(x^{\gamma\omega^{-1}\omega}) = g(y^{\gamma\omega^{-1}}).$$

Lastly we have that $\omega^{-1}Aut(X) = Aut(X)$ and therefore both sums are equal and $\hat{g}$ is constant on the vertices of X. If we consider any constant function $c\mathbf{1}$, then $c\mathbf{1}$ is a fractional clique if and only if $c \leq \alpha(X)^{-1}$, which results in the same weight as our bound. Thus the weight of g is equal to the weight of $\hat{g}$ and cannot surpass the bound of this lemma. □

2.3.4 Existence of fractional colorings

Now we show that the fractional chromatic number and the fractional clique number are well-defined.
Let B be the matrix with rows indexed by the vertices of X and columns indexed by the characteristic functions of the independent sets of X. We define B_{x_1,f_1} as 1 if x_1 is in the independent set belonging to f_1 and 0 otherwise. If f is a non negative vector, such that $Bf \geq \mathbf{1}$, meaning that every coordinate of Bf is at least 1, then such an f naturally defines a fractional coloring. On the other hand if g is a non negative vector, such that $g^T B \leq \mathbf{1}^T$, then g defines a fractional clique.

Lemma 2.9. *If a graph X has a fractional coloring f of weight w, then it also has a fractional coloring f' with weight at most w, such that $Bf' = \mathbf{1}$.*

Proof. Suppose f is a fractional coloring such that there exists a position j with $(Bf)_j = b > 1$ while $(Bf)_i = 1$ for all $i \neq j$. We want to define a new fractional coloring, that satisfies our requirements. Let $S_1, \ldots, S_t$ be the independent sets that contain x_j and are in the support of f. For each S_i we define a_i such that

$$a_i \leq f(S_i) \quad \text{and} \quad \sum_{i=1}^{t} a_i = b - 1.$$

We can choose these a_i and they exist, because $\sum_{i=1}^{t} f(S_i) = b$ by assumption. Now we define a function f' by

$$f'(S) = \begin{cases} f(S) - a_i, & \text{if } S = S_i, \\ f(S) + a_i, & \text{if } S = S_i \setminus x_j \text{ and } S \neq \emptyset, \\ f(S), & \text{else.} \end{cases}$$

This function is indeed a fractional coloring since

$$\sum_{S \in \mathcal{I}(X, x_j)} f'(S) = \sum_{i=1}^{t} f(S_i) - a_i = \sum_{i=1}^{t} f(S_i) - (b-1) = 1.$$

If we take the sum of independent sets of other vertices, where S_i is an addend then $S_i \setminus x_j$ is also an independent set of that vertex, thus canceling out the change. Therefore f' is a fractional coloring with no greater weight than f, satisfying $(Bf')_j = 1$ and $(Bf')_i = (Bf)_i$ for all $i \neq j$.
If $(Bf)_j > 1$ for more than one j we simply repeat this construction for those positions. This yields the proposed fractional coloring. □

Theorem 2.5. *Any graph X has a regular rational-valued fractional coloring with weight $\chi^*(X)$.*

Proof. We start by showing that given a fractional coloring f we can construct a fractional coloring f' such that the columns in B that are in $supp(f')$, the support of f', are linearly independent. Further $supp(f') \subset supp(f)$ and the weight of f' is not greater than the weight of f.
Let $B = (b_{ij})_{i=\{1,\ldots,n\}}^{j=\{1,\ldots,m\}}$ be the matrix consisting of characteristic vectors of independent sets of X, we use f to also denote the numerical vector representing the images of f. Suppose the columns of $supp(f)$ are linearly dependent and $(b_{il})_{i=\{1,\ldots,n\}}$ is a linear combination of some $(b_{ij})_{i=\{1,\ldots,n\}}$,

$$\mathbf{b}_l := \begin{pmatrix} b_{1l} \\ \vdots \\ b_{nl} \end{pmatrix} = \sum_{j \in J} \lambda_j \begin{pmatrix} b_{1j} \\ \vdots \\ b_{nj} \end{pmatrix}.$$

We want to set f_l to zero and thus remove $\mathbf{b}_l$ from $supp(f)$. Let $\mathbf{a}$ be the resulting vector of the multiplication $Bf = \mathbf{a}$,

$$\begin{pmatrix} b_{11} & b_{12} & \cdots & b_{1m} \\ \vdots & \ddots & & \vdots \\ \vdots & & \ddots & \vdots \\ b_{n1} & b_{n2} & \cdots & b_{nm} \end{pmatrix} \begin{pmatrix} f_1 \\ \vdots \\ f_l \\ \vdots \\ f_m \end{pmatrix} = \begin{pmatrix} a_1 \\ \vdots \\ \vdots \\ a_n \end{pmatrix} \geq \begin{pmatrix} 1 \\ \vdots \\ \vdots \\ 1 \end{pmatrix}.$$

We have to ensure that the value of f_l gets compensated by the other columns of the linear combination of $\mathbf{b}_l$. If we consider a_1 then we get

$$\begin{aligned} a_1 = \sum_{i \in I} b_{1i} f_i &= \sum_{i \in I \setminus (J \cup l)} b_{1i} f_i + b_{1l} f_l + \sum_{j \in J} b_{1j} f_j \\ &= \sum_{i \in I \setminus (J \cup l)} b_{1i} f_i + \sum_{j \in J} \lambda_j b_{1j} f_l + \sum_{j \in J} b_{1j} f_j \\ &= \sum_{i \in I \setminus (J \cup l)} b_{1i} f_i + \sum_{j \in J} b_{1j} (\lambda_j f_l + f_j). \end{aligned}$$

The last term shows that after setting f_l to zero we can compensate by adding the missing amount to f_j. Therefore we define f' as

$$f'(S) := \begin{cases} f(S) + \lambda_j f_l, & \text{if } S = S_j \text{ and } j \in J, \\ 0, & \text{if } S = S_l, \\ f(S), & \text{else.} \end{cases}$$

This function is indeed a fractional coloring of X with no more weight than f and $f'(S_l) = 0$, which implies that $supp(f') \subset supp(f)$. By iterating this procedure we can construct a fractional coloring with linearly independent support.
In conjunction with Lemma 2.9 we can transform any fractional coloring such that $Bf' = \mathbf{1}$ and $supp(f')$ is linearly independent. Given a graph we can easily define a fractional coloring. Therefore the solution space is not empty and we can always construct solutions on the boundary $Bf = \mathbf{1}$. Further B contains at least n linearly independent columns, since every vertex represents its own independent set. Thus if C is a matrix constructed from n independent columns of B, we get that $Cf = \mathbf{1}$ has a unique solution. This is essentially equal to a basic feasible solution in linear programming [2]. Thus the fractional chromatic number exists or is equal to $-\infty$, but the latter case cannot occur, since minimizing $\mathbf{1}^T f$, with $f \geq 0$ cannot yield $-\infty$. Construction of a solution involves inverting C which yields rational numbers and thus $\chi^*(X)$ is rational. □

As we have already hinted at, an alternative way to obtain $\chi^*(X)$ and $\omega^*(X)$ is via linear programming. Let B be the matrix defined as before, then the fractional chromatic number $\chi^*(X)$ is equal to the value of the linear optimization problem:

$$\begin{gathered} \min \mathbf{1}^T f \\ Bf \geq \mathbf{1} \\ f \geq 0. \end{gathered}$$

Further the fractional clique number $\omega^*(X)$ is the solution of the dual problem:

$$\begin{aligned} &\min g^T\mathbf{1} \\ &g^T B \leq \mathbf{1}^T \\ &g \geq 0. \end{aligned}$$

Lemma 2.10. *Let X be a graph, then $\omega^*(X) \leq \chi^*(X)$.*

Proof. Suppose that f is a fractional coloring of X and g is a fractional clique of X. We consider their difference

$$\mathbf{1}^T f - g^T\mathbf{1} = \mathbf{1}^T f - g^T Bf + g^T Bf - g^T\mathbf{1} = (\mathbf{1}^T - g^T B)f + g^T(Bf - \mathbf{1}).$$

Now the term $(\mathbf{1}^T - g^T B)$ is non negative in every coordiante, since g is a fractional clique. Similarly the term $(Bf - \mathbf{1})$ is also non negative. Scalar multiplication with a vector that is non negative yields a number that is non negative as well. Therefore we are adding two non negative numbers, resulting in $\mathbf{1}^T f - g^T\mathbf{1} \geq 0$. Finally $g^T\mathbf{1} \leq \mathbf{1}^T f$ for any fractional coloring f and fractional clique g implies that the solutions of the optimization problems satisfy $\omega^*(X) \leq \chi^*(X)$. □

For any fractional coloring and fractional clique, we now obtain the inequality

$$\omega(X) \leq \omega^*(X) \leq \chi^*(X) \leq \chi(X).$$

Combining the above lemma with the bound we obtained previously for $\omega^*(X)$ we conclude that for any graph X it holds that

$$\frac{|V(X)|}{\alpha(X)} \leq \omega^*(X) \leq \chi^*(X). \tag{9}$$

2.3.5 Fractional colorings of vertex transitive graphs

If X is a vertex transitive graph, it turns out that inequality (9) is actually an equality. To prove this we first state a couple of theorems and lemmata that will help us.

Lemma 2.11. *Let X be a graph, $x \in V$ a fixed vertex of X and G a group that acts transitively on V. Then there is a bijection from V to the right cosets of G_x and the action of G on V matches with right multiplication on the corresponding coset.*

Proof. Since G is transitive there exists $g \in G$ such that $x^g = y$. Then y corresponds to $G_x g$ by Lemma 1.5. If we consider the image y^v for $v \in G$, then there exists $h \in G$ such that $x^h = y^v$ and y^v corresponds to $G_x h$. But since $x^h = y^v$ if and only if $x^{hv^{-1}} = y$, we know that $hv^{-1} \in G_x g$ and that is exactly the case when $h \in G_x gv$. Therefore the action of G on V translates to right multiplication of the corresponding right coset. □

Theorem 2.6. *Let X be a connected vertex transitive graph. Then X is a retract of a Cayley graph.*

Proof. Suppose X is a connected vertex transitive graph and $x \in V(X)$ is a fixed vertex. We have to show that there exists a Cayley graph Y such that there is a homomorphism h from Y to X and the retraction $h \restriction X$ is the identity. We start by defining a group that will be used to construct the Cayley graph. Let G be the group generated by a certain set of automorphisms

$$G := \langle \{g \in Aut(X) : x \sim x^g\} \rangle.$$

Every element of G can be written as a product of elements of its generating set. Now G acts transitively on $V(X)$, which can be proven by induction over the length of the longest path in X. Therefore the above Lemma guarantees that there is a bijection from $V(X)$ to the right cosets of G_x. As a reminder the stabilizer G_x of x is defined as $G_x := \{g \in G : x = x^g\}$. Let C be the set defined by

$$C := \{g \in G : x \sim x^g\}.$$

Now C can be written as a union of right cosets of G_x, since $x^{hg} = (x^h)^g = x^g$, for $h \in G_x$ and $g \in G$. Further $C \cap G_x = \emptyset$, since $x \not\sim x$ and right cosets are either disjoint or equal. If $x^a \sim x^b$, then $x \sim x^{ba^{-1}}$ if and only if $ba^{-1} \in C$. Let $g \in C$ and $h_1, h_2 \in G_x$, then

$$x = x^{h_1} \sim x^{gh_1} = x^{h_2 g h_1}.$$

Therefore $h_2 g h_1 \in C$, which implies that $G_x C G_x \subseteq C$. Since G_x contains the identity we also get $C \subseteq G_x C G_x$ and thus $C = G_x C G_x$.
We can now define the Cayley graph, that will satisfy our requirements. Let Y be the Cayley graph $X(G, C)$. The right cosets of G_x partition $G = V(Y)$ and we can write every vertex of Y, which is an element of G, as ga for some $g \in G_x$ and $a \in G$. Let $g, h \in G_x$, then the two vertices $ga, hb \in V(Y)$ are adjacent in the Cayley graph if and only if

$$hb(ga)^{-1} = hba^{-1}g^{-1} \in C.$$

This is exactly the case if $ba^{-1} \in C$, since $C = G_x C G_x$. Thus adjacency in Y translates to $ba^{-1} \in C \Leftrightarrow x \sim x^{ba^{-1}} \Leftrightarrow x^a \sim x^b$. If there exists an edge between two elements of distinct right cosets of G_x then the subgraph induced by these two sets has to be a complete bipartite graph. If we assume that two vertices of the same coset are adjacent, then there exists $a \in G$ for $h, g \in G_x$, such that $ha(ga)^{-1} \in C$ which holds if and only if $aa^{-1} = e \in C$, but the identity is not in C.
Now we can take a representative of every distinct right coset of G_x and consider the subgraph Z of Y induced by this set of vertices. Due to the bijection between the right cosets and $V(X)$, the vertex set $V(Z)$ is isomorphic to $V(X)$. Let $y, z \in X$ with $y \sim z$. Since G is transitive on X there exist $g, h \in G$ such that $y = x^g \sim x^h = z$, which is the same adjacency property as in Y. Therefore Z is isomorphic to X and the map sending vertices of Y, belonging to a certain right coset, to vertices of X that correspond to this set is a homomorphism.
All that is left to do is to show that X is a retract of Y. Let $f : Y \to X$ be the homomorphism decribed above, where $f' := f \restriction X$ is an automorphism. Consider the function $f'^{-1} \circ f : Y \to X$, which maps $x \in X$ to $f'^{-1} \circ f(x) = x$. This function is a homomorphism and a witness that X is a retract of Y. □

Corollary 2.2. *Let X be a vertex transitive graph and $Y = X(G, C)$ the Cayley graph as defined in the theorem. Then*

$$\frac{|V(X)|}{\alpha(X)} = \frac{|V(Y)|}{\alpha(Y)}. \tag{10}$$

Proof. From the retraction that maps Y to X we can see that Y can be constructed by replacing each vertex in X with a right coset of G_x, such that all new vertices within a set are independent and if there exists an edge between two distinct right cosets, they form a complete bipartite subgraph. Therefore the largest independent set of X generates $|G_x|$ many independent sets in Y. Thus we get

$$\frac{|V(Y)|}{\alpha(Y)} = \frac{|V(X)||G_x|}{\alpha(X)|G_x|} = \frac{|V(X)|}{\alpha(X)}.$$

□

Let X and Y be graphs, $\varphi : X \to Y$ be a homomorphism and f a fractional coloring of Y. We define the *lift* $\hat{f} : \mathcal{I}(X) \to \mathbb{R}$ of f as the function

$$\hat{f}(S) := \sum_{T:\varphi^{-1}(T)=S} f(T).$$

Lemma 2.12. *Let X and Y be graphs, $\varphi : X \to Y$ be a homomorphism and f a fractional coloring of Y. Then the lift $\hat{f}$ of f is a fractional coloring of X with the same weight as f. The support of $\hat{f}$ consists of the preimages of the independent sets in the support of f.*

Proof. Let $T \in \mathcal{I}(Y)$, then the preimage $\varphi^{-1}(T)$ is an independent set in X, since φ is a homomorphism and thus maps adjacent vertices to adjacent vertices. The support of $f : \mathcal{I}(X) \to \mathbb{R}$ consists of independent sets where f is not zero. The preimages of these sets are exactly the independent sets in the support of $\hat{f}$. Let $u \in V(X)$, then

$$\sum_{T \in \mathcal{I}(X,u)} \hat{f}(T) = \sum_{S:u \in \varphi^{-1}(S)} f(S) = \sum_{S \in \mathcal{I}(Y,\varphi(u))} f(S) \geq 1.$$

Thus $\hat{f}$ is indeed a fractional coloring of X. It remains to show that $\hat{f}$ has the same weight as f. Let $T, U \in \mathcal{I}(Y)$ and suppose that

$$T \cap \varphi(X) = U \cap \varphi(X).$$

Then the preimage $\varphi^{-1}(T)$ is equal to the preimage $\varphi^{-1}(U)$, since the preimage of the image $\varphi(X)$ is X. Therefore two independent sets have the same preimage S if they have the same intersection with $\varphi(X)$. In the calculation of $\hat{f}(S)$ both S and T appear and thus every independent set in Y is part of the sum. The sum of $\hat{f}$ over all independent sets in X is equal to the sum of f over all independent sets in Y, which implies that f and $\hat{f}$ have the same weight. □

Since every fractional coloring of Y allows us to construct the lift, given a homomorphism, we get the following corollary.

Corollary 2.3. *Let X and Y be graphs and $\varphi : X \to Y$ a homomorphism. Then*

$$\chi^*(X) \leq \chi^*(Y).$$

Now we can prove equality of the fractional chromatic number and the fractional clique number for vertex transitive graphs.

Theorem 2.7. *Let X be a vertex transitive graph, then*

$$\omega^*(X) = \chi^*(X) = \frac{|V(X)|}{\alpha(X)}.$$

Proof. We will show that for any vertex transitive graph X holds

$$\chi^*(X) \leq \frac{|V(X)|}{\alpha(X)},$$

which combined with the previous inequality proves the theorem. We have already seen that $\chi^*(K_{v:r}) \leq v/r$ and we will use this as a bound for the chromatic fractional coloring of X. By Theorem 2.6 there exists a Cayley graph $Y = X(G, C)$ such that X is a retract and $|V(Y)|/\alpha(Y) = |V(X)|/\alpha(X)$. The resulting homomorphism ensures by the above corollary that $\chi^*(X(G, C)) \leq \chi^*(X)$. There is also a homomorphism from X to $X(G, C)$, mapping every vertex of X to its corresponding coset representative. Therefore we get $\chi^*(X(G, C)) = \chi^*(X)$. We will now show that there also exists a homomorphism from $Y = X(G, C)$ to the Kneser graph $K_{|V(Y)|:\alpha(Y)}$.
Let $S \subseteq G$ be an independent set of size $\alpha(Y)$ in Y and let $S^{-1} := \{s^{-1} : s \in S\}$ denote the set of inverse elements. We define a function $\varphi : V(Y) \to 2^{V(Y)}$ by

$$\varphi : g \mapsto (S^{-1}g).$$

Let $g, h \in G$ with $g \sim h$, then $\varphi(g) \cap \varphi(h) = \emptyset$. To see this assume that $y \in \varphi(g) \cap \varphi(h)$, then $y = a^{-1}g = b^{-1}h$ for some $a, b \in S$. This implies that $a^{-1}b = hg^{-1} \in C$, since $g \sim h$ and therefore $a \sim b$. But $a, b \in S$ and S is an independent set, which leads to a contradiction. Thus $\varphi(g) \cap \varphi(h) = \emptyset$, which means that these sets are fully disjoint and therefore adjacent in $K_{|V(Y)|:\alpha(Y)}$. Adjacent vertices in Y get mapped to adjacent vertices in the Kneser graph, resulting in φ being a homomorphism. Finally we have

$$\frac{|V(X)|}{\alpha(X)} \leq \omega^*(X) \leq \chi^*(X) = \chi^*(X(G, C)) \leq \frac{|V(Y)|}{\alpha(Y)} = \frac{|V(X)|}{\alpha(X)}.$$

□

Lemma 2.13. *Let X and Y be vertex transitive graphs such that $\chi^*(X) = \chi^*(Y)$ and let $\varphi : X \to Y$ be a homomorphism from X to Y. If S is an independent set in Y, then $\varphi^{-1}(S)$ is a maximum independent set in X.*

Proof. Theorem 2.7 implies that

$$\frac{|V(X)|}{\alpha(X)} = \frac{|V(Y)|}{\alpha(Y)}.$$

Now let f be a fractional coloring of X with weight $\chi^*(X)$ and define $g := \alpha(X)^{-1}\mathbf{1}$. As we have already seen in Lemma 2.8, g is a fractional clique of X with weight $\omega^*(X) = \chi^*(X)$. From the proof of Lemma 2.10 we can conclude that

$$(\mathbf{1}^T - g^T B)f = 0,$$

since the weight of g and f is the same. The sum of g over any independent set of size less than $\alpha(X)$ is less than 1 and therefore if we assume that $|\varphi^{-1}(S)| < \alpha(X)$, then $f(\varphi^{-1}(S)) = 0$. But Lemma 2.12 implies that X has a fractional coloring of weight $\chi^*(X)$, that has $\varphi^{-1}(S)$ in its support. Thus we get $|\varphi^{-1}(S)| = \alpha(X)$. □

2.3.6 Erdős-Ko-Rado

The Erdős-Ko-Rado theorem originates from extremal set theory and provides bounds for intersecting set families [8]. To provide an additional proof we make use of the theory we developed up until now. The equivalence to the Erdős-Ko-Rado theorem can be seen by considering that the elements of an independent set consist of r-subsets of $\{1, \ldots, n\}$.

For the proof of this theorem we make use of cyclic interval graphs. Let $\Omega = \{1, \ldots, n\}$ and $r \leq n$. Then we define the *cyclic interval graph* $C(v, r)$ as the graph with cyclical shifts of $\{1, \ldots, r\}$ modulo n as vertices and two vertices are adjacent if they are disjoint. Since vertices of $C(v, r)$ are r-sets it follows by definition that $C(v, r)$ is a subgraph of the Kneser graph $K_{v:r}$. Note that if $v < 2r$ then all vertices have at least one element in common, which means that $C(v, r)$ is empty. Therefore we assume $v \geq 2r$, similar to our assumptions for $J(v, k, i)$ graphs and thus Kneser graphs.

Lemma 2.14. *Let $v \geq 2r$ and $C(v, r)$ be a cyclic interval graph. Then the size of any independent set of $C(v, r)$ is at most r. For every independent set of size r there exists an element of $\{1, \ldots, n\}$, such that all vertices contain that element.*

Proof. Suppose that S is an independent set in $C(v, r)$. We have already shown that Kneser graphs are vertex transitive, which implies that $C(v, r)$ is vertex transitive. Automorphisms of independent sets are independent sets and coupled with vertex transitivity we can assume that the vertex $\beta = \{1, \ldots, r\}$ is in S. Let $S_1 := \{s \in S : 1 \in s\}$ and $S_r := \{s \in S : r \in s\}$ be the set of vertices in S that contain the elements 1 and r respectively. Now let j be the least integer, that is in all vertices of S_r. Since r is in all vertices of S_r, j exists. Therefore the least element of each set in S_r is at most j. Two distinct vertices of S_r have different least elements, which implies that $|S_r| \leq j$. Each vertex of S_1 has to have at least one element in common with every vertex of S_r, since they are both subsets of S. Now j is the least element that all vertices of S_r have in common, thus all cyclical shifts in S_1 also have to contain j.

This implies that the size of S_1 is at most $|S_1| \leq r - j + 1$, since there are r possible vertices in S_1, but j has to be contained in every single one. Due to our assumption, that $v \geq 2r$, we know that $S_1 \cap S_r = \{\beta\}$. Because S is independent in $C(v,r)$ and contains β it follows that

$$|S| = |S_1| + |S_r| - 1 \leq (r - j + 1) + j - 1 = r.$$

If $|S| = r$, then we know that every vertex in S has to contain j, which there are exactly r many. □

Corollary 2.4. *Let $v \geq 2r$ and $C(v,r)$ be a cyclic interval graph. Then $\chi^*(C(v,r)) = v/r$.*

Proof. Since $C(v,r)$ are vertex transitive graphs, by Theorem 2.7 and the above Lemma we get

$$\chi^*(C(v,r)) = \frac{|V(C(v,r))|}{\alpha(C(v,r))} = \frac{v}{r}.$$

□

Corollary 2.5. *Let $v \geq 2r$ and $K_{v:r}$ be the Kneser graph. Then $\chi^*(K_{v:r}) = v/r$.*

Proof. Since $C(v,r)$ is a subgraph of $K_{v:r}$, Theorem 2.7 implies that

$$\frac{v}{r} = \chi^*(C(v,r)) \leq \chi^*(K_{v:r}) \leq \frac{v}{r}.$$

□

We are now ready to prove that the independence number of Kneser graphs satisfies the Erdős-Ko-Rado bound.

Theorem 2.8 (Erdős-Ko-Rado)**.** *If $v > 2r$, then $\alpha(K_{v:r}) = \binom{v-1}{r-1}$. An independent set of size $\binom{v-1}{r-1}$ consists of r-subsets of $\Omega = \{1, \ldots, v\}$ that contain a certain element.*

Proof. Since Kneser graphs are vertex transitive, Theorem 2.7 implies that

$$\frac{v}{r} = \chi^*(K_{v:r}) = \frac{|V(K_{v:r})|}{\alpha(K_{v:r})} = \frac{\binom{v}{r}}{\alpha(K_{v:r})}.$$

This is equivalent to

$$\alpha(K_{v:r}) = \frac{r}{v}\binom{v}{r} = \binom{v-1}{r-1}.$$

Now suppose that S is an independent set of size $\binom{v-1}{r-1}$ in $K_{v:r}$. For any cyclic ordering of $\{1, \ldots, n\}$ we can construct the subgraph C of $K_{v:r}$, induced by the cyclic shifts of the first r elements of this ordering. Then C is isomorphic to $C(v,r)$ and there is an inclusion homomorphism $\iota : C \to K_{v:r}$ from C to $K_{v:r}$ mapping every element to itself. Further the preimage of S is $\iota^{-1}(S) = S \cap V(C)$ and by Lemma 2.13 it follows that $|S \cap V(C)| = r$. The vertices in $S \cap V(C)$ are exactly the cyclic shifts of some set of r consecutive elements in this ordering.

Lets consider the natural ordering $\{1, \dots, n\}$. Because of vertex transitivity we can relabel in a way such that we can assume that S contains the cyclic r sets of C of the form

$$\{1, 2, \dots, r\}, \{2, 3, \dots, r, r+1\}, \dots, \{r, r+1, \dots, 2r-1\}.$$

Since S is independent we can conclude that no vertex of the form $\{x, 1, \dots, r-1\}$ with $x \in \{2r, \dots, v\}$ is in S. Note that S contains exactly r cyclic shifts from any cyclic ordering.
All that is left to do is to show that every vertex of $K_{v:r}$ that contains r has to be in S and we do this by varying the cyclic ordering. Let g be a permutation in $Sym(v)$ that maps the subset $\{1, \dots, r-1\}$ to itself. We now consider a cyclic ordering for $x \in \{2r, \dots, v\}$ that starts with

$$\{x, 1^g, 2^g, \dots, (r-1)^g, r, (r+1)^g \dots\}.$$

Then by our previous considerations there exists a subgraph isomorphic to a cyclic interval graph, such that S has to contain $\beta = \{1^g, 2^g, \dots, (r-1)^g, r\}$, but does not contain $\{x, 1^g, 2^g, \dots, (r-1)^g\}$. Further S contains all r right shifts of β. For any r-subset γ, that contains r, there exists a cyclic ordering of the previous form, that has γ as one of these r cyclic shifts. The only exception is when γ contains all of the elements $\{2r, \dots, v\}$, since then there is no suitable choice for x.
Let $y \in \{r+1, \dots, 2r-1\}$. We now consider the natural ordering with $2r$ and y swapped

$$\{1, \dots, r, \dots, y-1, 2r, y+1, \dots, 2r-1, y, 2r+1, \dots\}.$$

Again $\beta = \{1, \dots, r\}$ is in S, but not $\{x, 1, \dots, r-1\}$ for $x \in \{y, 2r-1, \dots, v\}$, thus S contains all right shifts of β according to our ordering. With the same argument as above we conclude that all r-subsets that contain r are in S, except those that contain all of the elements of $\{y, 2r+1, \dots, v\}$. Therefore we only have to consider those that fulfill both cases and contain all the elements of $\{y, 2r, \dots, v\}$. Varying y yields that if there is any r-subset containing r that is not in S, then it has to contain all of the elements $\{r+1, \dots, v\}$. But our assumption was that $v > 2r$, thus these vertices do not exist. □

3 Spectral Graph Theory

Spectral graph theory is a mathematical field, that considers the spectrum of graphs and can be thought of as the union of graph theory and linear algebra. It has many applications in various fields as shown in the references of [21]. In this chapter we explore eigenvectors of graphs and their corresponding eigenvalues. We start by defining the adjacency matrix of a graph and its spectrum, but we will also look at other matrices that can be calculated for graphs.

The goal of this chapter is to classify the eigenvalues of Kneser graphs. It is based on chapter 8 and 9 of [12]. We will also take a look at the Laplacian of a graph, which is described in [5].

3.1 Adjacency Matrices

Let $X = (V, E)$ be a directed graph, then the *adjacency matrix* $A(X)$ of X is the matrix with rows and columns indexed by vertices of X, where $a_{x_i x_j}$ is equal to the number of arcs from x_i to x_j and 0 else. If X is an undirected simple graph, then $A(X)$ is 0 on the diagonal, symmetric and $A(X) \in \{0,1\}^{V \times V}$.

Our first observation is that isomorphic graphs may have different adjacency matrices. However given one adjacency matrix, the other can be easily obtained.

Theorem 3.1. *Let X and Y be directed graphs on the same vertex set. Then $X \cong Y$ if and only if there exists a permutation matrix P such that $P^T A(X) P = A(Y)$.*

Proof. Suppose $X \cong Y$, then the isomorphism that maps X to Y describes a permutation $\pi \in Sym(|V|)$. Thus π induces a permutation matrix $P \in \{0,1\}^{V \times V}$, such that right multiplication with P permutes the columns. If AP has permuted columns, then $(A^T P)^T = P^T A^{T^T} = P^T A$ has permuted rows. Therefore there exists a permutation matrix P such that $P^T A(X) P = A(Y)$.
Conversely if there exists a permutation matrix with these properties then the corresponding permutation is an isomorphism. □

If $X \cong Y$ then $A(X)$ is similar to $A(Y)$, since permutation matrices are orthogonal, $P^T = P^{-1}$. Every square matrix A allows the definition of its *characteristic polynomial* $\phi_A(x) = det(xI - A)$, where I is the identity matrix. Note that from now on we will mostly use x to denote variables of polynomials and use u and v for vertices of graphs.

The *spectrum* of a matrix is a list of eigenvalues coupled with their corresponding multiplicities. We define the *spectrum* of a graph X as the spectrum of its adjacency matrix $A(X)$. Eigenvectors and eigenvalues of X are the corresponding eigenvectors and eigenvalues of $A(X)$. Since similar matrices have the same eigenvalues, Theorem 3.1 implies that isomorphic graphs have the same spectrum.

Two graphs with the same spectrum are called *cospectral*, but they do not need to be

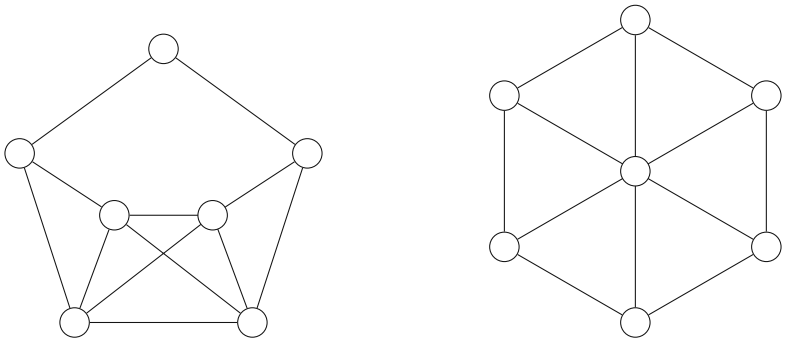

Figure 15: Two cospectral graphs

isomorphic. Consider the graphs in Figure 15. Since they share the same characteristic polynomial

$$(x+2)(x+1)^2(x-1)^2(x^2-2x-6),$$

their spectrum has to be equal. Thus the spectrum of both graphs is

$$\{-2, -1^{(2)}, 1^{(2)}, 1 \pm \sqrt{7}\}.$$

This example shows that vertex degree does not have to match for cospectral graphs. Also planarity can not be determined from looking at spectrums.

Let X be a directed graph, we call a sequence of r vertices with

$$v_0 \sim v_1 \sim \cdots \sim v_r,$$

a *walk* of length r in X. This definition is similar to a path, with the difference being, that walks are allowed to include the same vertices multiple times.

Lemma 3.1. *Let X be a directed graph with adjacency matrix A. Then $(A^r)_{uv}$ is precisely the number of walks from u to v in X with length r.*

Proof. Suppose A is the adjacency matrix for a graph X. Let $n = 1$, then A^1_{uv} is 1 if and only if u and v are adjacent. Suppose now that the theorem holds for n, we consider $A^{n+1} = A^n A$ and denote entries of A^n as a'_{ij}. Let $m = |V|$, then

$$(A^n A)_{uv} = \sum_{i=1}^{m} a'_{ui} a_{iv}.$$

By our induction hypothesis a'_{ij} denotes the number of walks from i to j of length m and a_{ij} is 1 if $i \sim j$ and 0 else. Thus we sum over all possibilities to get from u to i to v for all $i \in \{1, \ldots, n\}$. This equals the number of walks from u to v of length $n+1$. □

The *trace* of a square matrix A is the sum of its diagonal entries and we denote it by $tr(A)$. The above Lemma showed that $tr(A^r)$ is the number of walks that start and end in the same vertex.

Corollary 3.1. *Let X be a graph with $|E|$ edges and t triangles. Let A be the adjacency matrix of X, then*

1. $tr(A) = 0$,
2. $tr(A^2) = 2|E|$,
3. $tr(A^3) = t$.

Proof. There are no walks of length 1 that start and end in the same vertix for simple graphs. Walks of length 2 that start and end at the same vertex are all edges counted twice and similarly triangles are walks of length 3 with this property. □

From linear algebra [14] we know that the trace of a matrix is also equal to the sum of its eigenvalues and that the eigenvalues of A^r are the r-th powers of the eigenvalues of A. Therefore we can see that the trace of A^r is determined by the spectrum of A, which implies that the number of edges and triangles are determined by the spectrum of a graph. However this can not be easily generalized, since the graphs $K_{1,4}$ and $K_1 \cup C_4$ are cospectral, but do not share the same number of 4-cycles.

3.2 Symmetric Matrices

This section states many useful theorems regarding symmetric matrices. Many theorems that are stated here can be found in linear algebra books like [14]. If u and v are vectors of the same dimension, we also denote their scalar product $u^T v$ by $\langle u, v\rangle$.

Let A be a square matrix. Then A is *symmetric* if $A^T = A$. Symmetric matrices are especially important for graph theory, since adjacency matrices of undirected graphs are symmetric.

Lemma 3.2. *Let A be a real symmetric matrix and u and v be eigenvectors of A with different eigenvalues. Then u and v are orthogonal, meaning that $\langle u, v\rangle = 0$.*

Proof. Suppose u and v are eigenvectors with eigenvalues λ, τ, then $Au = \lambda u$ and $Av = \tau v$. Since A is symmetric we have

$$u^T A v = (v^T A^T u)^T = (v^T A u)^T.$$

Using the eigenvector property we get that

$$u^T \tau v = (v^T \lambda u)^T = u^T \lambda v,$$

which has to be 0 for $\lambda \neq \tau$. Therefore $u^T v = \langle u, v\rangle = 0$. □

Lemma 3.3. *Let A be a real symmetric matrix. Then the eigenvalues of A are real numbers.*

Proof. We have to show that the eigenvalues of A are not in $\mathbb{C} \setminus \mathbb{R}$. Let u be an eigenvector of A with eigenvalue λ. Thus $Au = \lambda u$ and if we consider the complex conjugate of this equation we get $A\bar{u} = \bar{\lambda}\bar{u}$, since A is a real matrix. Therefore $\bar{u}$ is also an eigenvector of A and $u^T \bar{u} \geq 0$, because $(a + ib)(a - ib) = a^2 + b^2$ and eigenvectors are not 0. By the previous Lemma u and $\bar{u}$ have the same eigenvalues, thus $\lambda = \bar{\lambda}$. □

Let U be a subspace of $\mathbb{R}^n$ and A a square matrix, then U is *A-invariant*, if $Au \in U$ for all $u \in U$.

Lemma 3.4. *Let $A \in \mathbb{R}^{n\times n}$ be symmetric. If U is an A-invariant subspace of $\mathbb{R}^n$, then $U^\perp := \{x \in \mathbb{R}^n : \langle x, u\rangle = 0,\ \forall u \in U\}$ is also A-invariant.*

Proof. Suppose u, v are vectors in $\mathbb{R}^n$, then

$$v^T(Au) = (Av^T)u.$$

Let $u \in U$ and $v \in U^\perp$, then $Au \in U$ implies that $v^T(Au) = 0$. By the above equation we get that $(Av^T)u = 0$. Thus $Av \in U^\perp$, which shows that $U^\perp$ is A-invariant. □

If A is a square matrix, then A has to have at least one eigenvalue, since the polynomial equation $det(xI - A) = 0$, must have at least one solution.

Lemma 3.5. *Let $A \in \mathbb{R}^{n\times n}$ be a symmetric matrix and U an A-invariant nonzero subspace of $\mathbb{R}^n$. Then U contains a real eigenvector of A.*

Proof. Since U is a subspace, there is an orthogonal basis that spans U and let R denote the matrix formed by this basis. Then we can find a square matrix B such that $AR = RB$, since A is U-invariant. Because of orthogonality we have $R^TR = I$ and

$$R^TAR = R^TRB = B.$$

Thus B is real and symmetric, because $B^T = R^TA^TR = R^TAR = B$. Every real and symmetric matrix has at least one eigenvalue. Let u be an eigenvector of B with eigenvalue λ, then $ARu = RBu = \lambda Ru$. Now $Ru \neq 0$, because $u \neq 0$ and all columns of R are linearly independent. Therefore Ru is an eigenvector of A and since A is U-invariant we also have that the eigenvector $Ru \in U$. □

Theorem 3.2. *Let $A \in \mathbb{R}^{n\times n}$ be a symmetric matrix. Then $\mathbb{R}^n$ has an orthonormal basis consisting of eigenvectors of A.*

Proof. Assume $\{u_1, \ldots, u_m\}$, with $m < n$ is a set of orthonormal eigenvectors of A. Let M be the supspace spanned by this set, then M is not 0, since A contains at least one eigenvector. Now M is A-invariant, therefore $M^\perp$ is A-invariant. Elements of $M^\perp$ are orthogonal to M and by normalizing a vector we obtain u_{m+1} such that $\{u_1, \ldots, u_m, u_{m+1}\}$ is an orthonormal set of $m+1$ eigenvectors of A. Thus by induction over n we obtain that we can construct an orthonormal basis of $\mathbb{R}^n$, consisting of eigenvectors of A. □

Corollary 3.2. *Let $A \in \mathbb{R}^{n\times n}$ be a symmetric matrix. Then there exist matrices L and D such that $L^TL = LL^T = I$ and $LAL^T = D$, and D is a diagonal matrix with entries equal to the eigenvalues of A.*

Proof. Let L be the matrix that represents an orthonormal basis of $\mathbb{R}^n$, where every row is an eigenvector of A, that exists by the previous theorem. Let $l_1, \ldots, l_n$ denote the

eigenvectors of A with eigenvalues $\lambda_1, \ldots, \lambda_n$, then the second multiplication equates to

$$AL^T = A(l_1|l_2|\ldots|l_n) = (\lambda_1 l_1|\lambda_2 l_2|\ldots|\lambda_n l_n|).$$

Thus an entry of LAL^T has the form

$$(LAL^T)_{ij} = \lambda_j \langle l_i, l_j \rangle,$$

where $\langle l_j, l_i \rangle = \delta_{ij}$, with δ_{ij} denoting the Kronecker delta, since L is an orthonormal basis. □

3.3 Eigenvectors and Eigenvalues

If we consider the adjacency matrix of a graph $X = (V, E)$, then $A(X)$ is a 01-matrix, where rows and columns are indexed by the vertex set V. Any vector u of length $|V|$ can therefore be interpreted as a function $f : V \to \mathbb{R}$ such that $v \mapsto u_v$. Thus every eigenvector can be visualized as a function that adds a label to every vertex of the graph. Consider a triangle graph T. The adjacency matrix of Figure 16, which in this

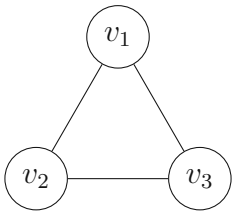

Figure 16: The triangle T

case does not depend on the order of vertices, is

$$A(T) = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}.$$

The characteristic polynomial of T is $-x^3+3x+2 = -(x+1)^2(x-2)$, which leads to the eigenvalues $\lambda_1 = -1^{(2)}$ and $\lambda_2 = 2$. An example of an eigenvector is $w := (1 \ -1 \ 0)^T$,

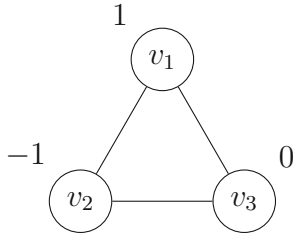

Figure 17: An eigenvector of T

depicted by adding labels to the graph which yields Figure 17. Calculating Aw yields $(-1\ 1\ 0)^T = -w$, proving that w is an eigenvector with eigenvalue -1.

There is an easy way to check if a given function is an eigenvector for a graph. Let $f : V \to \mathbb{R}$, then f can be interpreted as a vector of length $|V|$ and the image of Af of a vertex u is

$$(Af)(u) = \sum_{v \in V} A_{uv} f(v).$$

Now A is a 01-matrix that encodes adjacency and therefore we get

$$(Af)(u) = \sum_{v \sim u} f(v).$$

If f is an eigenvector of A, then $Af = \lambda f$, which results in

$$\lambda f(u) = \sum_{v \sim u} f(v). \tag{11}$$

The last equation says that the value of f at u has to be a multiple of the sum of values of its neighbors. Also any function f that satisfies this equation for all $u \in V$ necessarily has to be an eigenvector of A.

Another example is an eigenvector of the Petersen graph depicted in Figure 18. A different way to visualize an eigenvector of a graph is to write the value of a vertex as its vertex label, but we prefer to write the value to the side of the vertices. One can

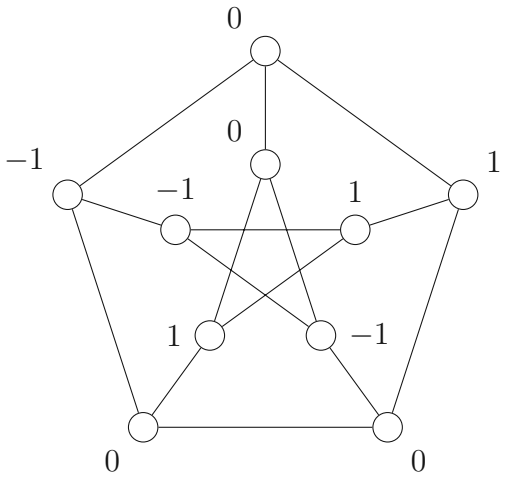

Figure 18: An eigenvector of the Petersen graph $J(5, 2, 0)$

quickly check that this eigenvector fulfills equation (11).

Let X be a graph, with $\theta_{max}(X)$ we denote the largest eigenvalue of $A(X)$ and $\theta_{min}(X)$ denotes the smallest eigenvalue of $A(X)$. Let $\lambda_1, \ldots, \lambda_m$ be the eigenvalues of a matrix A, then we define the *spectral radius* ρ of A as $\rho(A) := \max\{|\lambda_1|, \ldots, |\lambda_m|\}$. The spectral radius of a graph is the spectral radius of its adjacency matrix.

3.3.1 The Laplacian of a Graph

Closely related to the adjacency matrix of a graph is its Laplacian. Its eigenvalues prove to be very useful for other graph invariants. We follow the notation of Fan R. K. Chung [5]. Other authors usually refer to what we call Laplacian as *normalized* Laplacian of a graph. Let G be a graph and $d(v)$ denote the degree of vertex v, we begin by defining the matrix $L \in \mathbb{R}^{n\times n}$ for $u, v \in V$ as

$$L(u,v) = \begin{cases} d(v), & \text{if } u = v, \\ -1, & \text{if } u \text{ and } v \text{ are adjacent}, \\ 0, & \text{else.} \end{cases}$$

Let T denote the diagonal matrix with $T_{uu} := d(u)$ and A be the adjacency matrix of G, then we can write $L = T - A$. We define the *Laplacian* $\mathcal{L} \in \mathbb{R}^{n\times n}$ of G as

$$\mathcal{L}(u,v) = \begin{cases} 1, & \text{if } u = v, \\ -\frac{1}{\sqrt{d(u)d(v)}}, & \text{if } u \sim v, \\ 0, & \text{else.} \end{cases}$$

By defining $T_{vv}^{-1} := 0$ for $d(v) = 0$, we can write the Laplacian of G as

$$\mathcal{L} = T^{-1/2} L T^{-1/2}.$$

For easier clarification we will consider an example. Let C_5 be the cyclic graph with 5 vertices. We consider its adjacency matrix and Laplacian with their corresponding eigenvalues. Computing these objects for C_5 is comparatively easy.

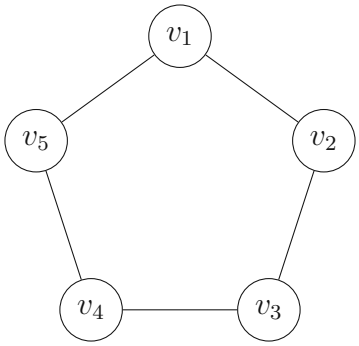

Figure 19: The cycle graph C_5

$$A(C_5) = \begin{pmatrix} 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 \end{pmatrix}.$$

The characteristic polynomial of A is

$$\chi_A(\lambda) = -\lambda^5 + 5\lambda^3 - 5\lambda + 2,$$

resulting in a spectrum of

$$\{2, \left(\frac{-1-\sqrt{5}}{2}\right)^{(2)}, \left(\frac{-1+\sqrt{5}}{2}\right)^{(2)}\}.$$

While the Laplacian of C_5 is of the form

$$\mathcal{L}(C_5) = \begin{pmatrix} 1 & -\frac{1}{2} & 0 & 0 & -\frac{1}{2} \\ -\frac{1}{2} & 1 & -\frac{1}{2} & 0 & 0 \\ 0 & -\frac{1}{2} & 1 & -\frac{1}{2} & 0 \\ 0 & 0 & -\frac{1}{2} & 1 & -\frac{1}{2} \\ -\frac{1}{2} & 0 & 0 & -\frac{1}{2} & 1 \end{pmatrix},$$

with characteristic polynomial

$$\chi_{\mathcal{L}}(\lambda) = -\lambda^5 + 5\lambda^4 - \frac{35\lambda^3}{4} + \frac{25\lambda^2}{4} - \frac{25\lambda}{16},$$

resulting in a spectrum of

$$\{0, \left(\frac{5+\sqrt{5}}{2}\right)^{(2)}, \left(\frac{5-\sqrt{5}}{4}\right)^{(2)}\}.$$

Generally the spectrum of the Laplacian of the cyclic graph C_n is $1 - \cos\left(\frac{2\pi k}{n}\right)$ for $k = 0, \ldots n-1$, also 0 is always an eigenvalue of the Laplacian [5].

3.4 Eigenvalues of Kneser graphs

Now we use the theory of eigenvectors of adjacency matrices, that we developed so far, to derive the eigenvalues of Kneser graphs. We do this by observing that certain partitions of the vertex set allow us to calculate eigenvalues for smaller graphs.

3.4.1 Equitable Partitions

Let X be a graph and let π be a partition of $V(X)$ into pairwise disjoint *cells* $C_1, \ldots, C_r$ such that $\bigcup C_i = V(X)$. A partition is *equitable* if the neighbors of $u \in C_i$ that are in C_j are a constant b_{ij}. Note that b_{ij} depends only on the cells and is independent of u. An equivalent definition is that the subgraph induced by each cell C_i is regular and the edges joining any two distinct cells form a semiregular bipartite graph, meaning that it has a proper 2-coloring such that all vertices with the same color have the same degree.
Such a partition defines a directed graph, where the cells are vertices with b_{ij} arcs from cell C_i to cell C_j. We call this directed graph the *quotient* of X over π and denote it with X/π. The adjecency matrix of the quotient of X over π is given by

$$A(X/\pi)_{ij} = b_{ij}.$$

Any automorphism group of a graph induces equitable partitions.

Lemma 3.6. *Let X be a graph and $G \subseteq Aut(X)$, then the orbits $x^G := \{x^g : g \in G\}$ of G are an equitable partition of V.*

Proof. The orbits of G partition V, since if $y \in x^G$, then $y^G = x^G$ and if $y \notin x^G$, then $y^G \cap x^G = \emptyset$. Suppose now that there exist vertices $u, v \in C_i$ in the same cell, that have a different amount of neighbors in C_j for some j. Since u and v are in the same orbit there is an automorphism g that maps u to v. Every automorphism maps adjacent vertices to adjacent vertices. Thus g has to map all neighbors of u that are in C_j to vertices in C_j, due to it being an orbit. Therefore u and v cannot have a different amount of neighbors in C_j. □

Consider the group of rotations of order 5 acting on the Petersen graph, see Figure 18. Then the 5 outer vertices are an orbit as well as the 5 inner vertices. Thus this group defines an equitable partition π_1 with adjacency matrix

$$A(X/\pi_1) = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}.$$

If π is a partition of V with r cells, then we define the *characteristic matrix* $P \in \{0,1\}^{|V| \times r}$ of π as the matrix, where columns of P are the characteristic vectors of the cells of π. Further P^TP is a diagonal matrix with entries $(P^TP)_{ii} = |C_i|$. Since π is a partition $C_i \neq \emptyset$ and therefore P^TP is invertible.

Lemma 3.7. *Let X be a graph and π be an equitable partition with characteristic matrix P. Let $B := A(X/\pi)$, then $AP = PB$ and $B = (P^TP)^{-1}P^TAP$.*

Proof. We start by considering the uj-entry of AP. Then $(AP)_{uj}$ is the number of neighbors of the vertex u that lie in cell C_j. If $u \in C_i$ then $(AP)_{uj} = b_{ij}$. If we take a look at $PA(X/\pi)$, we notice that every vertex can only lie in one cell and thus every row of $PA(X/\pi)$ has exactly one 1. Therefore this uj-entry also equals b_{ij} and we get

$$(AP)_{uj} = (PA(X/\pi))_{uj}.$$

Thus we have $AP = PB$ and left multiplying this equation with P^T yields

$$P^TAP = P^TPB.$$

Since P^TP is invertible this holds if and only if

$$B = (P^TP)^{-1}P^TAP.$$

□

This next theorem will provide the foundation to find the eigenvalues of Kneser graphs.

Theorem 3.3. *Let X be a vertex transitive graph, $G \subseteq Aut(X)$ a permutation group and π the equitable partition defined by orbits of G. If π has a singleton cell $\{u\}$, then every eigenvalue of X is an eigenvalue of the quotient X/π.*

Proof. Let f be a function on $V(X)$ and let $g \in Aut(X)$, then we define f^g as

$$f^g(x) := f(x^g).$$

Suppose f is an eigenvector of X. If g is an automorphism, then $A(X) = A(X^g)$, because $A(X)_{ij} = 1$ if and only if $i \sim j$ which happens if and only if $i^g \sim j^g$. Additionally, since $X \cong X^g$, we know that there exist permutation matrices, such that $A(X)$ and $A(X^g)$ are similar. Let Q be the permutation matrix induced by g, then

$$Af^g = AQf = QAQ^TQf = QAf = \lambda Qf = \lambda f^g,$$

which shows that f^g is an eigenvector of X with the same eigenvalue as f.
Now we define $\hat{f}$ as

$$\hat{f} := \frac{1}{|G|}\sum_{g\in G} f^g,$$

which is constant on the cells of π, since they are invariant under G. If $\hat{f}$ is not zero, then it is also an eigenvector of X with eigenvalue λ. Suppose that h is an eigenvector of X with eigenvalue λ. Since $h \neq 0$ there exists a vertex $v \in V(X)$, such that $h(v) \neq 0$. Let $g \in Aut(X)$ such that $u^g = v$ and set $f = h^g$. Then $f(u) \neq 0$ and we get

$$\hat{f}(u) = f(u) \neq 0.$$

This shows that $\hat{f} \neq 0$ and $\hat{f}$ is a candidate for an eigenvector of X/π with eigenvalue λ.
Let P be the characteristic matrix of π and let v be an eigenvector of $A(X/\pi)$. Lemma 3.7 implies that $AP = PA(X/\pi)$ and therefore we get

$$APv = PA(X/\pi)v = \theta Pv.$$

On the other hand, $AP = PA(X/\pi)$ further implies that the column space of P is A-invariant [9]. Thus P must have a basis consisting of eigenvectors of A. Now every eigenvector of A is constant on the cells of P and therefore has the form Pv, with $v \neq 0$. If $APv = \theta Pv$, then also $A(X/\pi)v = \theta v$, which implies that an eigenvalue of X is also an eigenvalue of X/π. □

Lemma 3.8. *The following binomial equation holds,*

$$\sum_{i=0}^{h}(-1)^{h-i}\binom{h}{i}\binom{a-i}{k} = (-1)^h\binom{a-h}{k-h}.$$

Proof. We start by defining

$$f(a,h,k) := \sum_{i=0}^{h}(-1)^{h-i}\binom{h}{i}\binom{a-i}{k}.$$

We can use a well known recursion for the binomal term in f,

$$\binom{a-i}{k} = \binom{a-i-1}{k} + \binom{a-i-1}{k-1}, \qquad 1 \leq k \leq a-i-1,$$

to obtain the following recursion for f,

$$f(a,h,k) = f(a-1,h,k) + f(a-1,h,k-1).$$

Now $f(a,h,0) = 0$ if $h > 0$ holds, because of the symmetry of binomial coefficients. Further we have $f(a,0,0) = 1$ and $f(k,h,k) = (-1)^h$, since both only sum over $i = 0$. Induction over a yields

$$\begin{aligned} f(a,h,k) &= f(a-1,h,k) + f(a-1,h,k-1) \\ &= (-1)^h \binom{a-1-h}{k-h} + (-1)^h \binom{a-1-h}{k-1-h} \\ &= (-1)^h \left(\binom{a-h-1}{k-h} + \binom{a-h-1}{k-h-1} \right) \\ &= (-1)^h \binom{a-h}{k-h}. \end{aligned}$$

□

Finally we can characterize the eigenvalues of the Kneser graphs.
Suppose $h(i,j)$ is a function of i and j and $0 \leq i,j \leq r$, then we define $[h(i,j)]$ as the $(r+1) \times (r+1)$ matrix with ij-entry $h(i,j)$.

Theorem 3.4. *The eigenvalues of the Kneser graph $K_{v:r}$ are the integers*

$$(-1)^i \binom{v-r-i}{r-i}, \qquad i \in \{0,1,\ldots,r\}.$$

Proof. Suppose $v \geq 2r$ and let X be the Kneser graph $K_{v:r}$, where vertices are r-subsets of $\Omega = \{1,\ldots,n\}$. Let $\alpha := \{1,\ldots,r\}$ and let $C_i \subseteq 2^\Omega$ be the set r-subsets that meet α in exactly $r-i$ elements. This defines a partition π consisting of the cells $C_0,\ldots,C_r$. Let $G \subseteq Sym(\Omega)$ such that $g \in G$ fixes α setwise. Then the orbits of G are the cells C_i, since the elements in the intersection with α of these vertices are the same. Thus π is an equitable partition, that satisfies the requirements of Theorem 3.3. Therefore if we can determine the eigenvalues of $A(X/\pi)$ then we know that all eigenvalues of X also appear. Suppose $\beta \in V(K_{v:r})$ and $|\beta \cap \alpha| = r-i$. The ij-entry of $A(X/\pi)$ is the number of r-subsets of Ω that are disjoint from β and meet α in exactly $r-j$ points, since π is an equitable partition. Thus we get

$$A(X/\pi)_{ij} = \binom{i}{r-j}\binom{v-r-i}{j}, \qquad 0 \leq i,j \leq r.$$

Let D be the diagonal matrix with entries

$$D_{ii} := (-1)^i \binom{v-r-i}{r-i}.$$

We will prove the following equality:

$$\left[(-1)^{i-j}\binom{i}{j}\right] A(X/\pi) \left[\binom{i}{j}\right] = D\left[\binom{r-i}{r-j}\right]. \tag{12}$$

If we assume that equation (12) holds, then A is similar to the product $D[\binom{r-i}{r-j}]$, since

$$\left[(-1)^{i-j}\binom{i}{j}\right]^{-1} = \left[\binom{i}{j}\right].$$

Now $[\binom{r-i}{r-j}]$ is an upper triangle matrix, with diagonal entries equal to 1 and therefore the eigenvalues of $A(X/\pi)$ have to be the diagonal entries of D.
All that is left to do is to prove equation (12). If we consider the ik-entry of $A(X/\pi)[\binom{i}{j}]$, we get

$$\left(\left[\binom{i}{r-j}\binom{v-r-i}{j}\right]\left[\binom{i}{j}\right]\right)_{ik} = \sum_{j=0}^{r}\binom{v-r-i}{j}\binom{i}{r-j}\binom{j}{k}.$$

Using a binomial identity,

$$\binom{a}{c}\binom{a-c}{b-c} = \frac{a!(a-c)!}{c!(a-c)!(b-c)!(a-c-b+c)!}\cdot\frac{b!}{b!} = \binom{a}{b}\binom{b}{c}, \tag{13}$$

we can calculate for the sum above:

$$\begin{aligned}&\sum_{j=0}^{r}\binom{v-r-i}{j}\binom{j}{k}\binom{i}{r-j}\\ =&\sum_{j=0}^{r}\binom{v-r-i}{k}\binom{v-r-k-i}{j-k}\binom{i}{r-j}\\ =&\binom{v-r-i}{k}\sum_{j=0}^{r}\binom{v-r-k-i}{j-k}\binom{i}{r-j}.\end{aligned}$$

We apply Vandermonde's identity

$$\binom{a+b}{n} = \sum_{j=0}^{n}\binom{a}{j}\binom{b}{n-j},$$

and an index shift, to simplify the last term of the equation to

$$\binom{v-r-i}{k}\binom{v-r-k}{r-k}.$$

With this partial solution we can calculate that the hk-entry of the product

$$\left[(-1)^{i-j}\binom{i}{j}\right]A(X/\pi)\left[\binom{i}{j}\right]$$

equals

$$\binom{v-r-k}{r-k}\sum_{i=0}^{r}(-1)^{h-i}\binom{h}{i}\binom{v-r-i}{k}.$$

Lemma 3.8 implies that the sum is equal to

$$(-1)^h \binom{v-r-k}{r-k} \binom{v-r-h}{k-h},$$

which can be simplified, using (13) and setting $a = v-r-h$, $b = r-h$ and $c = k-h$, to

$$(-1)^h \binom{v-r-h}{r-h} \binom{r-h}{r-k}.$$

But this is equal to

$$\left(D \left[\binom{r-i}{r-j} \right] \right)_{hk}.$$

□

4 EKR-Graphs

In this chapter we explore graphs that satisfy a certain Erdős-Ko-Rado property. We define what it means for a graph to be EKR, by generalizing the inequality of the Erdős-Ko-Rado Theorem. This chapter is based in large parts on the work of Holroyd and Talbot [16].

As we have already seen the sizes of independent sets on Kneser graphs are bounded by theorem 1.1. In this chapter the Erdős-Ko-Rado theorem will appear in a different context. Rather than looking at the largest independent set, we will look at the number of certain independent sets that all have a fixed size. We start by declaring more definitions that help us talk about these concepts.

Let G be a graph and $k \in \mathbb{N}^+$, then $\mathcal{I}^{(k)}(G)$ denotes the family of independent sets of G of size k. If $v \in V$, then $\mathcal{I}_v^{(k)}(G)$ denotes the family of independent sets of size k that contain the vertex v. The family $\mathcal{I}_v^{(k)}(G)$ is called a *k-star*, or just *star*, where the vertex v is called its *center*.

Definition 4.1. *A graph G is k-EKR if no intersecting subfamily of $\mathcal{I}^{(k)}(G)$ is bigger than the largest star.*

G is *strictly k-EKR* if every maximum size intersecting subfamily of $\mathcal{I}^{(k)}(G)$ is a k-star. This definition does indeed correlate with the Erdős-Ko-Rado Theorem, since the empty graphs are k-EKR.

Corollary 4.1. *The empty graph E_n is k-EKR for $2k \leq n$.*

Proof. The empty graph $E_n = (V, \emptyset)$ is isomorphic to the set $\Omega = \{1, \ldots, n\}$. Every subset of V is also an independent set, thus independent sets in $\mathcal{I}^{(k)}(G)$ are translated to k-subsets of Ω. To construct a largest star of E_n simply fix one out of n elements and choose $k-1$ other elements to generate a member. Thus the size of the largest star equals the bound of the Erdős-Ko-Rado Theorem 1.1. Now we can apply the theorem to see that every intersecting independent set has to be smaller than or equal to the size of the largest star. □

The empty graph being k-EKR can be seen as the baseline Erdős-Ko-Rado Theorem for sets. Thus being k-EKR is in a way a generalization of the Erdős-Ko-Rado Theorem to more complicated structures. To better understand this definition, we now consider some basic facts about k-EKR graphs.

Lemma 4.1. *Every graph is 1-EKR.*

Proof. Since there are no intersecting independent 1-sets, other than a trivial single element set, the condition is fulfilled. □

For bigger k this quickly becomes non-trivial. For $k = 2$ however there still is an easy way to calculate whether a graph is 2-EKR or not. Recall that the independence number $\alpha(G)$ is the size of the largest possible independent set for a graph G.

Theorem 4.1. *Let G be a non-complete graph, with minimum degree δ and $|V| = n$.*

1. *If $\alpha(G) = 2$, then G is strictly 2-EKR.*
2. *If $\alpha(G) \geq 3$, then G is 2-EKR if and only if $\delta \leq n - 4$. G is strictly 2-EKR if and only if $\delta \leq n - 5$.*

Proof. Let $\mathcal{A} \subseteq \mathcal{I}^{(2)}(G)$ be a family, that is not a subfamily of any star. We assume that the cardinality of $\mathcal{A}$ is at least 3, meaning that $\mathcal{A}$ contains at least 3 2-sets of independent vertices. If we denote them by v_i we know that $\{v_1, v_2\}$, $\{v_1, v_3\}$ and $\{v_2, v_3\} \in \mathcal{A}$. However no other 2-set can be in $\mathcal{A}$, since it has to intersect all of the other 3 2-sets. Therefore we know that $|\mathcal{A}| = 3$.
If we assume that $\alpha(G) = 2$, then an independent set of size 3 does not exist. If we assume $\alpha(G) \geq 3$, then families, that are not subfamilies of a star have size 3. However for any vertex v a star centered at v has a size of

$$|\mathcal{I}_v^{(2)}(G)| = n - 1 - d(v),$$

which proves the theorem. □

Next we will look at non-trivial examples of EKR graphs and also prove that lexicographic products of EKR graphs and complete graphs are again EKR graphs.

4.1 Trees and Claws

In this section we will prove EKR properties for a subfamily of trees. The complete bipartite graph $K_{1,3}$ is also called the *claw* graph, while $K_{1,n}$ is called a *claw*. An *elongated claw* is a rooted tree, where no non-root vertex has degree greater than 2. A *depth-two* claw is an elongated claw where the distance between every leaf to the root is 2. A *maximal* independent set of G is an independent set, that is not itself a proper subset of an independent set. We define $\mu(G)$ as the minimum size of a maximal independent set of G, some authors call it the *minimax independence number* of G. Note that these sizes can be equal, but in general $\mu(G) \leq \alpha(G)$ holds.

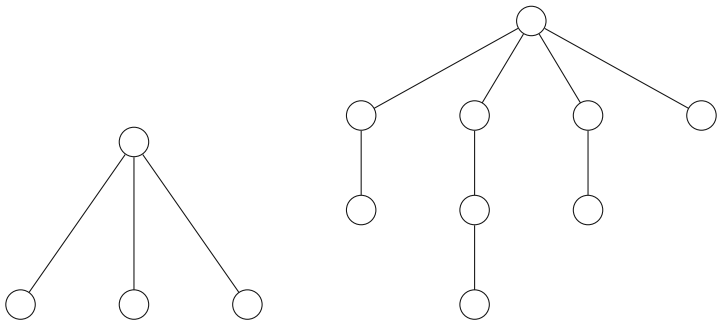

Figure 20: The claw graph $K_{1,3}$ and an elongated claw.

We will prove that if a graph G is a depth-two claw, then G is k-EKR and strictly

k-EKR, if $2k < \mu(G)$. Further if G is an elongated claw with n leaves and at least one leaf is adjacent to the root and $2k \leq n$, then G is k-EKR. This result is based on the work from Feghali, Johnson and Thomas [10].

We start with some easier observations about claws.

Theorem 4.2. *Let $n, k \in \mathbb{N}^+$, then the claw $K_{1,n}$ is k-EKR, if $2k \leq n$ and strictly k-EKR if $2k < n$.*

Proof. Independent sets of claws are easy to characterize. If the root is contained, then no other vertex can be contained. Otherwise independent sets can consist of up to n leafs. Any largest star therefore has to have a leaf as the center. Further we can delete the root vertex, since it cannot be part of any intersecting independent family, other than the set consisting of just this vertex. We get n vertices with no edges, which is isomorphic to the empty graph E_n. □

As we have seen, proving that certain graph classes are k-EKR usually amounts to finding a largest star.

4.1.1 Depth-two Claws

We again start by proving a few lemmata, which will help us with the final proof. A path from the root to a leaf is called a *limb* and we call it *short* limb if it only contains one edge.

Lemma 4.2. *Let G be an elongated claw and let $k \in \mathbb{N}^+$. Then there exists a largest k-star of G, with a leaf as the center.*

Proof. Let $v \in V$ be a vertex that is not a leaf and let L be the limb of G that contains v. In case v is the root, L can be any limb. Assume x is the unique leaf in L, we show that

$$|\mathcal{I}_v^{(k)}(G)| \leq |\mathcal{I}_x^{(k)}(G)|,$$

which proves this lemma. We do this by constructing an injection $f : \mathcal{I}_v^{(k)}(G) \to \mathcal{I}_x^{(k)}(G)$. Assume w is the unique neighbor of the leaf x and let $A \in \mathcal{I}_v^{(k)}(G)$. Then construct f as follows.

1. If $x \in A$ define $f(A) := A$.
2. If $x \notin A$ and $w \notin A$, then define $f(A) := (A \setminus \{v\}) \cup \{x\}$.
3. If $x \notin A$ and $w \in A$, then let $X := \{x = x_1, x_2, \ldots, x_m = v\}$ be the set of vertices of L from x to v. We define $Y := A \cap X$ and denote its elements with $Y = \{x_{i_1}, \ldots, x_{i_j}\}$, for some $1 \leq j < m$. Let $Z = \{x_{i_1-1}, \ldots, x_{i_j-1}\}$. Then $|Y| = |Z|$ and $x \in Z$, since $w \in Y$ and w was defined to be the unique neighbor of the leaf x. Then we define $f(A) := (A \cup Z) \setminus Y$.

Let us now consider $f(A_1)$ and $f(A_2)$ for distinct $A_1, A_2 \in \mathcal{I}_v^{(k)}(G)$. If they fall into the same case, then it is easy to verify that their images also have to be distinct. If they are part of different cases, then note that case 1 always contains v and x. Case 2 does contain x but never v, or any of its neighbors, because they could not have been in A originally. Images from case 3 do not contain v, but at least one neighbor of v. Thus the images of A_1 and A_2 have to be distinct, proving that f is injective. Also f maps independent sets to independent sets. □

This is a very useful theorem, since it allows us to assume that the center of a largest star is at a leaf, simplifying the calculation of its size.

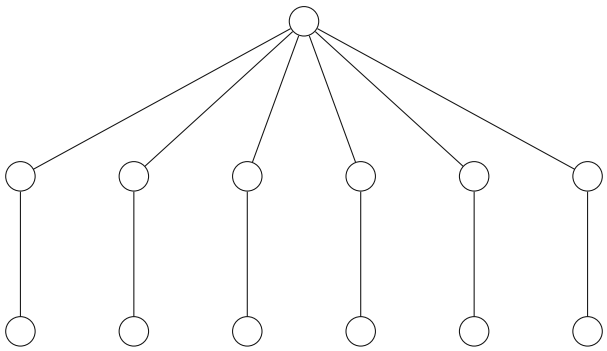

Figure 21: A depth-two claw with 6 leafs.

Lemma 4.3. *Let $k, n \in \mathbb{N}^+$ with $k \leq n$ and let G be a depth-two claw with n leaves. Then the largest k-star of G has size*

$$\binom{n-1}{k-1}2^{k-1} + \binom{n-1}{k-2},$$

and is centered at a leaf.

Proof. The previous lemma guarantees us that a largest star centered at a leaf can be found. Since limbs in depth-two claws are isomorphic, we can assume that the star is centered at any leaf. Let v be a leaf of G and let r be the root of G. We define a partition of $\mathcal{I}_v^{(k)}(G)$ into $\mathcal{B} = \{A \in \mathcal{I}_v^{(k)}(G) : r \notin A\}$ and $\mathcal{C} = \{A \in \mathcal{I}_v^{(k)}(G) : r \in A\}$. Now we can calculate that

$$|\mathcal{B}| = \binom{n-1}{k-1}2^{k-1},$$

since each set in $\mathcal{B}$ contains v and intersects $k-1$ of the $n-1$ other limbs. For each of those $k-1$ limbs there are 2 possible choices to add a vertex, because the root is not contained. Every set in $\mathcal{C}$ has to contain v and the root r, thus we can choose $k-2$ leafs out of the $n-1$ possible leafs. Therefore we can calculate the size of $\mathcal{C}$ to be

$$|\mathcal{C}| = \binom{n-1}{k-2}.$$

□

Next we state a theorem that will help us prove what we proposed. Let $\mathcal{A}$ be a family of sets and $s \in \mathbb{N}^+$, then the *s-shadow* of $\mathcal{A}$ is defined as

$$\partial_s \mathcal{A} := \{S : |S| = s, \exists A \in \mathcal{A} \wedge S \subseteq A\}.$$

Lemma 4.4 (Katona [19])**.** *Let $a, b \in \mathbb{N}^+$ and let $\mathcal{A}$ be a family of sets of size a such that $|A \cap A'| \geq b \geq 0$ for all $A, A' \in \mathcal{A}$. Then $|\mathcal{A}| \leq |\partial_{a-b}\mathcal{A}|$.*

This lemma is actually a special case of Theorem 2 in [19]. To see this set $g = l - k$. Then, using Katona's notation, we get

$$|\mathcal{A}| \leq \frac{\binom{2l-k}{l-k}}{\binom{2l-k}{l}} |\mathcal{A}^{l-k}| = |\mathcal{A}^{l-k}|.$$

With this we can prove the proposed theorem about depth-two claws.

Theorem 4.3. *Let $k \in \mathbb{N}^+$ and let G be a depth-two claw, with $\mu(G) \geq 2k - 1$. Then G is strictly k-EKR.*

Proof. Let r denote the root of G and assume that G has n leafs and $\mu(G) \geq 2k - 1$. Now $n = \mu(G)$, since a maximal independent set can be constructed by taking the unique neighbor of every leaf. This set is maximal and any other maximal set cannot be smaller, since then you could add a vertex of a missing limb. Thus we have $n \geq 2k-1$. Let $\mathcal{A} \subseteq \mathcal{I}^{(k)}(G)$ be any intersecting family. We define a partition, similarly to Lemma 4.3, of $\mathcal{A}$ into $\mathcal{B} := \{A \in \mathcal{A} : r \notin A\}$ and $\mathcal{B} := \{A \in \mathcal{A} : r \in A\}$.
Now any vertex $v \in A \in \mathcal{B}$ has to be either a leaf or a neighbor of a leaf. Since every set in $\mathcal{B}$ has k elements, it intersects exactly k distinct limbs. Thus for $B \in \mathcal{B}$ we can define M_B as the set of k leafs that either belong to B or are adjacent to a vertex in B. We say that M_B *represents* B and define $\mathcal{M} := \{M_B : B \in \mathcal{B}\}$. Since $\mathcal{B}$ is intersecting we have that $\mathcal{M}$ is also intersecting. Note that any M_B can represent multiple different B, since we have two choices of vertices to include for each limb. In theory any M_B can represent 2^k many distinct B, however since $\mathcal{B}$ is an intersecting family the number of $B \in \mathcal{B}$ that are represented by M_B is at most 2^{k-1}. Let s_M be the number of sets in $\mathcal{B}$ represented by M. Since $\mathcal{M}$ is an intersecting family where every set consists of exactly k elements, the Erdős-Ko-Rado Theorem implies that $|\mathcal{M}| \leq \binom{n-1}{k-1}$. Thus we get the inequality

$$|\mathcal{B}| = \sum_{M \in \mathcal{M}} s_M \leq \binom{n-1}{k-1} 2^{k-1}.$$

To find a bound of $\mathcal{C}$, we first define N_B to be the set of $n - r$ leafs that are neither in B nor adjacent to a vertex in B for $B \in \mathcal{B}$. Then M_B and N_B form a partition of the leafs of our graph G. Let $\mathcal{N} := \{N_B : B \in \mathcal{B}\}$. Since $\mathcal{M}$ is intersecting we know that M_{B_1} and M_{B_2} also have non-empty intersection for any $B_1, B_2 \in \mathcal{B}$. Therefore we get $|M_{B_1} \cup M_{B_2}| \leq 2k - 1$. All leafs not in this union are members of N_{B_1} and N_{B_2} and there are at least $n - (2k - 1) > 0$ many of them. This implies that $\mathcal{N}$ is also intersecting and we can apply Lemma 4.4 with $a = n - r$, $b = n - (2k - 1)$ to get

$$|\mathcal{N}| \leq |\partial_{k-1}\mathcal{N}|.$$

Every set $S \in \partial_{k-1}\mathcal{N}$ is a collection of sets of $k-1$ leafs such that $S \subseteq N_B$, for some $B \in \mathcal{B}$. Thus S is disjoint to M_B and also disjoint to B.
Let $C \in \mathcal{C}$. Since the root vertex $r \in C$ and $|C| = k$, we know that C contains $k-1$ leafs. Because $\mathcal{A}$ is intersecting we know that C has non-empty intersection with every $B \in \mathcal{B}$. By our above observation C is not a member of $\partial_{k-1}\mathcal{N}$. This yields an upper bound of

$$|\mathcal{C}| \leq \binom{n}{k-1} - |\partial_{k-1}\mathcal{N}|.$$

Since $|\mathcal{N}| \leq |\partial_{k-1}\mathcal{N}|$ we get

$$|\mathcal{C}| \leq \binom{n}{k-1} - |\mathcal{N}|.$$

We already showed that every $B \in \mathcal{B}$ defines M_B and N_B, which partition the leafs of G. Therefore we have $|\mathcal{M}| = |\mathcal{N}|$, which results in

$$|\mathcal{C}| \leq \binom{n}{k-1} - |\mathcal{M}|.$$

Finally we can add our bounds, to get an upper bound for $|\mathcal{A}|$.

$$\begin{aligned}
|\mathcal{A}| &= |\mathcal{B}| + |\mathcal{C}| \\
&\leq |\mathcal{M}|2^{k-1} + \binom{n}{k-1} - |\mathcal{M}| \\
&\leq \binom{n-1}{k-1}2^{k-1} + \binom{n}{k-1} - \binom{n-1}{k-1} \\
&= \binom{n-1}{k-1}2^{k-1} + \binom{n-1}{k-2}.
\end{aligned}$$

Note that strict inequality applies if $n > 2k$. If we apply Lemma 4.3 we get that the size of an independent set is always less than or equal to the size of the largest k-star of G, which is what we wanted to prove. □

One of our assumptions was that $n = \mu(G) \geq 2k-1$. For $k = n$, we will show that G is not n-EKR, by constructing an intersecting family of independent sets, with bigger size than the largest n-star of G. Let r be the root of G then $G' := G - r$ is the graph obtained from G by removing the root. Thus G' is isomorphic to n copies of K_2. Now G' contains 2^n independent sets of size n, that can be partitioned into complementary pairs. For each such pair we choose the independent set that contains more leafs. If a pair has an equal number of leafs, then we choose one arbitrarily. We define $\mathcal{B}$ to be the family of exactly these sets. We picked exactly half of the independent sets, thus we know that $|\mathcal{B}| = 2^{n-1}$. Further $\mathcal{B}$ has to be an intersecting family, since every set contains n vertices and more than $n/2$ leafs, while the leaf partitioning partner is not contained. But $\mathcal{B}$ is not a star. Now we define $\mathcal{C} = \{C \in \mathcal{I}^{(n)}(G) : r \in C\}$. Every $C \in \mathcal{C}$ contains the root and therefore contains $n-1$ leafs. Therefore $|\mathcal{C}| = \binom{n}{n-1} = n$

and for every $C \in \mathcal{C}$ and $B \in \mathcal{B}$ we have $C \cap B \neq \emptyset$ and $B \neq C$. Let $\mathcal{A} = \mathcal{B} \cup \mathcal{C}$, then $\mathcal{A}$ is intersecting and

$$|\mathcal{A}| = |\mathcal{B}| + |\mathcal{C}| = 2^{n-1} + n.$$

But this is larger than the largest n-star of G which, by Lemma 4.3, has size equal to

$$\binom{n-1}{n-1}2^{n-1} + \binom{n-1}{n-2} = 2^{n-1} + n - 1.$$

4.1.2 Elongated Claws with Short Limbs

In this section we consider elongated claws. Contrary to depth-two claws, a limb can now consist of more than two vertices, making the proofs slightly more complex. Our goal is to show that an elongated claw with a short limb is k-EKR, if $n \geq 2k$.

Let $G = (V, E)$ be a graph and $v \in V$, we define $G - v$ to be the graph obtained by deleting v and all of its incident edges from G. We define $G \downarrow v$ to be the graph obtained by deleting v and all of its adjacent vertices from G with all of their incident edges.

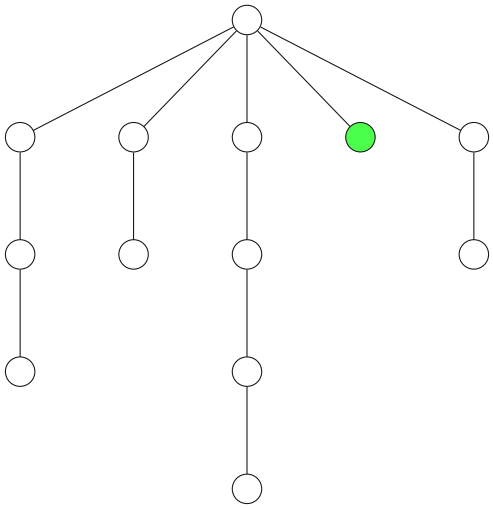

Figure 22: An elongated claw with a short limb.

We start by stating some lemmata that will help us with the main theorem.

Lemma 4.5. *Let $k \in \mathbb{N}^+$ and let $G = (V, E)$ be a graph. Let $v \in V$ and let u be a vertex of $G \downarrow v$. Then*

$$|\mathcal{I}_u^{(k)}(G)| = |\mathcal{I}_u^{(k)}(G - v)| + |\mathcal{I}_u^{(k-1)}(G \downarrow v)|.$$

Proof. We define a partition of $\mathcal{I}_u^{(k)}(G)$ into $\mathcal{B} = \{A \in \mathcal{I}_u^{(k)}(G) : v \notin A\}$ and $\mathcal{C} = \{A \in \mathcal{I}_u^{(k)}(G) : v \in A\}$. This is in fact a partition. But now we get that $\mathcal{B} = \mathcal{I}_u^{(k)}(G - v)$. On the other hand $|\mathcal{C}| = |\mathcal{I}_u^{(k-1)}(G \downarrow v)|$, since no neighbor of v can be in the independent set and v is fixed. Thus we get $|\mathcal{I}_u^{(k)}(G)| = |\mathcal{B}| + |\mathcal{C}|$. □

Lemma 4.6. *Let $k \in \mathbb{N}^+$ and let $G = (V, E)$ be an elongated claw with a short limb and root r. If $x \in V$ is a leaf that is adjacent to r, then x is the center of a largest k-star of G.*

Proof. Let $v \in V$ be a non-leaf that is adjacent to r. We have to show that $|\mathcal{I}_v^{(k)}(G)| \leq |\mathcal{I}_x^{(k)}(G)|$. Suppose $v = r$, then these sizes are equal since you could swap r and v to get the other independent set.
Otherwise let L be the limb of G that contains v. We will now construct an injection $f : \mathcal{I}_v^{(k)}(G) \to \mathcal{I}_x^{(k)}(G)$, which proves the statement. Let $A \in \mathcal{I}_v^{(k)}(G)$, then we define the image of A as:

1. If $x \in A$, then $f(A) := A$.
2. If $x \notin A$ and $r \notin A$, then $f(A) := (A \setminus \{v\}) \cup \{x\}$.
3. If $x \notin A$ and $r \in A$, let $X = \{v = x_1, \ldots, x_m\}$ be the set of vertices from v to x_m, the unique neighbor of r in L. Let $Y = A \cap X = \{x_{i_1}, \ldots, x_{i_j}\}$, for some $1 \leq j < m$. Let $Z = \{x_{i_1+1}, \ldots, x_{i_j+1}\}$. Then $|Y| = |Z|$ and $Y \cap Z = \emptyset$. We define $f(A) := (A \cup Z \cup \{x\}) \setminus (Y \cup \{r\})$.

This function is indeed an injection, which proves the lemma. □

Now we can prove the proposed theorem for elongated claws with short limbs.

Theorem 4.4. *Let $n, k \in \mathbb{N}^+$, $2k \leq n$ and let G be an elongated claw with n leafs and a short limb. Then G is k-EKR.*

Proof. Let r be the root of G. Let $\mathcal{A} \subseteq \mathcal{I}^{(r)}(G)$ be any intersecting family. We have to show that the size of $\mathcal{A}$ is less than or equal to the size of the largest k-star. We do this by induction over k. For $k = 1$ this holds trivially, since an intersecting family, where every member is a single element set, can only have a single member.
Suppose now the theorem holds for $k - 1$, we want to show that it also holds for k. Again we use induction over the number vertices $|V|$ of G. Suppose $G = K_{1,n}$ is a claw. By Theorem 4.2 the claw $K_{1,n}$ is k-EKR if $2k \leq n$. Thus let $|V| \geq n + 2$ and suppose that the result is true for elongated claws with fewer vertices.
Let x be a leaf adjacent to the root r and v be a leaf not adjacent to r. Let w denote the unique neighbor of v and z the other neighbor of w.
The situation is depicted in Figure 23. Now we define a function $f : \mathcal{A} \to \mathcal{I}^{(k)}(G)$ that maps each $A \in \mathcal{A}$ to

$$f(A) = \begin{cases} (A \setminus \{v\}) \cup \{w\}, & \text{if } v \in A, z \notin A, (A \setminus \{v\}) \cup \{w\} \notin \mathcal{A}, \\ A, & \text{else.} \end{cases}$$

This function effectively swaps v with w for an independent set, if independence is preserved. Note that $|\mathcal{A}| = |f(\mathcal{A})|$, since f is injective by construction. Now we define three families:

- $\mathcal{A}' := \{f(A) : A \in \mathcal{A}\}$,

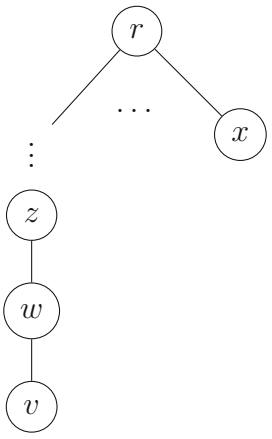

Figure 23: G with both leafs v and x and their unique neighbors.

- $\mathcal{B} := \{A : v \notin A, A \in \mathcal{A}'\}$,
- $\mathcal{C} := \{A \setminus v : v \in A, A \in \mathcal{A}'\}$.

By the definition of these families and f we get

$$|\mathcal{A}| = |\mathcal{A}'| = |\mathcal{B}| + |\mathcal{C}|. \tag{14}$$

Lemma 4.7. *Both $\mathcal{B}$ and $\mathcal{C}$ are intersecting families.*

Proof. Our claim is that both families $\mathcal{B}$ and $\mathcal{C}$ are intersecting. We start by defining a partition of $\mathcal{B}$ into $\mathcal{B}_1 := \{B \in \mathcal{B} : B \in \mathcal{A}\}$ and $\mathcal{B}_2 := \{B \in \mathcal{B} : (B \setminus \{w\}) \cup \{v\} \in \mathcal{A}\}$. This is indeed a partition, since every member of $\mathcal{B}_1$ is a fixpoint of f that does not contain v and every member of $\mathcal{B}_2$ represents a set that is not a fixpoint of f. Now $\mathcal{B}_1$ is an intersecting family, because $\mathcal{A}$ is intersecting and every member of $\mathcal{B}_2$ has to contain w, thus $\mathcal{B}_2$ is also intersecting. We want to show that $\mathcal{B}$ is intersecting. Therefore let $B_1 \in \mathcal{B}_1$ and $B_2 \in \mathcal{B}_2$. Since B_1 and $(B_2 \setminus \{w\}) \cup \{v\}$ are both in $\mathcal{A}$, they intersect. But v is not part of that intersection, because v is not in B_1. Thus we get

$$\emptyset \neq B_1 \cap ((B_2 \setminus \{w\}) \cup \{v\}) \subseteq B_1 \cap B_2,$$

showing that B_1 and B_2 intersect, which proves that $\mathcal{B}$ is intersecting.
Now we show that $\mathcal{C}$ is intersecting. Let $C \in \mathcal{C}$, then $C \cup \{v\} \in \mathcal{A}'$ by definition and also $C \cup \{v\} \in \mathcal{A}$, since every $A \in \mathcal{A}'$ that contains v has to be a fixpoint of f. Now either $C \cup \{w\}$ is in $\mathcal{A}$ or z has to be in C. Let $C_1, C_2 \in \mathcal{C}$, then there are the following cases. Either both contain z or both contain $\{w\}$ or one contains the element that is not in the other. Applying the previous consideration and assuming C_1 does not contain z, then $C_1 \cup \{w\}$ is in $\mathcal{A}$. But $C_2 \cup \{v\}$ is also in $\mathcal{A}$, which is intersecting. Thus we get that $C_1 \cup \{w\}$ and $C_2 \cup \{v\}$ have non-empty intersection. Since all members in $\mathcal{A}$ are independent, this intersection contains neither v nor w. Therefore we get

$$\emptyset \neq (C_1 \cup \{w\}) \cap (C_2 \cup \{v\}) = C_1 \cap C_2,$$

which proves that $\mathcal{C}$ is also intersecting. □

Coming back to the proof, we note that $G - v$ is an elongated claw with a short limb, fewer vertices than G and still n leafs. Since every set in $\mathcal{A}$ consists of k independent vertices, we know that every member of $\mathcal{B}$ contains k vertices of $G - v$ and $\mathcal{B}$ is intersecting, by the lemma above. Using the induction hypothesis on $G - v$, we get that $G - v$ is k-EKR, implying that the largest intersecting families are k-stars. Also Lemma 4.6 guarantees that $\mathcal{I}_x^{(k)}(G - v)$ is a largest k-star of $G - v$ rooted at the leaf x. Therefore we get

$$|\mathcal{B}| \leq |\mathcal{I}_x^{(k)}(G - v)|. \tag{15}$$

Similarly $G \downarrow v$ is an elongated claw with a short limb, fewer vertices than G and either n or $n - 1$ leafs. Each member of $\mathcal{C}$ now contains $k - 1$ vertices of $G \downarrow v$ and by the previous lemma, we know that $\mathcal{C}$ is also intersecting. Using the induction hypothesis again yields that $G \downarrow v$ is $(k-1)$-EKR and the largest intersecting families are $(k-1)$-stars. Applying Lemma 4.6, we know that $\mathcal{I}_x^{(k-1)}(G \downarrow v)$ is a largest $(k-1)$-star of $G \downarrow v$. Thus we get

$$|\mathcal{C}| \leq |\mathcal{I}_x^{(k-1)}(G \downarrow v)|. \tag{16}$$

Now we combine equations (14), (15) and (16) to obtain

$$|\mathcal{A}| = |\mathcal{B}| + |\mathcal{C}| \leq |\mathcal{I}_x^{(k)}(G - v)| + |\mathcal{I}_x^{(k-1)}(G \downarrow v)|.$$

Using Lemma 4.5 we can finish the proof with

$$|\mathcal{A}| \leq |\mathcal{I}_x^{(k)}(G - v)| + |\mathcal{I}_x^{(k-1)}(G \downarrow v)| = |\mathcal{I}_x^{(k)}(G)|.$$

□

4.2 Lexicographic Products with Complete Graphs

This section is based on the work of Fred Holroyd and John Talbot [16]. We will show that the lexicographic product of a k-EKR graph with a complete graph is again k-EKR.

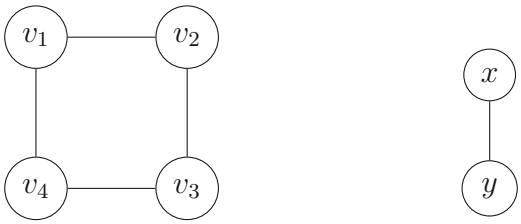

Figure 24: The square graph Q_4 and K_2

Let G and H be graphs, then the *lexicographic product* $G[H]$ is defined as the graph with vertex set $V(G[H]) = V(G) \times V(H)$, where $(v, w) \sim (x, y)$ if and only if either $v \sim x$ in G or $v = x$ and $w \sim y$ in H. This definition is inspired by the lexicographic ordering, where you order pairs by the first entry and only if these are equal, then order by the second entry. The resulting graph can be thought of as a copy of G where

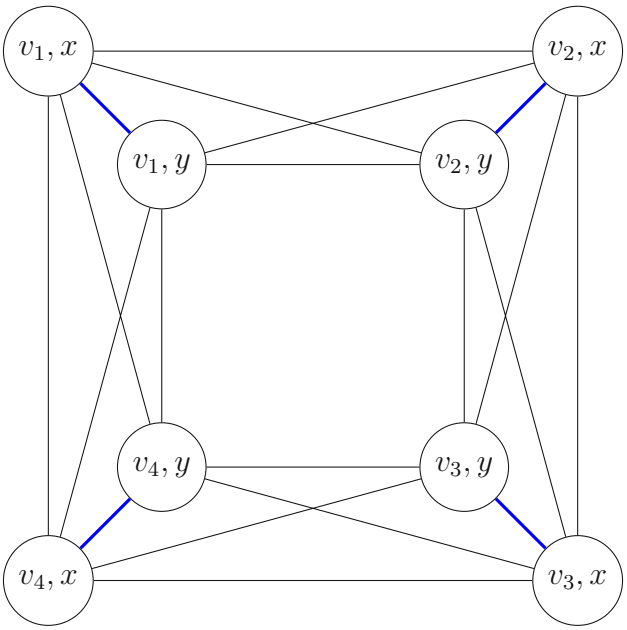

Figure 25: The lexicographic product $Q_4[K_2]$

every vertex $g \in V(G)$ is replaced by a copy of H.

An example is given in Figure 24. We start with the square Graph Q_4 and K_2. Then the lexicographic product $Q_4[K_2]$ is depicted in Figure 25. We color the edges of every copy of the right graph of the product to better distinguish the construction of the edge set.

A natural question to consider is whether the lexicographic product is commutative. In general this is not the case, as illustrated by the product $K_2[Q_4]$ as seen in Figure 26.

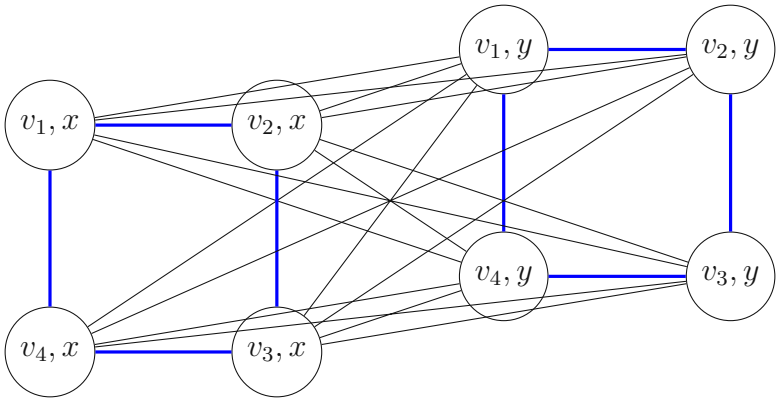

Figure 26: The lexicographic product $K_2[Q_4]$

If we redraw Figure 26 and place one square graph inside the other, then we obtain Figure 27. This way it is easier to see, that the lexicographic product of two graphs

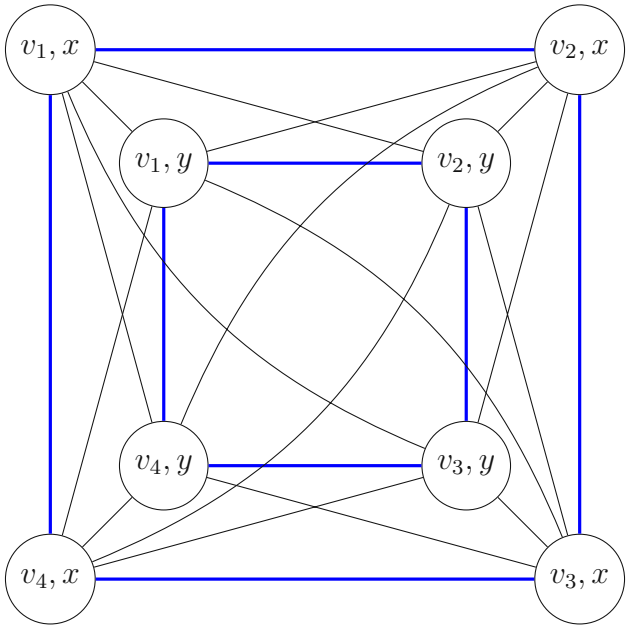

Figure 27: An isomorphic drawing of $K_2[Q_4]$

is not commutative. Thus in general $G[H] \ncong H[G]$. An easier way to verify this is by considering the number of edges of the product. This calculation can be broken down into two cases. Let $G = (V_1, E_1)$, $H = (V_2, E_2)$ and (v, x) be a vertex in the product $G[H]$. If we fix the first coordinate v and only consider vertices such that $(v, x) \sim (v, y)$, then we get a copy of H. In this case we get $|E_2||V_1|$ many edges, since for every vertex in V_1 we get a copy of H. For the second case, $(v, x) \sim (w, y)$ if and only if $v \sim w$ in G, we get that every copy of H is fully connected to every other copy of H, as long as the respective first coordinates are adjacent in G. Thus for every edge in G we get $|V_2|^2$ many new edges. Therefore the total amount of edges in $G[H]$ is

$$|E(G[H])| = |E_1||V_2|^2 + |E_2||V_1|.$$

Since this is not symmetric in general we have $G[H] \ncong H[G]$.

It can be useful to allow an even more general definition of the lexicographic product. Instead of replacing every vertex with the same copy of a graph, we can allow certain vertices to be replaced by different graphs. For example let G, J and H be graphs, then

$$G[v_1, \ldots, v_k : H; w_1, \ldots w_q : J]$$

denotes the lexicographic product of G, where all vertices $v_1, \ldots, v_k$ are replaced by H and all vertices $w_1, \ldots w_q$ are replaced by J. With this definition we can write the disjoint union of the graphs $G_1, \ldots, G_n$ as $E_n[G_1, \ldots, G_n]$. Since the vertices of E_n are interchangeable, we will not mention them specifically.

We start by introducing two lemmata, that help us deal with k-centers in lexicographic products. A family of subsets of a set S is called a *q-covering* of S if each element of S is contained in exactly q sets of the family.

Lemma 4.8. *Let $\mathcal{F}$ be a family of k-subsets of S and let $\Gamma \subseteq \mathcal{P}(\mathcal{F})$ be a family of subfamilies of $\mathcal{F}$. Let $x \in S$. If for some q the following holds*

- *Γ is a q-covering of $\mathcal{F}$;*
- *x is a k-center for each $\mathcal{G} \in \Gamma$,*

then x is a k-center of $\mathcal{F}$.

Proof. Let $\mathcal{A}$ be any intersecting subfamily of $\mathcal{F}$ and suppose that Γ is a q-covering of $\mathcal{F}$. Then Γ is also a q-covering of $\mathcal{A}$, which implies

$$q|\mathcal{A}| = \sum_{\mathcal{G}\in\Gamma} |\mathcal{A} \cap \mathcal{G}|. \tag{17}$$

Since this holds for any intersecting subfamily, it also holds for $\mathcal{F}_x := \{A \in \mathcal{F} : x \in A\}$. Therefore we have

$$q|\mathcal{F}_x| = \sum_{\mathcal{G}\in\Gamma} |\mathcal{F}_x \cap \mathcal{G}| = \sum_{\mathcal{G}\in\Gamma} |\mathcal{G}_x|. \tag{18}$$

The last equality holds, because x is a k-center for each $G \in \Gamma$. For an intersecting subfamily $\mathcal{A}$ of $\mathcal{F}$ and $\mathcal{G} \in \Gamma$ we get that $\mathcal{A} \cap \mathcal{G}$ is an intersecting subfamily of $\mathcal{G}$. Thus we get

$$|\mathcal{A} \cap \mathcal{G}| \leq |\mathcal{G}_x|. \tag{19}$$

Using the equations (17), (18) and (19), we get

$$q|\mathcal{A}| = \sum_{\mathcal{G}\in\Gamma} |\mathcal{A} \cap \mathcal{G}| \leq \sum_{\mathcal{G}\in\Gamma} |\mathcal{G}_x| = q|\mathcal{F}_x|.$$

For any intersecting subfamily $\mathcal{A}$ of $\mathcal{F}$ we get

$$|\mathcal{A}| \leq |\mathcal{F}_x|,$$

proving, that x is a k-center of $\mathcal{F}$. □

Lemma 4.9. *Let v be a k-center of a graph G and let $m \in \mathbb{N}^+$. Then each vertex of the form (v, x), with $x \in V(K_m)$ is a k-center of the lexicographic product $G[K_m]$.*

Proof. We will assume that $m > 1$, since for $m = 1$ the statement is trivial. To make our life easier for this proof, we will identify the vertices of G with the elements $\{1, \ldots, n\}$ and the vertices of K_m with the elements of the cyclic group $\mathbb{Z}_m$. Let $\mathcal{F}$ be the family of functions $f : \{1, \ldots, n\} \to \mathbb{Z}_m$. There are $|\mathcal{F}| = m^n$ many such functions. Then for each $X \in \mathcal{I}^{(k)}(G)$ and $f \in \mathcal{F}$, we define

$$X \circ f := \{(v, f(v)) : v \in X\}.$$

We now define an equivalence relation $\sim$ on our family of functions $\mathcal{F}$, by identifying $f \sim g$, if there exists a $z \in \mathbb{Z}_m$, such that $f(v) = g(v) + z$ for all $v \in \{1, \dots, n\}$. Let Ψ be the family of equivalence classes of this relation. For each $\psi \in \Psi$ we define $\mathcal{I}_\psi$ as

$$\mathcal{I}_\psi := \{X \circ f : X \in \mathcal{I}^{(k)}(G),\ f \in \psi\}.$$

This is a subfamily of $\mathcal{I}^{(k)}(G[K_m])$, since $X \circ f$ is an independent set for $X \in \mathcal{I}^{(k)}(G)$. Let $y \in \mathcal{I}^{(k)}(G[K_m])$, then $|y| = k$. Further since y is independent we know that for all $(v_1, x_1), (v_2, x_2) \in y$ we have $v_1 \neq v_2$. Thus we can represent y as $X \circ f$ for exactly one $X \in \mathcal{I}^{(k)}(G)$ and exactly m^{n-k} functions f. This is because we have to fix k elements and we can map the other vertices freely, also we can choose these functions from distinct equivalence classes. Therefore the family $\{\mathcal{I}_\psi : \psi \in \Psi\}$ is a m^{n-k}-covering of $\mathcal{I}^{(k)}(G[K_m])$. Our goal is to use Lemma 4.8, meaning we have to show that each (v, x) is a k-center of $\mathcal{I}_\psi$, for each $\psi \in \Psi$.

Let $\mathcal{A}$ be an intersecting subfamily of $\mathcal{I}_\psi$, for some $\psi \in \Psi$. We define

$$\mathcal{B} := \{X \in \mathcal{I}^{(k)}(G) : X \circ f \in \mathcal{A} \text{ for some } f \in \psi\}.$$

Then $\mathcal{B}$ is an intersecting subfamily of $\mathcal{I}^{(k)}(G)$ and also $|\mathcal{B}| \leq |\mathcal{I}_v^{(k)}(G)|$. If $X \in \mathcal{I}^{(k)}(G)$ and $f \neq g$ are in ψ, then $X \circ f \cap X \circ g = \emptyset$. But $\mathcal{A}$ is intersecting, thus any two distinct elements of $\mathcal{A}$ correspond to distinct elements of $\mathcal{B}$. Therefore we get $|\mathcal{A}| = |\mathcal{B}|$. Since we have already shown the inequality for $\mathcal{B}$, we get

$$|\mathcal{A}| \leq |\mathcal{I}_v^{(k)}(G)|.$$

Let $x \in \mathbb{Z}_m$ and let (v, x) be a vertex in $G[K_m]$. For each $\psi \in \Psi$ and $X \in \mathcal{I}^{(k)}(G)$, we have $(v, x) \in X \circ f$, for some $f \in \psi$ if and only if $X \in \mathcal{I}_v^{(k)}(G)$ and there is exactly one f that satisfies this condition. This can be seen by considering, that $(v, x) \in X$ has to hold for all X and secondly that there are no two functions f in the same equivalency class that both map $v \mapsto x$. Thus for $(\mathcal{I}_\psi)_{(v,x)}$, which is also a subfamily of $\mathcal{I}^{(k)}(G[K_m])$, we get

$$|(\mathcal{I}_\psi)_{(v,x)}| = |\mathcal{I}_v^{(k)}(G)|.$$

Using our above result we can conclude, that for any intersecting subfamily $\mathcal{A}$ of $\mathcal{I}_\psi$, we have

$$|\mathcal{A}| \leq |(\mathcal{I}_\psi)_{(v,x)}|,$$

which shows that each (v, x) is a k-center of $\mathcal{I}_\psi$. Now we can apply Lemma 4.8, to obtain the result that each vertex (v, x) is a k-center of the lexicographic product $G[K_m]$. □

With these tools we can prove the theorem, all we need to do is to assemble the proof.

Theorem 4.5. *If G is k-EKR and $m \geq 1$, then $G[K_m]$ is also k-EKR.*

Proof. Since G is k-EKR, we know that no independent set is larger, than the largest k-star. Let v be that k-center. By Lemma 4.9, we know that (v, x) is a k-center for $G[K_m]$, which proves the theorem. □

4.3 Further Developments

There have been a number of additional proofs that certain types of graphs are in fact EKR graphs. We will give a brief summary of additional discoveries.

Theorem 4.6 (Borg and Holroyd)**.** *Let G be a disjoint union of complete multipartite graphs, copies of **mnd** graphs, powers of cycles, modified powers of cycles, trees and at least one singleton. Let $k \leq \mu(G)/2$, then G is k-EKR.*

This theorem by Borg and Holroyd [4] proves that a large number of graph classes are EKR. However they require the existence of a singleton vertex. The reason for that is that it simplifies stars, since a singleton is a natural choice for the center of a star.

Another result is from Hurlbert and Kamat [17] that explores chordal graphs.

Theorem 4.7. *Let G be a disjoint union of chordal graphs, including at least one singleton. If $k \leq \mu(G)/2$, then G is k-EKR.*

We end this chapter with a conjecture by Holroyd and Talbot [16], that inspired a lot of researchers to expand the current list of EKR graphs and so far has held true.

Conjecture 4.1 (Holroyd and Talbot)**.** *Let G be a graph and $1 \leq k \leq \mu(G)/2$, then G is k-EKR and strictly k-EKR if $2 < k < \mu(G)/2$.*

One way to gain insights for this bound is to analyze $\alpha(G)$, since we know that $\mu(G) \leq \alpha(G)$. But calculating $\alpha(G)$ for a general graph G is a well-known strongly NP-complete problem [20]. However a computationally easier bound to consider is from a rather recent paper by J. Harant and S. Richter [13]. Let $m = |E|$ be the size of the graph G, d_v be the degree of vertex v and δ be the minimum degree of G. For any graph G the (*normalized*) *Laplacian* is defined as the matrix $\mathcal{L} = (l_{ij}) \in \mathbb{R}^{n \times n}$, with $l_{ij} = 1$ if $i = j$, $l_{ij} = -\frac{1}{\sqrt{d_i d_j}}$ if $ij \in E$ and $l_{ij} = 0$ else. Let σ be the largest eigenvalue of the $\mathcal{L}$ of G. Then

$$\mu(G) \leq \alpha(G) \leq \frac{2\sigma - 2}{\sigma\delta}. \tag{20}$$

Thus if the conjecture turns out to be true, it provides us with an upper bound for a graph to be k-EKR that is significantly easier to compute, although also more lenient.

References

[1] Prathik Anandhan, K. Uma, and J. Anuradha. An overview of application of graph theory. 9:242–248, 01 2016.

[2] Dimitris Bertsimas and John N Tsitsiklis. *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA, 1997.

[3] Béla Bollobás. *Modern graph theory*, volume 184. Springer Science & Business Media, 2013.

[4] Peter Borg and Fred Holroyd. The erdős-ko-rado properties of various graphs containing singletons. *Discrete Mathematics*, 309(9):2877–2885, 2009.

[5] Fan RK Chung. Lectures on spectral graph theory. *CBMS Lectures, Fresno*, 6(92):17–21, 1996.

[6] M Deza and P Frankl. Erdös–ko–rado theorem—22 years later. *SIAM Journal on Algebraic Discrete Methods*, 4(4):419–431, 1983.

[7] David Eppstein. Twenty proofs of euler's formula: V-E+F=2. `https://www.ics.uci.edu/~eppstein/junkyard/euler/`. Accessed: 2021-07-15.

[8] P Erdos. Intersection theorems for systems of finite sets. *Quart. J. Math. Oxford Ser.(2)*, 12:313–320, 1961.

[9] Mark Farrell. Invariant column spaces. `https://m4farrel.quora.com/Invariant-Column-Spaces`. Accessed: 2021-09-30.

[10] Carl Feghali, Matthew Johnson, and Daniel Thomas. Erdős–ko–rado theorems for a family of trees. *Discrete Applied Mathematics*, 236:464–471, 2018.

[11] Peter Frankl and Zoltán Füredi. Non-trivial intersecting families. *Journal of Combinatorial Theory, Series A*, 41(1):150–153, 1986.

[12] Chris Godsil and Gordon F Royle. *Algebraic graph theory*, volume 207. Springer Science & Business Media, 2001.

[13] Jochen Harant and Sebastian Richter. A new eigenvalue bound for independent sets. *Discrete Mathematics*, 338(10):1763–1765, 2015.

[14] Hans Havlicek. *Lineare Algebra für Technische Mathematiker*. Heldermann Verlag, 2012.

[15] Anthony JW Hilton and Eric C Milner. Some intersection theorems for systems of finite sets. *The Quarterly Journal of Mathematics*, 18(1):369–384, 1967.

[16] Fred Holroyd and John Talbot. Graphs with the erdős–ko–rado property. *Discrete mathematics*, 293(1-3):165–176, 2005.

[17] Glenn Hurlbert and Vikram Kamat. Erdös-ko-rado theorems for chordal and bipartite graphs. *arXiv preprint arXiv:0903.4203*, 2009.

[18] G.O.H Katona. A simple proof of the erdös-chao ko-rado theorem. *Journal of Combinatorial Theory, Series B*, 13(2):183–184, 1972.

[19] Gyula Katona. Intersection theorems for systems of finite sets. *Acta Mathematica Academiae Scientiarum Hungaricae*, 15(3-4):329–337, 1964.

[20] CE Larson. The independence number project: α-bounds. *The Conjecturing Project, Virginia Commonwealth University, https://independencenumber. wordpress. com*, 2017.

[21] Daniel Spielman. Spectral graph theory. *Combinatorial scientific computing*, 18, 2012.

[22] Robin J Wilson. *Introduction to graph theory*. Pearson Education India, 1979.