

© 2010 by Kevin G. Milans. All rights reserved.

EXTREMAL PROBLEMS ON EDGE-COLORINGS, INDEPENDENT SETS, AND
CYCLE SPECTRA OF GRAPHS

BY

KEVIN G. MILANS

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Mathematics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2010

Urbana, Illinois

Doctoral Committee:

Professor Alexandr V. Kostochka, Chair
Professor Douglas B. West, Director of Research
Professor Carl G. Jockusch
Research Assistant Professor Sujith Vijay

Abstract

We study problems in extremal graph theory with respect to edge-colorings, independent sets, and cycle spectra. In Chapters 2 and 3, we present results in Ramsey theory, where we seek Ramsey host graphs with small maximum degree. In Chapter 4, we study a Ramsey-type problem on edge-labeled trees, where we seek subtrees that have a small number of path-labels. In Chapter 5, we examine parity edge-colorings, which have connections to additive combinatorics and the minimum dimension of a hypercube in which a tree embeds. In Chapter 6, we prove results on the chromatic number of circle graphs with clique number at most 3. The tournament analogue of an independent set is an acyclic set. In Chapter 7, we present results on the size of maximum acyclic sets in k -majority tournaments. In Chapter 8, we prove a lower bound on the size of the cycle spectra of Hamiltonian graphs.

Acknowledgments

I thank my Mom, my Dad, my sister Kim, and my extended family for constant love and support. I also thank my good friends, whose company I have enjoyed immensely: Bryan K. Clark, Kelly Clark, Zhou Dong, P. Brighten Godfrey, Ellen Juhlin, Ben Kane, the Ann Lewis, Elster McSkelterson, James Pistole, Jacob Richman, Matthew Smith, and Jane E. Valentine.

I wish to extend a special thanks to my advisor, Douglas B. West, who was and is instrumental in shaping my academic growth, and his wife, Ching Muyot, for the many parties they hosted with incredible food, particularly Ching's famous flan. I would also like to thank my dissertation committee members: Carl G. Jockusch, Alexandr V. Kostochka, Sujith Vijay, and Douglas B. West. This work would not have been possible without their guidance.

I am also sincerely grateful for my friendship and mathematical discussions with many others that I met during my time as a graduate student: József Balogh, Michael D. Barus, Csaba Biró, Timothy Brauch, David P. Bunde, Jane V. Butterfield, Erin W. Chambers, Jehanzeb H. Chaudhry, Chandra Chekuri, Jeong-Ok Choi, Joshua Cooper, Daniel W. Cranston, Rosona Eldred, Alina R. Ene, Jeff Erickson, Michael J. Ferrara, Zoltán Füredi, Tracy Grauman, Jerrold R. Griggs, Sarel Har-Peled, Stephen G. Hartke, Michael S. Jacobson, Kyle Jao, Tao Jiang, Adam Jobson, Gwenaël Joret, Jeong-Hyun Kang, Hemanshu Kaul, Ida Š. Kantor, André Kézdy, Hal A. Kierstead, Seog-Jin Kim, Bill Kinnersley, Nitish Korula, Daniel Král, Mohit Kumbhat, Timothy D. LeSaulnier, Ryan R. Martin, Dhruv Mubayi, Suil O, Megan Owen, Michael J. Pelsmajer, Noah Prince, Gregory Puleo, Mike Rosulek, Wojciech Samotij, James Shook, Chris Stocker, Craig Tennenhouse, Shripad Thite, William T. Trotter, Jennifer Vandenbussche, Paul S. Wenger, Lesley W. Wiglesworth, Hehui Wu, and Gexin Yu.

The author acknowledges support of the National Science Foundation through a fellowship funded by the grant “EMSW21-MCTP: Research Experience for Graduate Students” (NSF DMS 08-38434).

Table of Contents

| | | |
|------------------|---|-----------|
| Chapter 1 | Overview | 1 |
| 1.1 | Degree Ramsey theory | 1 |
| 1.2 | Online degree Ramsey theory | 3 |
| 1.3 | Subtrees with few labeled paths | 4 |
| 1.4 | Parity edge-colorings of graphs | 4 |
| 1.5 | Chromatic number of circle graphs | 6 |
| 1.6 | Acyclic sets in k -majority tournaments | 6 |
| 1.7 | Cycle spectra of Hamiltonian graphs | 7 |
| 1.8 | Background material | 8 |
| Chapter 2 | Degree Ramsey Theory | 13 |
| 2.1 | Introduction | 13 |
| 2.2 | Trees | 15 |
| 2.3 | R_{Δ} -bounded Families | 23 |
| 2.4 | Cycles | 25 |
| Chapter 3 | Online Degree Ramsey Theory | 31 |
| 3.1 | Introduction | 31 |
| 3.2 | The greedy Painter, stars, and kites | 33 |
| 3.3 | Graphs with small online degree Ramsey number | 35 |
| 3.4 | Weighted graphs | 36 |
| 3.5 | The consistent Painter | 38 |
| 3.6 | Large odd cycles | 41 |
| Chapter 4 | Subtrees with Few Labeled Paths | 45 |
| 4.1 | Introduction | 45 |
| 4.2 | Some facts and a question about f | 46 |
| 4.3 | An upper bound on $f(n)$ | 50 |
| 4.4 | A lower bound on $f(n)$ | 55 |
| 4.5 | An application to computability theory | 57 |
| Chapter 5 | Parity Edge-Colorings of Graphs | 65 |
| 5.1 | Introduction | 65 |
| 5.2 | Trees, hypercubes, paths, and cycles | 66 |
| 5.3 | Complete graphs | 69 |

| | | |
|-------------------|--|------------|
| 5.4 | Complete bipartite graphs | 73 |
| 5.5 | Open problems | 75 |
| Chapter 6 | The Chromatic Number of Circle Graphs | 77 |
| 6.1 | Introduction | 77 |
| 6.2 | Clean circle graphs | 78 |
| 6.3 | Chromatic number of K_4 -free circle graphs | 86 |
| Chapter 7 | Acyclic Sets in k-majority Tournaments | 90 |
| 7.1 | Introduction | 90 |
| 7.2 | Small odd k | 92 |
| 7.3 | General odd k | 94 |
| Chapter 8 | Cycle Spectra of Hamiltonian Graphs | 98 |
| 8.1 | Introduction | 98 |
| 8.2 | Chords of a spanning path | 100 |
| 8.3 | Cycle lengths in Hamiltonian graphs | 103 |
| References | | 105 |

Chapter 1

Overview

People have a natural fascination with everyday objects that are extreme with respect to some property like height or speed. In much the same way, the mathematical objects that achieve the minimum or maximum value of a parameter are often interesting. Extremal problems have also inspired the development of deep techniques, such as the probabilistic method and the regularity method.

Extremal problems have applications beyond discrete mathematics. For example, the problem of register allocation in computer science can be cast as a graph coloring problem. The existence of certain extremal structures have important consequences for algorithmic performance guarantees, such as the time complexity of matrix multiplication.

Extremal problems come in many different flavors. Chapters 2 through 5 discuss extremal problems involving edge-colorings of graphs. Chapter 6 explores an extremal problem involving proper colorings of circle graphs, where the vertices are partitioned into independent sets. Acyclic sets in tournaments are analogous to independent sets in graphs; in chapter 7, we study the size of maximum acyclic sets in a special class of tournaments. Chapter 8 discusses an extremal problem involving the number of distinct lengths of cycles that occur in Hamiltonian graphs with a specified number of edges.

The following sections give detailed overviews of the remaining chapters. Background material and basic definitions about graphs and digraphs are provided in Section 1.8.

1.1 Degree Ramsey theory

Ramsey theory consists of a diverse array of results that have a common theme: large objects must contain smaller parts that are highly structured. Consider the classic context of Ramsey's Theorem. Let H be a "host" graph, and let G be a "target" graph. If every s -edge-coloring of H contains a monochromatic copy of G as a subgraph, then we write $H \xrightarrow{s} G$. When $s = 2$, we omit the superscript and simply write $H \rightarrow G$. For example, $K_3 \rightarrow P_3$ because a 2-edge-coloring of the triangle K_3 uses the same color on two of its three edges. Here K_n is the complete graph with n vertices, and P_n is the path with n vertices.

Ramsey's Theorem implies that for every graph G and every s , there exists an n such that $K_n \xrightarrow{s} G$. More generally, when the host graph is large enough and dense enough, a

monochromatic copy of the target graph is forced in every s -edge-coloring of the host's edges. The *Ramsey number* of a graph G , denoted $R(G)$, is $\min\{|V(H)| : H \rightarrow G\}$; similarly, we define $R(G; s)$ to be $\min\{|V(H)| : H \xrightarrow{s} G\}$. The Ramsey number is a classic graph parameter and has been studied extensively. The principal question is, given a graph G , how large does the host graph H need to be before $H \xrightarrow{s} G$?

We examine a related question. If we are willing to allow the host graph H to have arbitrarily many vertices, how much degree is needed at vertices in H before in order to have $H \xrightarrow{s} G$? If H has small maximum degree, then H is sparse. The *degree Ramsey number* of a graph G , denoted $R_\Delta(G)$, is $\min\{\Delta(H) : H \rightarrow G\}$, and we define $R_\Delta(G; s)$ to be $\min\{\Delta(H) : H \xrightarrow{s} G\}$.

The degree Ramsey number has been established for some graphs. Burr, Erdős, and Lovász [21] showed that $R_\Delta(K_n) = R(K_n) - 1$ for each n . More recently, Kurek and Ruciński [73] proved the stronger statement that if $H \rightarrow K_n$, then H contains a subgraph with average degree at least $R(K_n) - 1$. Burr, Erdős, and Lovász [21] also characterized the set of host graphs H such that $H \rightarrow K_{1,n}$, where $K_{1,n}$ consists of n edges with one common endpoint. Their characterization immediately implies that $R_\Delta(K_{1,n})$ is $2n - 2$ when n is even and $2n - 1$ when n is odd. In Section 2.2, we generalize this by obtaining the degree Ramsey number of the double-star $S_{a,b}$, which is the tree that consists of two adjacent vertices u and v with $a - 1$ and $b - 1$ leaf neighbors respectively. Note that $K_{1,n} = S_{1,n}$.

In studying the size of monochromatic components of 2-edge-colored graphs, Alon *et al.* [4] gave a short argument that implies $R_\Delta(P_n) \leq 4$ for all n . On the other hand, Thomassen [109] proved that every graph with maximum degree at most 3 can be 2-edge-colored so that all monochromatic components are subgraphs of P_6 . Consequently, $R_\Delta(P_n) = 4$ for $n \geq 7$.

The most intriguing problem about the degree Ramsey parameter is whether $R_\Delta(G; s)$ is bounded by a function of $\Delta(G)$ and s . Jiang observed that the argument of Alon *et al.* extends to show that $R_\Delta(T; s) \leq 2s\Delta(T)$ for each tree T ; a *tree* is a connected graph with no cycles. In Section 2.3, we extend this to prove the following (joint with Jiang and West).

Theorem ([60]). *Let \mathcal{F} be the family of graphs that can be obtained from a tree T by replacing each vertex in T with an independent set and each edge in T with a complete bipartite graph. There exists a function f such that $R_\Delta(G) \leq f(\Delta(G))$ for each graph G in \mathcal{F} .*

In Section 2.4, we specialize these techniques to the n -vertex cycle C_n .

Theorem ([60]). *If n is even, then $R_\Delta(C_n) \leq 96$. For all n , $R_\Delta(C_n) \leq 3458$.*

Previously, a result of Haxell *et al.* [58] implicitly showed that $R_\Delta(C_n)$ is bounded; because their proof uses Szemerédi's regularity lemma, their bound is very large.

1.2 Online degree Ramsey theory

Traditional Ramsey theory can be viewed as a game between two players called Builder and Painter. Builder presents a host graph H to Painter, and then Painter colors the edges of H . Builder wins if the coloring contains a monochromatic copy of some target graph G ; otherwise Painter wins. The statement $R(G) \leq k$ is equivalent to the statement that when G is the target graph, Builder has a winning strategy that presents only graphs with at most k vertices.

An online variant in which Builder and Painter take turns was introduced by Grytczuk, Hałuszczak, and Kierstead [53]. The host graph H starts empty. Builder selects a pair $\{u, v\}$ of non-adjacent vertices (one or both of which may be new) and adds the edge uv to H . Painter responds by coloring uv red or blue. We require Builder to keep the presented graph in a specified family \mathcal{H} . Builder wins if Painter ever completes a monochromatic copy of the target graph G , and Painter wins otherwise. This defines the *online Ramsey game* (G, \mathcal{H}) . The fundamental problem of online Ramsey theory is to characterize the games (G, \mathcal{H}) in which Builder has a winning strategy.

Each parameter Ramsey number corresponding to a monotone parameter has an online analogue. The *online degree Ramsey number*, denoted $\mathring{R}_\Delta(G)$, is $\min\{k: \text{builder wins } (G, \mathcal{H}_k)\}$, where \mathcal{H}_k is the family of graphs with maximum degree at most k . Butterfield, Grauman, Kinnersley, Milans, Stocker, and West [22] studied the online degree Ramsey number of trees and cycles.

Theorem ([22]). *If T is a tree, then $\mathring{R}_\Delta(T) \leq 2\Delta(T) - 1$, with equality whenever T has an adjacent pair of vertices with maximum degree.*

In Section 3.6, we obtain the following result for cycles (joint with Butterfield, Grauman, Kinnersley, Stocker, and West).

Theorem ([22]). *For each n , $\mathring{R}_\Delta(C_n) \in \{4, 5\}$. If n is even, $n = 3$, $337 \leq n \leq 514$, or $n \geq 689$, then $\mathring{R}_\Delta(C_n) = 4$.*

The lower bound $\mathring{R}_\Delta(C_n) > 3$ follows from our characterization of the graphs G with online degree Ramsey number at most 3. Recently, Rolnick [97] proved that $\mathring{R}_\Delta(C_n) = 4$ always. It would be interesting to know if the online degree Ramsey number is bounded by a function of the maximum degree. Of course, if the (offline) degree Ramsey number is bounded by a function of the maximum degree, then so is the online degree Ramsey number. It is known that the online degree Ramsey number of graphs with maximum degree 2 is bounded.

Theorem ([22]). *If $\Delta(G) \leq 2$, then $\mathring{R}_\Delta(G) \leq 6$.*

1.3 Subtrees with few labeled paths

One goal in computability theory is to find algorithmic solutions to combinatorial problems, and solutions to problems in computability theory often require combinatorial proofs. Consequently, computability theory and combinatorics enjoy a history of successful collaboration. We consider a question from computability theory that led to a Ramsey-type problem of finding highly structured binary subtrees in large ternary trees; the resulting problem is of independent combinatorial interest.

A rooted tree is *complete* if all leaves have the same distance from the root, and the *depth* of a complete tree is the common distance between the leaves and the root. A rooted tree is *q-ary* if all non-leaves have q children. If T is a complete tree of depth n in which each edge is labeled with 0 or 1, then reading the edge labels along a path from the root to a leaf yields a *path-label* in $\{0, 1\}^n$. Let $L(T)$ be the set of all path-labels that occur along such paths in T .

When T is a complete ternary tree of depth n and S is a complete binary tree of depth n contained in T , we write $S \sqsubset T$. Given a $\{0, 1\}$ -edge-labeled complete ternary tree T , we seek a binary subtree S with as few path-labels as possible. Let $f(T) = \min \{|L(S)| : S \sqsubset T\}$. Of course, since $L(S) \subseteq \{0, 1\}^n$, always $f(T) \leq 2^n$. We define

$$f(n) = \max \{f(T) : T \text{ is a } \{0, 1\}\text{-edge-labeled complete ternary tree of depth } n\}.$$

The computability theory application requires that $\lim_{n \rightarrow \infty} f(n)/2^n = 0$. In fact, we prove the following.

Theorem ([30]). *There exist positive constants c_1, c_2, c_3 such that $c_1 2^{\frac{1}{\log_2 3} n} \leq f(n) \leq c_2 2^{n - c_3 \sqrt{n}}$ for each n .*

A relatively simple argument shows that $\lim_{n \rightarrow \infty} (f(n))^{1/n}$ exists; our bounds on $f(n)$ yield $1.548 \leq 2^{\frac{1}{\log_2 3}} \leq \lim_{n \rightarrow \infty} (f(n))^{1/n} \leq 2$.

1.4 Parity edge-colorings of graphs

A *parity walk* in an edge-colored graph is a walk that traverses each color an even number of times. An edge-coloring is a *strong parity edge-coloring* if every parity walk is closed (starts and ends at the same vertex). A strong parity edge-coloring is *optimal* if it uses as few colors as possible. The *strong parity edge-chromatic number* of a graph G , denoted $\widehat{p}(G)$, is the number of colors in an optimal strong parity edge-coloring.

Although Bunde, Milans, West, and Wu [19] introduced the strong parity edge-chromatic number as a general graph parameter, related concepts were studied earlier. A 1972 result

of Havel and Movárek [57] essentially implies when T is a tree, $\widehat{p}(T)$ is the minimum k such that T is a subgraph of the k -dimensional hypercube.

Complete graphs admit strong parity edge-colorings with nice structure. Let \mathbb{F}_2^k denote the k -dimensional vector space over the field with 2 elements, and let $A \subseteq \mathbb{F}_2^k$. The *canonical edge-coloring* of the complete graph $K(A)$ with vertex set A assigns each edge uv the color $u + v$. The canonical edge-coloring is a strong parity edge-coloring. Indeed, if u_0, \dots, u_t are the vertices of a walk in which the j th edge has color c_j , then $u_0 + \sum_{j=1}^t c_j = u_0 + \sum_{j=1}^t (u_{j-1} + u_j) = u_t$. Also, if u_0, \dots, u_t is a parity walk, then $\sum_{j=1}^t c_j = 0$. Hence, every parity walk is closed.

The canonical edge-coloring of $K(A)$ does not use the vector $\vec{0}$. Consequently, if n is an integer and m is the smallest power of 2 that is at least n , then $\widehat{p}(K_n) \leq m - 1$. In Section 5.3, we prove that equality holds (joint work with Bunde, West, and Wu).

Theorem ([20]). *For each n , $\widehat{p}(K_n) = 2^{\lceil \log_2 n \rceil} - 1$. In fact, every optimal strong parity edge-coloring of K_n is isomorphic to the canonical edge-coloring of $K(A)$ for some $A \subseteq \mathbb{F}_2^k$, where $k = \lceil \log_2 n \rceil$.*

These results strengthen a special case of Yuzvinsky's Theorem from additive combinatorics. For each $r, s \geq 1$, the *Hopf–Stiefel function* $r \circ s$ is the minimum integer n such that $(x + y)^n$ is in the ideal of $\mathbb{F}_2[x, y]$ generated by x^r and y^s . Equivalently, $r \circ s$ is the minimum integer n such that $\binom{n}{k}$ is even for each k with $n - s < k < r$. Yuzvinsky [113] proved that if $A, B \subseteq \mathbb{F}_2^k$ with $|A| = r$, $|B| = s$, and $C = \{a + b \mid a \in A \text{ and } b \in B\}$, then $|C| \geq r \circ s$; furthermore, this bound on $|C|$ is sharp. When $A = B$, Yuzvinsky's Theorem states that the set of colors $\{a + b \mid a, b \in A \text{ and } a \neq b\}$ used by the canonical edge-coloring of $K(A)$ has size at least $(r \circ r) - 1$. Because $r \circ r = 2^{\lceil \log_2 r \rceil}$ and the family of strong parity edge-colorings is more general than the family of canonical edge-colorings, our result strengthens this case of Yuzvinsky's Theorem.

Canonical edge-colorings extend to complete bipartite graphs in a natural way. If $A, B \subseteq \mathbb{F}_2^k$, then the canonical edge-coloring of the complete bipartite graph $K(A, B)$ with partite sets A and B assigns each edge uv the color $u + v$. The set of colors used by the canonical edge-coloring of $K(A, B)$ is $\{a + b \mid a \in A \text{ and } b \in B\}$. In its full generality, Yuzvinsky's Theorem states that the canonical edge-coloring of $K(A, B)$ uses $r \circ s$ colors, where $r = |A|$ and $s = |B|$. We conjecture that every strong parity edge-coloring of $K(A, B)$ also needs $r \circ s$ colors, or equivalently that $\widehat{p}(K_{r,s}) = r \circ s$, where $K_{r,s}$ denotes the complete bipartite graph with partite sets of sizes r and s . Proving this conjecture would strengthen all cases of Yuzvinsky's Theorem.

1.5 Chromatic number of circle graphs

A *clique* is a set of vertices that are pairwise adjacent, and the *clique number* of a graph G , denoted $\omega(G)$, is the maximum size of a clique in G . The *chromatic number* of a graph G , denoted $\chi(G)$, is the minimum size of a partition of the vertices into independent sets called *color classes*. Because the vertices of a clique must be placed into distinct color classes, it follows that $\chi(G) \geq \omega(G)$ for every graph. While $\chi(G)$ may be arbitrarily large even when $\omega(G) = 2$, there are a number of interesting classes of graphs in which $\chi(G)$ is bounded by a function of $\omega(G)$.

A *circle graph* is a graph G whose vertices are chords of a circle drawn in the plane, with chords u and v adjacent if they cross. Kostochka and Kratochvíl [69] showed that $\chi(G) \leq 50 \cdot 2^{\omega(G)} - 32\omega(G) - 64$ for each circle graph G and Kostochka [70] showed that there are circle graphs with $\omega(G) = k$ and $\chi(G) \geq c \cdot k \log k$ for a constant c . The exponential gap between the two bounds has been open for over a decade. More is known when the clique number is small. If G is a circle graph with $\omega(G) = 2$, then $\chi(G) \leq 5$ [63], and this bound is best possible [1]. In Section 6.3, we investigate the case $\omega(G) \leq 3$ (joint work with Kostochka).

Theorem ([71]). *If G is a circle graph with $\omega(G) \leq 3$, then $\chi(G) \leq 44$.*

Prior to our work, the best known bound for the case $\omega(G) \leq 3$ was $\chi(G) \leq 120$, which follows from an earlier result of Kostochka [69]. Our proofs use a lemma stating that when G belongs to the subfamily of circle graphs not containing a certain structure, $\chi(G) \leq 2\omega(G) - 1$.

1.6 Acyclic sets in k -majority tournaments

Let Π be a set of linear orders of a ground set X . The *majority digraph* of Π is the directed graph D on vertex set X such that $uv \in E(D)$ if more than half of the orders in Π rank u before v . If Π has size k , then D is a *k -majority digraph*. When k is odd, D is an orientation of a complete graph, and so D is a *k -majority tournament*. A *dominating set* in D is a set S of vertices such that for each vertex $v \in V(D)$, either $v \in S$ or $uv \in E(D)$ for some vertex $u \in S$. The *domination number* of D , denoted $\gamma(D)$, is the smallest size of a dominating set in D . Alon *et al.* [3] introduced k -majority tournaments and showed that there is a constant c such that $\gamma(D) \leq ck \log k$ for every k -majority tournament D . Moreover, they construct a family of k -majority tournaments $\{D_k\}$ such that $\gamma(D_k) \geq c'k / \log k$ for some constant k .

Let $a(D)$ denote the maximum size of an acyclic set in D , and let $f_k(n)$ be the minimum of $a(D)$ over all n -vertex k -majority tournaments D . In Section 7.2, we prove the following (joint with Schreiber and West).

Theorem ([78]). *For each n , $f_3(n) \geq \sqrt{n}$. When n is a perfect square, $f_3(n) \leq 2\sqrt{n} - 1$.*

We also show $f_5(n) \geq n^{1/4}$ and prove the following for general k .

Theorem ([78]). *If $c_k = 3^{-(k-1)/2}$ and $d_k = O(\log \log k / \log k)$, then $n^{c_k} \leq f_k(n) \leq n^{d_k}$.*

Because adding a permutation and its reverse to Π does not change the majority digraph, the construction for 3-majority tournaments yields $f_k(n) \leq 2\sqrt{n} - 1$ whenever $k \geq 3$ and n is a perfect square. This bound is the best known when $k = 5$.

1.7 Cycle spectra of Hamiltonian graphs

The *cycle spectrum* of a graph G is the set of lengths of cycles in G . A graph on n vertices is *pancyclic* if its cycle spectrum contains all lengths from 3 to n . Let $s(G)$ denote the size of the cycle spectrum of G . In 1960, Ore [87] showed that every n -vertex graph in which every pair of non-adjacent vertices has degree sum at least n is Hamiltonian. Subsequently, Bondy [17] showed that every graph satisfying Ore's condition is pancyclic unless it is $K_{n/2, n/2}$. On this basis, Bondy proposed the meta-conjecture that natural sufficient conditions for Hamiltonicity are often also sufficient for pancyclicity.

When G is $n/2$ -regular, Bondy's result shows that G is pancyclic or is $K_{n/2, n/2}$. Conditions that are not strong enough to imply pancyclicity may nevertheless imply a large cycle spectrum. Jacobson and Lehel asked how small the cycle spectrum can be when G is a k -regular Hamiltonian graph. For n divisible by $2k$, they constructed k -regular Hamiltonian graphs G with $s(G) = \frac{k-2}{2k}n + k$. Jacobson and Lehel, and independently Jiang, proved (but did not publish) the result that $s(G) \geq \sqrt{a(m-n)}$ when G is an n -vertex Hamiltonian graph with m edges, where a is a positive constant. It follows that $s(G) \geq \sqrt{\frac{a}{2}n}$ when G is a 3-regular Hamiltonian graph. Neither Jacobson and Lehel nor Jiang attempted to optimize the constant a .

Using ideas from a paper of Faudree *et al.* [42], we prove the following in Section 8.3 (joint with Rautenbach, Regen, and West).

Theorem ([77]). *If G is a Hamiltonian graph with n vertices and m edges, then $s(G) \geq \sqrt{\frac{4}{7}(m-n)}$.*

On the other hand, if n is even and $G = K_{n/2, n/2}$, then $s(G) \leq \sqrt{m-n} + 1$. Hence, the largest constant a such that $s(G) \geq \sqrt{a(m-n)} - O(1)$ satisfies $4/7 \leq a \leq 1$. When G is 3-regular, our lower bound yields the following corollary.

Corollary ([77]). *If G is a 3-regular Hamiltonian graph, then $s(G) \geq \sqrt{\frac{2}{7}n}$.*

It is believed that a linear lower bound exists, matching the order of growth in the example of Jacobson and Lehel. However, obtaining better bounds in the case that G is 3-regular or even k -regular with $3 < k < n/2$ will require different ideas from those developed in our paper. This question is a direction for future research.

1.8 Background material

This section provides a brief introduction to common terms and concepts.

Graphs

A *graph* G consists of a set $V(G)$ of *vertices* and a set $E(G)$ of *edges*; each edge is an unordered pair of vertices. Two vertices are *adjacent* if together they form an edge. Unless otherwise stated, all graphs are finite. To improve readability, when u and v are adjacent vertices in G , we write $uv \in E(G)$ or $vu \in E(G)$ instead of $\{u, v\} \in E(G)$. The *neighborhood* of a vertex u in a graph G , denoted $N_G(u)$ (or $N(u)$ when G is clear from context), is the set of all vertices adjacent to u . The *degree* of u in G , denoted $d_G(u)$ (or $d(u)$ when G is clear from context), is $|N(u)|$. The *closed neighborhood* of u , denoted $N[u]$, is $N(u) \cup \{u\}$. The *maximum degree* of a graph G , denoted $\Delta(G)$, is the maximum among all vertex degrees in G ; similarly, the *minimum degree* is denoted by $\delta(G)$. If all vertices in G have degree k , then we say that G is *k -regular* or simply *regular*.

A graph F is a *subgraph* of G if there is an injection $f: V(F) \rightarrow V(G)$ such that $uv \in E(F)$ implies that $f(u)f(v) \in E(G)$. We say that F is a *spanning subgraph* if f is also a bijection. If f has the stronger property that $uv \in E(F)$ if and only if $f(u)f(v) \in E(G)$, then F is an *induced subgraph* of G .

Another name for a spanning subgraph of G is a *factor*. A *k -factor* of G is a k -regular spanning subgraph of G . A *matching* in G is a subgraph with maximum degree at most 1. A *perfect matching* in G is a 1-factor.

An *isomorphism* from a graph G to a graph H is a bijection $f: V(G) \rightarrow V(H)$ such that $uv \in E(G)$ if and only if $f(u)f(v) \in E(H)$. We say that G and H are *isomorphic* if there is an isomorphism from G to H . Note that isomorphism defines an equivalence relation on the class of all graphs. We often use the same notation for a graph and its isomorphism class. When G denotes an isomorphism class, we say that G is a subgraph of a graph H if some member of G is a subgraph of H .

A *path* with n vertices is a graph whose vertices can be labeled v_1, \dots, v_n so that the edge set is $\{v_i v_{i+1} : 1 \leq i < n\}$. We specify a path with vertices v_1, \dots, v_n in order as $\langle v_1, \dots, v_n \rangle$. The *endpoints* of $\langle v_1, \dots, v_n \rangle$ are v_1 and v_n ; the other vertices are *internal* vertices. We use P_n to denote the isomorphism class of n -vertex paths. A *cycle* with n vertices is a graph

whose vertices can be labeled v_1, \dots, v_n so that the edge set is $\{v_i v_{i+1} : 1 \leq i < n\} \cup \{v_n v_1\}$. We specify a cycle with vertices v_1, \dots, v_n in order as $[v_1, \dots, v_n]$. We use C_n to denote the isomorphism class of n -vertex cycles. An n -vertex graph is *Hamiltonian* if it contains a spanning subgraph that is a cycle. The *length* of a path or cycle is the number of edges it contains. A graph is *acyclic* if it does not contain a cycle as a subgraph.

A *walk* W in a graph G is a list u_1, \dots, u_k of vertices (not necessarily distinct) such that $u_j u_{j+1}$ is an edge in G for $1 \leq j < k$. We say that the walk u_1, \dots, u_k *traverses* the edges of the form $u_j u_{j+1}$. A u, v -*walk* is a walk that starts at u and ends at v . The *length* of a walk is one less than the length of its list. A walk is *open* if it starts and ends at different vertices and is *closed* otherwise. A *circuit* is a closed walk. An *Eulerian circuit* is a circuit that traverses each edge in G exactly once.

A graph G is *connected* if for all $u, v \in V(G)$, there is a u, v -walk in G . The *components* of G are the maximal connected subgraphs of G . If u and v are vertices in the same component of G , the *distance* between u and v , denoted $\text{dist}(u, v)$, is the length of a shortest u, v -walk in G . Note that a shortest u, v -walk does not repeat vertices and is therefore a path. The *diameter* of a connected graph G , denoted $\text{diam}(G)$, is $\max\{\text{dist}(u, v) : u, v \in V(G)\}$. When G contains a cycle as a subgraph, the *girth* of G is the length of a shortest cycle in G .

A vertex u is a *leaf* in G if u has degree 1. When u is a leaf, its neighborhood consists of one vertex v , and uv is called a *pendant edge*. A *tree* is an acyclic connected graph.

The *complete graph* K_n is the graph on n vertices in which each of the $\binom{n}{2}$ unordered pairs forms an edge. A graph G is *bipartite* if the vertices can be partitioned into two sets X and Y (possibly empty) such that each edge has an endpoint in X and an endpoint in Y . When we want to name the vertex partition, we may introduce G as an X, Y -*bigraph*. The *complete bipartite graph* $K_{m,n}$, or *biclique*, is the X, Y -bigraph with $|X| = m$ and $|Y| = n$ such that x and y are adjacent whenever $x \in X$ and $y \in Y$.

An *independent set* in a graph is a set of pairwise nonadjacent vertices. A graph G is r -*partite* if the vertices can be partitioned into r (possibly empty) independent sets. Note that a graph is bipartite if and only if it is 2-partite. The *complete r -partite graph* K_{n_1, \dots, n_r} consists of r disjoint independent sets X_1, \dots, X_r with $|X_j| = n_j$ such that x and y are adjacent whenever x and y are not contained in the same set.

The *chromatic number* of G , denoted $\chi(G)$, is the least r such that it is possible to color the vertices of G with r colors so that adjacent vertices receive different colors. Equivalently, $\chi(G)$ is the least r such that G is an r -partite graph. An s -*edge-coloring* is a function that assigns to each edge a color from a set of size s .

A *clique* in a graph is a set of pairwise adjacent vertices. Cliques and independent sets are complementary objects. The *clique number* of G , denoted $\omega(G)$, is the size of a largest clique in G , and the *independence number*, denoted $\alpha(G)$, is the size of a largest independent set in G .

Graph Operations

It is often useful to construct new graphs by modifying specified graphs in particular ways. The more common modifications have special notation. The complement of G , denoted \overline{G} , is the graph with the same vertex set as G such that uv is an edge in \overline{G} if and only if u and v are non-adjacent in G .

If G is a graph and $R \subseteq V(G)$, then the graph obtained from G by *deleting* R , denoted $G - R$, is the subgraph of G with vertex set $V(G) - R$ and edge set $\{xy \in E(G) : \{x, y\} \cap R = \emptyset\}$. When R is a singleton $\{v\}$, we use $G - v$ for $G - \{v\}$. If $S \subseteq V(G)$, then the graph obtained by *deleting* the vertices outside of S from G is the *subgraph of G induced by S* , denoted $G[S]$.

When e is an edge, the graph obtained from G by *deleting* e , denoted $G - e$, has the same vertex set and has edge set $E(G) - \{e\}$. Note that a graph F is a subgraph of G if it is obtainable from G by deleting vertices and/or edges, F is an induced subgraph of G if it is obtainable by deleting only vertices, and F is a spanning subgraph if it is obtainable from G by deleting only edges.

When G_1 and G_2 are graphs, the *union* $G_1 \cup G_2$ is the graph with vertex set $V(G_1) \cup V(G_2)$ and edge set $E(G_1) \cup E(G_2)$. The *disjoint union* $G_1 + G_2$ is the union of disjoint copies of G_1 and G_2 . A *decomposition* of a graph G is a family of edge-disjoint subgraphs whose union is G .

An *automorphism* of a graph G is an isomorphism from G to itself. An automorphism is an expression of a symmetry in the graph; if the vertices of G are unlabeled and f is an automorphism, then u and $f(u)$ are indistinguishable in the sense that every statement about u applies equally well to $f(u)$. We say that G is *edge-transitive* if, for every pair of edges e_1 and e_2 , there is an automorphism of G that maps the vertex set of e_1 to the vertex set of e_2 .

If G is edge-transitive and $e \in E(G)$, then the isomorphism class of $G - e$ does not depend on e . In this case, we write G^- for $G - e$. Similarly, when \overline{G} is edge-transitive, adding an edge uv when u and v are non-adjacent yields a graph whose isomorphism class does not depend on the choice of the pair. This graph is denoted by G^+ .

Directed Graphs

A *directed graph* (or *digraph*) D consists of a set $V(D)$ of vertices and a set $E(D)$ of edges. Each edge is an ordered pair of vertices. When $(u, v) \in E(D)$, we say that uv is an edge *from* u *to* v , and we write $uv \in E(D)$ instead of $(u, v) \in E(D)$, for readability. An *orientation* of a graph G is a directed graph obtained from G by replacing each edge uv in G with either uv or vu in D . A *tournament* is an orientation of a complete graph.

In the context of directed graphs, a *path* with n vertices is a digraph whose vertices can be

labeled v_1, \dots, v_n so that the edge set is $\{v_i v_{i+1} : 1 \leq i < n\}$. As with undirected graphs, we specify a path with vertices v_1, \dots, v_n in order as $\langle v_1, \dots, v_n \rangle$. The *endpoints* of $\langle v_1, \dots, v_n \rangle$ are v_1 and v_n ; the other vertices are *internal* vertices. A *cycle* with n vertices is a digraph whose vertices can be labeled v_1, \dots, v_n so that the edge set is $\{v_i v_{i+1} : 1 \leq i < n\} \cup \{v_n v_1\}$. We specify a cycle with vertices v_1, \dots, v_n in order as $[v_1, \dots, v_n]$. We use C_n to denote the isomorphism class of n -vertex cycles. A digraph is *acyclic* if it does not contain a cycle as a subgraph.

Other Conventions and Notation

We use $\lg(x)$ for $\log_2(x)$. It is often useful to compare the order of growth of real-valued functions; a common convention for expressing these relationships is called “Big-Oh notation”. If f and g are functions, we write $f(x) = O(g(x))$ if there are positive constants c and x_0 such that $|f(x)| \leq c|g(x)|$ whenever $x \geq x_0$. This is a common abuse of notation, since technically $O(g(x))$ represents a set of functions. Similarly, we write $f(x) = \Omega(g(x))$ if there are positive constants c and x_0 such that $|f(x)| \geq c|g(x)|$ whenever $x \geq x_0$. When $f(x) = O(g(x))$ and $f(x) = \Omega(g(x))$, we write $f(x) = \Theta(g(x))$. We write $f(x) = o(g(x))$ if $\lim_{x \rightarrow \infty} |f(x)|/|g(x)| = 0$.

Computational complexity theory is the study of how much of a resource (such as time or memory) is needed to solve a computational problem. The complexity class \mathbf{P} is the set of all decision problems that can be solved in time that is bounded by a polynomial function of the size of the input. The complexity class \mathbf{NP} is the set of all decision problems that can be solved with a non-deterministic algorithm in time that is bounded by a polynomial function of the size of the input. A *nondeterministic algorithm* is a (deterministic) algorithm $\mathcal{A}(x, s)$ that takes as input a problem instance x and a seed s and returns either “YES” or “NO”. The *seed* is a bitstring; typically, a nondeterministic algorithm interprets the seed as an encoding of some object, such as a graph or the prime factorization of an integer. A nondeterministic algorithm solves a problem A if

- x is a “YES” instance of A implies that $\mathcal{A}(x, s)$ outputs “YES” for some seed s , and
- x is a “NO” instance of A implies that $\mathcal{A}(x, s)$ outputs “NO” for every seed s .

A nondeterministic algorithm solves a problem in polynomial time if $\mathcal{A}(x, s)$ runs in time that is bounded by a polynomial in the sizes of x and s , and, whenever x is a “YES” instance, there is a seed s whose size is bounded by a polynomial function of the size of x such that $\mathcal{A}(x, s)$ outputs “YES”. Since a deterministic algorithm $\mathcal{A}(x)$ lifts to a nondeterministic algorithm $\mathcal{A}'(x, s)$ that makes no use of its seed, it is clear that $\mathbf{P} \subseteq \mathbf{NP}$. Most researchers believe the inclusion is proper. It remains a fundamental problem to determine whether $\mathbf{P} = \mathbf{NP}$.

A problem A is *NP-hard* if a polynomial-time algorithm for A can be used to obtain polynomial-time algorithms for each problem in \mathbf{NP} . Since most researchers believe $\mathbf{P} \neq$

NP, a proof that a problem is **NP**-hard is considered to be strong evidence that it does not admit a polynomial-time algorithm for its solution. If A is also in **NP**, then A is **NP**-complete. Cook [27] proved that a particular problem in **NP**, called SAT, is **NP**-complete. Hence, to resolve the question of $\mathbf{P} = \mathbf{NP}$, it suffices to determine if SAT admits a polynomial-time algorithm. Since Cook’s groundbreaking result, many interesting and natural problems have been shown to be **NP**-complete. In the context of graph theory, for example, the problem of deciding if a given graph has chromatic number at most 3 is **NP**-complete. The most commonly used technique for establishing that a problem A is **NP**-hard is to *reduce* a problem B that is already known to be **NP**-hard to A , meaning that there is a polynomial-time algorithm for B that uses a hypothetical polynomial-time algorithm for A as a subroutine. When A is **NP**-complete, showing that A is in **NP** is typically (although not always) easier than showing that A is **NP**-hard.

A *partially ordered set* (or *poset*) is a set X on which a binary relation \leq is defined that is *reflexive* ($x \leq x$ for all $x \in X$), *antisymmetric* ($x \leq y$ and $y \leq x$ imply $x = y$), and *transitive* ($x \leq y$ and $y \leq z$ imply $x \leq z$). Elements x and y are *comparable* if $x \leq y$ or $y \leq x$, and they are *incomparable* otherwise. A *chain* is a set of pairwise comparable elements, and an *antichain* is a set of pairwise incomparable elements. The *height* of a poset is the size of a largest chain, and the *width* is the size of a largest antichain. We sometimes use the terminology *partial order* in preference to “partially ordered set”. A *pre-partial order* is a set X on which a binary relation \leq is defined that is reflexive and transitive.

Let X be a ground set. A *partition* of X is a family of nonempty disjoint subsets whose union is X . When \mathcal{F}_1 and \mathcal{F}_2 are partitions of X , we say that \mathcal{F}_1 is a *refinement* of \mathcal{F}_2 if each $A \in \mathcal{F}_1$ is contained in some $B \in \mathcal{F}_2$. We also say that \mathcal{F}_1 *refines* \mathcal{F}_2 when \mathcal{F}_1 is a refinement of \mathcal{F}_2 .

Chapter 2

Degree Ramsey Theory

Ramsey theory is the study of the principle that large objects must contain smaller highly structured objects. This principle holds in a surprisingly diverse set of contexts, and accordingly Ramsey theory is an active area of research in several different areas of mathematics.

In classical graph Ramsey theory, we are given a target graph G and seek a graph H such that every 2-edge-coloring of H produces a monochromatic copy of G , in which case we write $H \rightarrow G$ and say that H *arrows* or *forces* G and that H is a *Ramsey host* for G . (A 2-edge-coloring of H assigns one of two colors, traditionally red and blue, to each edge in H .) More generally, when every s -edge-coloring of $E(H)$ produces a monochromatic copy of G , we write $H \xrightarrow{s} G$; when some s -edge-coloring of $E(H)$ avoids monochromatic copies of G , then we write $H \not\xrightarrow{s} G$.

Ramsey's Theorem guarantees for every G that such a graph H exists. For a fixed target G , we seek Ramsey hosts that are extremal with respect to a particular property. Classical Ramsey theory asks for a host with as few vertices as possible. In "size Ramsey theory", host graphs are sought with as few edges as possible. We focus on "degree Ramsey theory", where we seek host graphs with small maximum degree.

This chapter is based on joint work with B. Kinnersley and D. B. West that appears in [67] and joint work with T. Jiang and D. B. West that appears in [60].

2.1 Introduction

The *Ramsey number* $R(G)$ (or $R(G; s)$) is the smallest number of vertices in a graph H such that $H \rightarrow G$ (or $H \xrightarrow{s} G$). Exact Ramsey numbers are known only for a handful of graphs. Obtaining bounds on the Ramsey number of the complete graph K_n is a classic problem. It is known that $\Omega(n2^{n/2}) < R(K_n) < (4 + o(1))^n$, and improving the base of the exponential function in either bound would constitute a major advance. The growth of $R(G)$ is quite different when the maximum degree of vertices in G , denoted $\Delta(G)$, is bounded. Chvátal, Rödl, Szemerédi, and Trotter [24] proved that for each k , there exists a constant c_k such that for each n -vertex target graph G with $\Delta(G) \leq k$, the Ramsey number of G satisfies $R(G) \leq c_k n$.

The *size Ramsey number* of a graph G , denoted $\hat{R}(G)$, is $\min \{|E(H)| : H \rightarrow G\}$. Clearly,

$\hat{R}(G) \leq \binom{R(G)}{2}$ for each graph G . In a paper by Erdős, Faudree, Rousseau, and Schelp [37], a proof credited to Chvatál shows that equality holds when G is a complete graph. Beck [9] showed that $\hat{R}(P_n) \leq cn$, where c is a constant and P_n is the path on n vertices. Beck then asked whether the size Ramsey number of graphs of bounded maximum degree grows linearly, analogously to the Ramsey number. This is the case for trees [48] and for cycles [58]. However, Rödl and Szemerédi [95] answered Beck’s question in the negative by exhibiting an infinite family \mathcal{F} of graphs with maximum degree 3 such that if G is an n -vertex graph in \mathcal{F} , then $\hat{R}(G) \geq c_1 n (\log n)^{c_2}$, where c_1 and c_2 are positive constants. (See [10, 29, 96] for further results on size Ramsey number.)

More generally, given a graph parameter ρ , the ρ -Ramsey number, denoted $R_\rho(G; s)$, is $\min\{\rho(H) : H \xrightarrow{s} G\}$. Folkman [46] proved that in the 2-color setting, the clique Ramsey number of a graph G equals its clique number $\omega(G)$, and Nešetřil and Rödl [85] extended this to s colors. Burr, Erdős, and Lovász [21] showed that the chromatic Ramsey number $R_\chi(G)$ equals the Ramsey number of the family of homomorphic images of G , where the Ramsey number of a family \mathcal{G} is the minimum number of vertices in a graph H such that every 2-coloring of $E(H)$ produces a monochromatic copy of some graph in \mathcal{G} . Since every homomorphic image of K_n contains K_n , it follows that $R_\chi(K_n) = R(K_n)$. They also conjectured that $\min\{R_\chi(G; s) : \chi(G) = k\}$ equals the easy lower bound $k^s + 1$. When $s = 2$, Zhu [112] proved this for $k \leq 5$, but otherwise it remains open.

In this chapter, we study the *degree Ramsey number* R_Δ , where $\Delta(G)$ denotes the maximum degree of G . Note that $R_\Delta(G)$ is $\min\{\Delta(H) : H \rightarrow G\}$. Establishing the degree Ramsey number of a graph appears to be a more difficult problem than establishing its Ramsey number. Indeed, for each graph G and each integer k , it is a finite problem to check whether $R(G) > k$. By contrast, deciding if $R_\Delta(G) > k$ may not even be a computable problem.

Nevertheless, the degree Ramsey number is known for some graphs. The result of Burr, Erdős, and Lovász [21] about R_χ implies that $R_\Delta(K_n) = R(K_n) - 1$. More recently, Kurek and Ruciński [73] proved the stronger statement that if $H \rightarrow K_n$, then H contains a subgraph with average degree at least $R(K_n) - 1$. Burr, Erdős, and Lovász [21] also characterized the set of graphs H such that $H \rightarrow K_{1,n}$. Their characterization immediately implies that $R_\Delta(K_{1,n})$ is $2n - 2$ when n is even and $2n - 1$ when n is odd. In Section 2.2, we generalize this by obtaining the degree Ramsey number of the *double-star* $S_{a,b}$, the tree with adjacent vertices u and v of degrees a and b and no other non-leaf vertices. The graph $K_{1,n}$ arises as $S_{1,n}$.

In studying the size of monochromatic components of 2-edge-colored graphs, Alon, Ding, Oporowski, and Vertigan [4] gave a short argument that proves $R_\Delta(P_n; 2s) \leq 2s$ for all n . In Section 2.2, we present a simple extension of the upper bound to show that $R_\Delta(T; s) \leq 2s(\Delta(T) - 1)$ for each tree T . On the other hand, Thomassen [109] proved that every

graph with maximum degree at most 3 can be 2-edge-colored so that all monochromatic components are subgraphs of P_6 . Consequently, $R_\Delta(P_n) = 4$ for all $n \geq 7$. In establishing the linearity of the size Ramsey numbers of cycles [58], Haxell *et al.* also showed that $R_\Delta(C_n)$ is at most a constant c , where C_n denotes the cycle on n vertices. Because their proof uses Szemerédi's Regularity Lemma, c is quite large. In Section 2.4, we show that $R_\Delta(C_n) \leq 96$ when n is even and $R_\Delta(C_n) \leq 3458$ in general. For small cycles, exact results are known. It follows from [21] that $R_\Delta(C_3) = 5$, and $R_\Delta(C_4) = 5$ [67].

When G is a tree or a cycle, the degree Ramsey number of G is bounded by a function of $\Delta(G)$. It is natural to ask about other families in which the degree Ramsey number is bounded by a function of the maximum degree. In Section 2.3, we explore this question.

2.2 Trees

Burr, Erdős, and Lovász [21] began the study of R_Δ by characterizing the set of all Ramsey hosts for the star $K_{1,m}$, and they applied their characterization to obtain the degree Ramsey numbers of stars. We extend these results in two ways: first to multiple colors, and second to double-stars.

Multicolored Stars

Burr, Erdős, and Lovász [21] proved that $H \rightarrow K_{1,m}$ if and only if $\Delta(H) \geq 2m - 1$ or H is $(2m - 2)$ -regular with an odd number of edges; it follows that $R_\Delta(K_{1,m})$ is $2m - 2$ when m is even and $2m - 1$ otherwise. When m is odd, the characterization of Ramsey hosts for stars extends naturally to the s -color setting.

Proposition 2.2.1. *If m is odd, then $H \xrightarrow{s} K_{1,m}$ if and only if $\Delta(H) \geq s(m - 1) + 1$.*

Proof. By the pigeonhole principle, $\Delta(H) \geq s(m - 1) + 1$ implies that $H \xrightarrow{s} K_{1,m}$. For the converse, suppose that m is odd and $\Delta(H) \leq s(m - 1)$. Construct a graph H' that contains H as a subgraph and is $s(m - 1)$ -regular. We claim that $H' \not\xrightarrow{s} K_{1,m}$. By Petersen's Theorem [88], H' decomposes into 2-factors. Group the 2-factors into s sets of size $(m - 1)/2$, and use one color for the edges of the 2-factors in a single group. Because each color class is $(m - 1)$ -regular, no monochromatic copy of $K_{1,m}$ occurs. \square

While the characterization of multicolor Ramsey hosts for $K_{1,m}$ appears more difficult when m is even, computing the multicolor degree Ramsey number of stars is easy.

Proposition 2.2.2. *If $s \geq 2$, then $R_\Delta(K_{1,m}; s) = \begin{cases} s(m - 1) & \text{if } m \text{ is even} \\ s(m - 1) + 1 & \text{if } m \text{ is odd.} \end{cases}$*

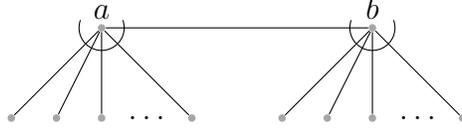


Figure 2.1: The double-star $S_{a,b}$.

Proof. When m is odd, the result follows from Proposition 2.2.1. Suppose that m is even. First, we show that $R_{\Delta}(K_{1,m}; s) > s(m-1) - 1$. Let H be a graph with $\Delta(H) \leq s(m-1) - 1$; we show that $H \xrightarrow{s} K_{1,m}$. By Vizing's Theorem [110], H decomposes into a family of matchings M_1, \dots, M_t where $t \leq \Delta(H) + 1 = s(m-1)$. Group the matchings into s sets of size at most $m-1$, and color edges of H according to the parts. Because the subgraph induced by each color class has maximum degree at most $m-1$, no monochromatic copy of $K_{1,m}$ occurs.

Bollobás, Saito, and Wormwald proved that if t is odd and $r > t$, then there are r -regular graphs that do not have a t -factor [15]. Let H be an $s(m-1)$ -regular graph that does not contain an $(m-1)$ -factor. An s -edge-coloring of H that avoids a monochromatic copy of $K_{1,m}$ requires that every vertex have degree $m-1$ in each color. Hence, each color class would yield an $(m-1)$ -factor of H , which is impossible. It follows that $H \xrightarrow{s} K_{1,m}$. \square

Double-Stars

The *double-star* $S_{a,b}$ is the tree whose non-leaves are adjacent vertices of degrees a and b ; see Figure 2.1. Because $K_{1,m} = S_{1,m}$, the family of double-stars generalizes the family of stars. Our second extension of the star result of Burr, Erdős, and Lovász [21] is the determination of the degree Ramsey number of double-stars. Our upper bound requires the concept of a graph blowup.

Definition 2.2.3. A graph G' is a *blowup* of G if G' is obtained from G by replacing the vertices of G with independent sets and replacing the edges of G with complete bipartite graphs. The independent sets in G' that replace vertices of G are called *clusters*. If each vertex in G is replaced by an independent set of size k , then we say that G' is the *k -blowup* of G .

Theorem 2.2.4. If $a \leq b$, then $R_{\Delta}(S_{a,b}) = \begin{cases} 2b - 2 & \text{if } a < b \text{ and } b \text{ is even} \\ 2b - 1 & \text{otherwise.} \end{cases}$

Proof. Because $K_{1,b}$ is a subgraph of $S_{a,b}$, it follows that $R_{\Delta}(S_{a,b}) \geq R_{\Delta}(K_{1,b})$. By Proposition 2.2.2 with $s = 2$, we have that $R_{\Delta}(K_{1,b})$ is $2b - 2$ when b is even and $2b - 1$ when b is odd. This suffices for the lower bound except when $a = b$ and b is even. Let H be a

connected graph with $\Delta(H) \leq 2b - 2$, and let H' be a $(2b - 2)$ -regular connected graph that contains H . Because H' is a connected graph in which every vertex has even degree, H' contains an Eulerian circuit C . Starting with an arbitrary vertex v , color the edges of H' by alternating red and blue along C . If u is a vertex in H' and $u \neq v$, then different colors are assigned when entering and leaving u , so u is incident to $b - 1$ red edges and $b - 1$ blue edges. Similarly, except at the beginning and end of C , different colors are assigned when entering and leaving v . If H' has an even number of edges, different colors are assigned to the first and last edges in C , and v also has $b - 1$ incident red edges and $b - 1$ incident blue edges. If H' has an odd number of edges, then the first and last edges are both colored red, and v is incident to b red edges and $b - 2$ blue edges. Because every vertex except v has degree at most $b - 1$ in each color, it follows that $H' \not\rightarrow S_{b,b}$.

For the upper bound, we first show that $R_\Delta(S_{a,b}) \leq 2b - 1$. Let H be the complete bipartite graph $K_{2b-1,2b-1}$ with bipartition (X, Y) . We claim that $H \rightarrow S_{a,b}$; in fact, we show that $H \rightarrow S_{b,b}$. Suppose for a contradiction that there is a $\{\text{red, blue}\}$ -edge-coloring of H that does not contain a monochromatic copy of $S_{b,b}$. We label the vertices of H red or blue according to the majority color of incident edges. We claim that X does not contain vertices of different colors. Indeed, if $u \in X$ has at least b incident red edges and $v \in X$ has at least b incident blue edges, then because $|Y| = 2b - 1$, there is a vertex $w \in Y$ such that uw is red and vw is blue. If w is red, then uw is the center edge in a red copy of $S_{b,b}$, and if w is blue, then vw is the center edge of a blue copy of $S_{b,b}$. It follows that all vertices in X have the same color, say red. Therefore more than half the edges in H are red, and Y contains at least one red vertex w . Now each red edge incident to w is the center edge of a red copy of $S_{b,b}$, a contradiction.

When $a < b$ and b is even, we improve the upper bound. Let H be the $(b - 1)$ -blowup of an odd cycle of length at least 5. Let A_1, \dots, A_t be the clusters of H , indexed in order around the cycle. Note that H is $2(b - 1)$ -regular. We claim that $H \rightarrow S_{a,b}$. Suppose for a contradiction that some $\{\text{red, blue}\}$ -edge-coloring of H avoids a monochromatic copy of $S_{a,b}$. Each vertex v in G is labeled red (if v is incident to at least b red edges), blue (if v is incident to at least b blue edges), or tied (if v has incident to $b - 1$ edges of each color). Note that not all vertices are tied, because then the red subgraph of H would be $(b - 1)$ -regular with an odd number of vertices, which is impossible since $b - 1$ is odd.

Because H does not contain triangles, if uv is red when u is red and v is red or tied, then uv is the center edge of a red copy of $S_{a,b}$. Therefore if uv is red and u is red, it follows that v is blue. Because each cluster has size $b - 1$ and a red vertex has b incident red edges, it follows that if cluster A_j contains a red vertex, then both neighboring clusters A_{j-1} and A_{j+1} contain at least one blue vertex. After following the clusters around the cycle, it follows that some cluster A_j contains a red vertex u and a blue vertex v . Because u has at least b incident red edges, v has at least b incident blue edges, and the common neighborhood of all

vertices in A_i has size $2b - 2$, there exists a vertex w such that uw is red and vw is blue. If w is red or tied, then uw is the center of a red copy of $S_{a,b}$. If w is blue or tied, then vw is the center of a blue copy of $S_{a,b}$. The contradiction implies that $H \rightarrow S_{a,b}$. \square

The multicolor degree Ramsey number of double-stars is studied in [67].

Paths

Paths have maximum degree 2. As the length of a path grows, how does the required degree in a Ramsey host grow? As it turns out, the required degree in the host graph is bounded by a constant in terms of the number of colors.

Alon *et al.* proved that $R_\Delta(P_n; s) = 2s$ when n is sufficiently large in terms of s [4]. The upper bound $R_\Delta(P_n; s) \leq 2s$ holds for all n ; we present a generalization to trees in Theorem 2.2.10. In fact, almost every graph with maximum degree at most $2s$ is an s -color Ramsey host for a long path.

Theorem 2.2.5. *Let s and m be fixed integers, let \mathcal{G}_n be the set of all labeled, n -vertex graphs with maximum degree at most $2s$, and let \mathcal{F}_n be the set of all graphs $H \in \mathcal{G}_n$ such that $H \xrightarrow{s} P_m$. As $n \rightarrow \infty$, the ratio $|\mathcal{F}_n|/|\mathcal{G}_n|$ tends to zero.*

Proof. We obtain a lower bound on $|\mathcal{G}_n|$ and an upper bound on $|\mathcal{F}_n|$ in terms of auxiliary families.

Let A be the set of nonempty subsets of $[2s]$, and let \mathcal{H}_n be the family of all labeled n -vertex A -edge-colored graphs with maximum degree at most $2s$. First, we give a lower bound on $|\mathcal{H}_n|$ when n is even. Fix a ground set X of size n . For each $2s$ -tuple (M_1, \dots, M_{2s}) of matchings on X , we obtain a graph H in \mathcal{H}_n by adding an edge uv to H whenever uv is an edge in some matching and coloring uv with the set $\{j: uv \in M_j\}$. It follows that $|\mathcal{H}_n| \geq z^{2s}$, where z is the number of perfect matchings on X . It is well known that $z = \prod_{j=1}^{n/2} (2j - 1) = \frac{n!}{2^{n/2}(n/2)!}$, so $z \sim \sqrt{2} \left(\frac{n}{2e}\right)^{n/2}$ by Stirling's formula. Hence $|\mathcal{H}_n| \geq z^{2s} \geq e^{sn \ln(n/2e)}$ for sufficiently large n . Moreover, $|\mathcal{H}_n| \leq |\mathcal{G}_n| |A|^{ns}$ because each graph in \mathcal{H}_n is uniquely obtained by selecting a graph in \mathcal{G}_n and choosing a color from A for each edge, and there are at most ns edges. It follows that

$$|\mathcal{G}_n| \geq \frac{|\mathcal{H}_n|}{|A|^{ns}} \geq \frac{e^{sn \ln(n/2e)}}{4^{s^2 n}} \geq e^{sn \ln(n/2e) - (s^2 \ln 4)n} \geq e^{sn \ln n - \alpha n},$$

where α is a constant.

Next, we give an upper bound on $|\mathcal{F}_n|$. When the maximum degree is bounded, the diameter of a graph must grow with the number of vertices. Let c be the maximum number of vertices in a graph with maximum degree at most $2s$ and diameter less than $m - 1$. For each graph $H \in \mathcal{F}_n$, there is an s -edge-coloring of H witnessing $H \xrightarrow{s} P_m$. It follows that

every monochromatic component contains at most c vertices, or else some monochromatic component has diameter at least $m - 1$ and therefore contains P_m .

We use the small monochromatic components to assign distinct codes to each graph in \mathcal{F}_n . The code of a graph $H \in \mathcal{F}_n$ is an s -tuple (H_1, \dots, H_s) , where each H_j is a spanning subgraph of H whose components have at most c vertices. Moreover, the H_j are chosen so that their union is H , implying that distinct graphs in \mathcal{F}_n are assigned distinct codes. Let \mathcal{Q} be the set of all graphs on X with maximum degree at most $2s$ whose components have at most c vertices. We have $|\mathcal{F}_n| \leq |\mathcal{Q}|^s$.

We assign distinct codes to graphs in \mathcal{Q} . Let $H' \in \mathcal{Q}$. The first part of the code for H' is a composition of n that records the sizes of the components of H' . A *composition* of n is a list (n_1, \dots, n_k) of positive integers where $n = \sum n_j$. Fix an indexing $\{v_1, \dots, v_n\}$ of X . Given H' , we obtain its composition by indexing the components C_1, \dots, C_k in order of the least indexed vertex they contain, and set $n_j = |V(C_j)|$. It is well known that there are 2^{n-1} compositions of n .

The second part of the code for H' records the distribution of vertices to components using a single list of $n - k$ vertices. The list begins with all vertices in C_1 except for the least indexed vertex in C_1 , followed by all vertices in C_2 except for the least index vertex in C_2 , and continues until all components are recorded. The first and second parts of the code for H' are sufficient to recover the partition of X into vertex sets of components of H' . There are at most n^{n-k} possible choices for the second part of the code. Moreover, each component has size at most c , so $k \geq n/c$. It follows that there are at most $e^{(1-1/c)n \ln n}$ choices for the second part.

The third part of the code for H' records the edge relation. For each vertex v in H' , we record a $2s$ -tuple (u_1, \dots, u_{2s}) where each u_j is either one of at most $c - 1$ other vertices in the same component as v (indicating that vu_j is an edge in H') or a special indicator value designed to fill coordinates when v has degree less than $2s$. There are at most c^{2s} possible $2s$ -tuples for each vertex, and therefore at most c^{2sn} possible codes for the third part.

Multiplying all choices for the three parts together, we have

$$|\mathcal{Q}| \leq 2^{n-1} \cdot e^{(1-1/c)n \ln n} \cdot c^{2sn} \leq e^{(n-1) \ln 2 + (1-1/c)n \ln n + 2sn \ln c},$$

and it follows that $|\mathcal{F}_n| \leq e^{(1-1/c)sn \ln n + \beta n}$, where β is a constant. Combining our bounds on $|\mathcal{G}_n|$ and $|\mathcal{F}_n|$, we have

$$\frac{|\mathcal{F}_n|}{|\mathcal{G}_n|} \leq \frac{e^{(1-1/c)sn \ln n + \beta n}}{e^{sn \ln n - \alpha n}} \leq e^{(\alpha + \beta)n - (s/c)n \ln n}.$$

The bound tends to zero as n grows. □

For the lower bound $R_\Delta(P_n; s) \geq 2s$, Alon *et al.* [4] proved that for each s , there exists

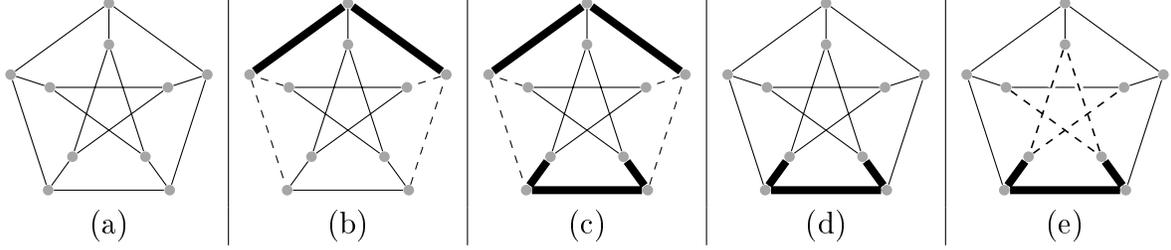


Figure 2.2: Petersen graph as a Ramsey host for P_3 , P_4 , and P_5 . Red edges are thick, and blue edges are dashed.

a constant c such that whenever $\Delta(H) \leq 2s - 1$, there is an s -edge-coloring of H in which every monochromatic component has at most c edges. When $s = 2$, more precise results are known. Thomassen [109] showed that if $\Delta(H) \leq 3$, then there is a 2-edge-coloring of H in which every monochromatic component is a subgraph of P_6 . Consequently, $R_\Delta(P_n) = 4$ when $n \geq 7$. Combining the results of [4] and [109] with some observations about the Petersen graph, we find $R_\Delta(P_n)$ exactly for all paths except P_6 .

Theorem 2.2.6.
$$R_\Delta(P_n) = \begin{cases} n - 1 & n \leq 4 \\ 3 & n \in \{4, 5\} \\ 4 & n \geq 7. \end{cases}$$

Proof. The results for $n \leq 3$ are clear. For $n \geq 7$, the result follows from the theorems of Alon *et al.* [4] and Thomassen [109]. For $n \in \{4, 5\}$, the lower bound $R_\Delta(P_n) \geq 3$ holds because if $\Delta(H) \leq 2$, then H is the disjoint union of paths and cycles, and therefore H can be 2-edge-colored so that monochromatic components are subgraphs of P_3 .

It remains to show that $R_\Delta(P_n) \leq 3$ for $n \in \{4, 5\}$. In fact, we prove that the Petersen graph (see Figure 2.2a) is a Ramsey host for P_4 and P_5 . Let H be the Petersen graph. Note that H is P_k -transitive for each $k \in \{1, 2, 3, 4\}$, meaning that if G_1 and G_2 are copies of P_k in H , then there is an automorphism that restricts to an isomorphism from G_1 to G_2 . This is well known and is easy to prove using the algebraic description of H as the disjointness graph on $\binom{[5]}{2}$.

First, we show that $H \rightarrow P_4$. Consider a {red, blue}-edge-coloring of H , and suppose for a contradiction there is no monochromatic copy of P_4 . Because H contains cycles of odd length, it has a monochromatic copy of P_3 , which we may assume is red. Because H is P_3 -transitive, we may assume that the red copy of P_3 is at the top of the outer 5-cycle. The red P_3 extends to a red P_4 unless the edges incident to its endpoints are blue (see Figure 2.2b). Now the blue copies of P_3 extend to a blue P_4 unless the edges incident to its endpoints on the outer cycle are red, but this forces a red P_4 (see Figure 2.2c). It follows that every {red, blue}-edge-coloring of H contains a monochromatic copy of P_4 .

Next, we show that $H \rightarrow P_5$. Consider a {red, blue}-edge-coloring of H , and suppose for a contradiction there is no monochromatic copy of P_5 . We have seen that there is a monochromatic copy of P_4 . Because H is P_4 -transitive, we may assume it occurs in red as shown in Figure 2.2d. The red P_4 extends to a red P_5 unless the edges incident to its endpoints are blue, but this forces a blue P_5 (see Figure 2.2e). It follows that every {red, blue}-edge-coloring of H contains a monochromatic copy of P_5 . \square

It remains open whether $R_\Delta(P_6)$ is 3 or 4. A graph H is *subcubic* if $\Delta(H) \leq 3$. Let H be a subcubic graph. Thomassen's proof that H admits a 2-edge-coloring in which every monochromatic component is a subgraph of P_6 is long. Here, we give a short proof of the weaker result that H has a 2-edge-coloring in which every monochromatic component has at most 32 vertices.

Lemma 2.2.7. *Let F be the graph obtained from P_7 by adding a pendant leaf to the center vertex. If G is a subcubic bipartite graph, then there is a 2-edge-coloring of G under which all monochromatic components are subgraphs of F .*

Proof. Because every bipartite graph with maximum degree k is a subgraph of a k -regular bipartite graph, we may assume that G is 3-regular. Because G is a 3-regular bipartite graph, Hall's Theorem implies that G contains a perfect matching M . Let X and Y be the partite sets of G , with $X = \{x_1, x_2, \dots, x_n\}$ and $Y = \{y_1, y_2, \dots, y_n\}$, indexed so that $M = \{x_j y_j : 1 \leq j \leq n\}$.

Let $H = G - M$. Because H is 2-regular, it consists of disjoint cycles. We 2-color the edges of H as follows. If C is a cycle with length divisible by four, then we color the edges of C so that the monochromatic components are copies of P_3 with endpoints at vertices in Y . Otherwise, if C has length $4k + 2$, then we color the edges of C so that the monochromatic components are copies of P_3 with endpoints at vertices in Y , except that one monochromatic component is a copy of P_5 with endpoints at vertices in Y .

We extend the edge-coloring of H to an edge-coloring of G as follows. Consider a vertex $x_j \in X$. By construction, both edges in H incident to x_j receive the same color; we assign the opposite color to $x_j y_j$. Monochromatic components in H may only grow to larger monochromatic components in G via edges in M , which can only add pendant leaves to vertices in Y . Hence, monochromatic copies of P_3 can grow to copies of P_4 or P_5 , and monochromatic copies of P_5 can grow to as large as F if it grows at each vertex in Y in the copy of P_5 . Hence each monochromatic component is a subgraph of F . \square

Theorem 2.2.8. *If G is a subcubic graph, then there is a 2-edge-coloring of G in which each monochromatic component has at most 32 vertices.*

Proof. We may assume that G is connected. It follows from Brooks' Theorem that $\chi(G) \leq 4$, with equality only if $G = K_4$. If $G = K_4$, then there is a 2-edge-coloring of G in which each

monochromatic component is a copy of P_4 . Hence, we may assume that $\chi(G) \leq 3$. By Lemma 2.2.7, we may assume that $\chi(G) = 3$.

Let V_1, V_2, V_3 be a partition of $V(G)$ into three independent sets such that if $u \in V_j$ and $i < j$, then u has a neighbor in V_i . Hence, each vertex $u \in V_3$ has a neighbor $u_1 \in V_1$ and a neighbor $u_2 \in V_2$. Let H be the graph on $V_1 \cup V_2$ obtained from G by deleting V_3 and adding the edges of the form u_1u_2 for $u \in V_3$ that are not already present in G . Because H is a subcubic bipartite graph, Lemma 2.2.7 implies that there is a 2-edge-coloring f of H such that each monochromatic component is a subgraph of the graph F specified in the lemma.

We color the edges of G as follows. Edges in $[V_1, V_2]$ are given the same color they have under f . It remains to color edges incident to vertices in V_3 . If $u \in V_3$, then we color uu_1 and uu_2 in G with the same color as u_1u_2 in H , and we use the opposite color on a third edge incident to u , if present in G .

In progressing from H to G , monochromatic components may grow via edges incident to vertices in V_3 . Note that for each vertex $u \in V_3$, our coloring only assigns the same color, say red, to a pair of edges $\{uu_1, uu_2\}$ incident to u if u_1u_2 is a red edge in H . Therefore monochromatic components in H can grow to become a monochromatic component in G either by attaching pendant leaves to vertices in H or by adding a common neighbor to vertices that were adjacent in H . In particular, no two components of H having the same color can combine in G . Because each vertex in H has at most three neighbors in V_3 , the monochromatic component in G corresponding to a monochromatic component of H with k vertices has at most $4k$ vertices. Because each monochromatic component in H has order at most 8, it follows that each monochromatic component in G has at most 32 vertices. \square

General Trees

T. Jiang observed that the argument of Alon *et al.* generalizes to give upper bounds on the degree Ramsey numbers of trees. We include this short proof due to Jiang because we will extend these ideas to give upper bounds on the degree Ramsey numbers of cycles in Section 2.4. We need a short, well-known lemma.

Lemma 2.2.9. *If r is an integer with $r \geq 2$, and H has average degree at least $2(r - 1)$, then H contains a subgraph with minimum degree at least r .*

Theorem 2.2.10 ([60]). *If T is a tree and H is a graph with average degree at least $2s(\Delta(T) - 1)$ and girth larger than $2|V(T)|$, then $H \xrightarrow{s} T$.*

Proof. Let $r = \Delta(T)$, and consider an s -edge-coloring of H . By the pigeonhole principle, some color class has average degree at least $2(r - 1)$, and by Lemma 2.2.9, H contains a monochromatic subgraph H' with minimum degree at least r . Let u be a vertex in H' . Because H' has minimum degree at least r and girth larger than $2|V(T)|$, the subgraph of

H' induced by $\{v: \text{dist}_{H'}(v, u) \leq |V(T)|\}$ is a complete r -ary tree of depth $|V(T)|$. Hence H' contains a monochromatic copy of T . \square

Corollary 2.2.11 ([60]). *If T is a tree, then $R_\Delta(T; s) \leq 2s(\Delta(T) - 1)$.*

Proof. For each k and each g , there is a k -regular graph with girth at least g [39]. \square

The result of Alon *et al.* [4] that $R_\Delta(P_n; s) = 2s$ shows that Corollary 2.2.11 is sharp when T is a path. In [67], it is shown that for each $\varepsilon > 0$, there is an s such that for all sufficiently large k , there exists a tree T with $\Delta(T) = k$ such that $R_\Delta(T; s) \geq (2 - \varepsilon)s(k - 1)$.

By Corollary 2.2.11, within the family of trees, the degree Ramsey number is bounded by a function of the number of colors and the maximum degree. In Section 2.3, we consider generalizations to larger families.

2.3 R_Δ -bounded Families

Perhaps the most interesting question about the degree Ramsey parameter is the following.

Question 2.3.1. Is $R_\Delta(G; s)$ bounded by some function of $\Delta(G)$ and s ?

Question 2.3.1 seems difficult. We have seen that when G is restricted to the family of trees, the answer is yes. By way of analogy with Gyárfás' concept of χ -bounded graph families, we introduce R_Δ -bounded families. A family \mathcal{G} of graphs is R_Δ -bounded if there is a function f such that $R_\Delta(G; s) \leq f(\Delta(G), s)$ for each $G \in \mathcal{G}$, and \mathcal{G} is *weakly R_Δ -bounded* if there is a function f such that $R_\Delta(G) \leq f(\Delta(G))$ for each $G \in \mathcal{G}$. By Corollary 2.2.11, the family of trees is R_Δ -bounded. Question 2.3.1 is equivalent to asking whether the family of all graphs is R_Δ -bounded.

Which other graph families are R_Δ -bounded? In other words, which graph families have sparse Ramsey hosts? Positive progress toward Question 2.3.1 is made by proving that a larger graph family is R_Δ -bounded. Alternately, Question 2.3.1 can be refuted by exhibiting a single family that fails to be R_Δ -bounded.

First, we adapt the proof of Burr *et al.* on the chromatic Ramsey number [21] to show that the family of all blowups of graphs in an R_Δ -bounded family is R_Δ -bounded. (See Definition 2.2.3 for the definition of d -blowup.) We need a bipartite version of Ramsey's Theorem: for each d and s , there is an m such that $K_{m,m} \xrightarrow{s} K_{d,d}$. Let $B_s(d)$ be the minimum m such that $K_{m,m} \xrightarrow{s} K_{d,d}$. Currently, the best known bound for $s = 2$ is that $B(d) \leq (1 + o(1))2^{d+1} \log_2 d$ [26].

Let B_s^k be the iterated composition of B_s with itself k times. For example, $B_s^2(d) = B_s(B_s(d))$ and $B_s^0(d) = d$.

Theorem 2.3.2. *If $k = R_\Delta(G; s)$ and G' is the d -blowup of G , then $R_\Delta(G'; s) \leq kB_s^{k+1}(d)$.*

Proof. Let H be a graph with $\Delta(H) = k$ such that $H \xrightarrow{s} G$, let $m_j = B_s^j(d)$ for $j \geq 0$, and let H' be the m_{k+1} -blowup of H . We show that $H' \xrightarrow{s} G'$. Fix an s -edge-coloring of H' . By Vizing's Theorem [110], the edges of H can be partitioned into $k+1$ matchings M_1, \dots, M_{k+1} . Each edge uv in H corresponds to a copy of $K_{m_{k+1}, m_{k+1}}$ in H' . Because $m_{k+1} = B_s(m_k)$, each s -edge-coloring of $K_{m_{k+1}, m_{k+1}}$ contains a monochromatic copy of K_{m_k, m_k} . For each edge uv in M_{k+1} , find a monochromatic copy of K_{m_k, m_k} and delete all vertices in the u -cluster and v -cluster that are not involved. For each vertex w in H not saturated by M_{k+1} , arbitrarily select a set of m_k vertices and delete the others. After processing M_{k+1} , all clusters contain m_k vertices. Iterate this process for the remaining matchings to obtain a blowup H'_0 of H in which each cluster contains d vertices and the bicliques joining clusters are monochromatic.

Construct an s -edge coloring of H by assigning to uv the common color of the biclique joining the u -cluster and v -cluster in H'_0 . Because $H \xrightarrow{s} G$, a monochromatic copy of G occurs, and this copy lifts to a monochromatic copy of G' in H'_0 . It follows that $H' \xrightarrow{s} G'$. \square

Corollary 2.3.3. *Let \mathcal{F} be a family of graphs, and let \mathcal{G} be the family of blowups of graphs in \mathcal{F} . If \mathcal{F} is R_Δ -bounded, then so is \mathcal{G} . If \mathcal{F} is weakly R_Δ -bounded, then so is \mathcal{G} .*

Proof. Let $G \in \mathcal{G}$ be a blowup of a graph $F \in \mathcal{F}$. Because adding isolated vertices does not change the degree Ramsey number, we may assume that G and F contain no isolated vertices. Note that if d is the maximum size of a cluster in G , then G is a subgraph of the d -blowup G' of F . Because F has no isolated vertices, $\Delta(G') \geq d$. By Theorem 2.3.2, $R_\Delta(G; s) \leq R_\Delta(G'; s) \leq kB_s^{k+1}(d) \leq kB_s^{k+1}(\Delta(G'))$, where $k = R_\Delta(F; s)$. \square

With Corollary 2.2.11 and Corollary 2.3.3, we obtain the following result.

Corollary 2.3.4. *The family of blowups of trees is R_Δ -bounded.*

While Corollary 2.3.4 provides a general family of graphs that is R_Δ -bounded, every graph in the family is bipartite. Establishing R_Δ -boundedness for families of non-bipartite graphs seems challenging. In Section 2.4, we show that the family of all cycles is weakly R_Δ -bounded.

The iterated bipartite Ramsey number $B_s^k(d)$ grows very rapidly with k and d , and the bounding function obtained for the family of blowups of trees in Corollary 2.3.4 depends on the iterated bipartite Ramsey numbers. Adapting the strategy of Corollary 2.2.11 yields better bounds in this case. We need the following well-known lemma.

Lemma 2.3.5. *For each g and k , there is a k -regular bipartite graph with girth at least g .*

Theorem 2.3.6. *Let T be a tree, and let G be the d -blowup of T . If $m = B_s(d)$, then $R_\Delta(G; s) \leq 2sm \binom{m}{d}^2 (\Delta(T) - 1)$.*

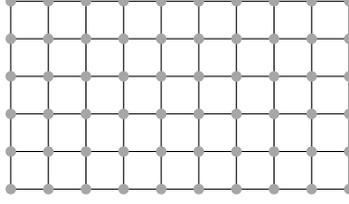


Figure 2.3: The grid $P_6 \square P_{10}$.

Proof. Let H_0 be an (X, Y) -bigraph with girth larger than $2|V(T)|$ that is regular of degree $2s \binom{m}{d}^2 (\Delta(T) - 1)$. Let H be the m -blowup of H_0 . For each cluster U in H corresponding to a vertex u in H_0 , we fix an arbitrary indexing $\{u_1, \dots, u_m\}$ of U .

We claim that $H \xrightarrow{s} G$. Consider an s -edge-coloring of H . We construct an $s \binom{m}{d}^2$ -edge-coloring of H_0 . Let uv be an edge in H_0 . Note that uv in H_0 corresponds to a complete (U, V) -bigraph in H , where U and V are the clusters of size m in H corresponding to u and v in H_0 . Because $m = B_s(d)$, there are subsets $U_0 \subseteq U$ and $V_0 \subseteq V$ of size d that induce a monochromatic copy of $K_{d,d}$. We color uv in H_0 with a 3-tuple. The first coordinate records the color in which $K_{d,d}$ appears; there are s possibilities. The second and third coordinates record the subsets U_0 and V_0 . Assuming without loss of generality that $u \in X$ and $v \in Y$, we record the set $\{j : u_j \in U_0\}$ in the second coordinate and the set $\{j : v_j \in V_0\}$ in the third. Since H_0 is regular of degree $2s \binom{m}{d}^2 (\Delta(T) - 1)$ and has girth larger than $2|V(T)|$, Theorem 2.2.10 implies that every $s \binom{m}{d}^2$ -edge-coloring of H_0 contains a monochromatic copy of T . Using the monochromatic bicliques encoded in the $s \binom{m}{d}^2$ -edge-coloring of H_0 , we obtain a monochromatic copy of G in H . \square

A *grid* is the Cartesian product $P_m \square P_n$ of two paths (see Figure 2.3). Each grid in the family $\{P_n \square P_n : n \in \mathbb{N}\}$ has maximum degree 4, but these graphs do not all appear as subgraphs of d -blowups of trees for any fixed d . Consequently, Theorem 2.3.6 does not give a constant upper bound on the degree Ramsey number of grids.

Question 2.3.7. Is the family of grids R_Δ -bounded or even weakly R_Δ -bounded?

2.4 Cycles

The proof of a result on the size Ramsey number due to Haxell, Kohayakawa, and Łuczaket *al.* [58] describes Ramsey hosts with few edges that force cycles. Because these Ramsey hosts have bounded degree, it follows from their work that there exists a constant c such that $R_\Delta(C_n) \leq c$ for every cycle C_n . Their proof uses Szemerédi's Regularity Lemma [107], and thus the value c that they obtain is very large.

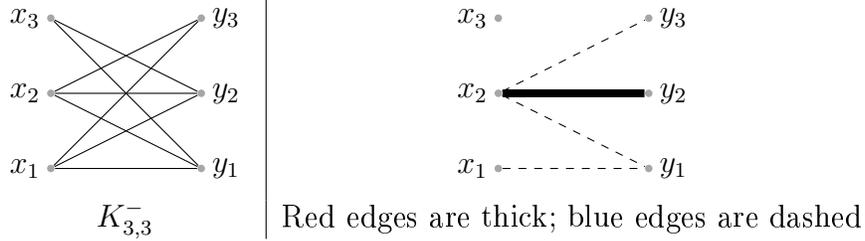


Figure 2.4: $K_{3,3}^-$

The even cycle C_{2n} is a subgraph of the 2-blowup of P_{n+1} . Because $B_2(2) = 5$ [12], it follows from Theorem 2.3.6 that $R_\Delta(C_{2n}) \leq 20\binom{5}{2}^2 = 2000$. In this section, we improve the upper bound for even cycles to $R_\Delta(C_{2n}) \leq 96$. We also show that $R_\Delta(C_n) \leq 3458$ for general n .

Even Cycles

The main improvement to our upper bound on $R_\Delta(C_{2n})$ results from the observation that, in the argument of Theorem 2.3.6, we do not need the full strength of obtaining monochromatic copies of $K_{2,2}$ joining the clusters. It suffices to obtain monochromatic copies of P_4 , and paths are easier to force than $K_{2,2}$. Recall that if G is an edge-transitive graph, then G^- denotes the graph obtained from G by deleting an edge.

Lemma 2.4.1. $K_{3,3}^- \rightarrow P_4$.

Proof. Let H be a {red, blue}-edge-colored copy of $K_{3,3}^-$ with bipartition (X, Y) and missing edge joining $x_3 \in X$ and $y_3 \in Y$ (see Figure 2.4 left).

We say that a vertex is red if it is incident to at least two red edges, and blue if it is incident to at least two blue edges. Note that if two vertices u and v in X are red, then they have a common neighbor $y \in Y$ that is red. Hence, uyv is a red copy of P_3 , which extends to a red copy of P_4 using the other red edge incident to v . Therefore we obtain a monochromatic copy of P_4 unless $\{x_1, x_2\}$ contains one red vertex and one blue vertex, and, by symmetry, $\{y_1, y_2\}$ contains one red vertex and one blue vertex. Assume without loss of generality that x_1 and y_1 are red, and x_2 and y_2 are blue. If x_1y_1 is red, then this edge extends to a red copy of P_4 using the other red edges at x_1 and y_1 , so x_1y_1 must be blue. By symmetry, x_2y_2 is red. Since x_2 is blue, the other two edges incident to x_2 are blue (see Figure 2.4 right). Now $x_1y_1x_2y_3$ is a blue copy of P_4 . \square

Lemma 2.4.1 is sharp in the sense that if H is a proper subgraph of $K_{3,3}^-$, then $H \not\rightarrow P_4$. In passing from a graph to its 3-blowup, each edge expands to a copy of $K_{3,3}$; the maximum degree triples. Our arguments that bound the degree Ramsey number of cycles require only that edges expand to copies of $K_{3,3}^-$. If F is an (X, Y) -bigraph with $|X| = |Y| = d$, we say

that H is an F -blowup of a graph G if H is obtained from G by replacing each vertex in G with an independent set of size d and replacing each edge in G with a copy of F .

In general, a graph can have many different $K_{3,3}^-$ -blowups, depending on the selection of which edges are deleted from copies of $K_{3,3}$ to obtain copies of $K_{3,3}^-$. Our next lemma shows how to distribute the deleted edges evenly to save on the maximum degree.

Lemma 2.4.2. *If G is a k -regular graph and 18 divides k , then there is a $\frac{8}{3}k$ -regular $K_{3,3}^-$ -blowup of G .*

Proof. Because k is even, Petersen's Theorem [88] implies that G decomposes into 2-factors. Group these into nine sets of size $k/18$ to obtain a decomposition of G into nine $k/9$ -factors there is a copy of P_4 . Index them as F_{ij} for $1 \leq i \leq 3$ and $1 \leq j \leq 3$.

Let G' be the 3-blowup of G , and for each vertex u in G , fix an arbitrary indexing of the cluster $\{u^1, u^2, u^3\}$ in G' corresponding to u . Obtain a graph H from G' by removing one edge from each copy of $K_{3,3}$ as follows. For each edge uv in G that belongs to F_{ij} , we remove $u^i v^j$ from G' . Note that G' is 3-regular and that each vertex in G' is incident to exactly $3 \cdot k/9$ removed edges. It follows that H is a $(3 - \frac{1}{3})k$ -regular $K_{3,3}^-$ -blowup of G . \square

To obtain even cycles, it suffices to obtain a P_4 -blowup of a path.

Lemma 2.4.3. *If H is a P_4 -blowup of P_{t+1} , then H contains C_{2t} .*

Proof. Denote the path by $\langle z_1, \dots, z_{t+1} \rangle$ with corresponding clusters Z_1, \dots, Z_{t+1} in H . Note that joining Z_j and Z_{j+1} , there is a copy of P_4 and hence a matching of size 2. Linking these matchings from Z_2 through Z_t yields two vertex-disjoint copies of P_{t-1} . Because there is a copy of P_4 joining Z_1 and Z_2 , some vertex in Z_1 is adjacent to both vertices in Z_2 . Similarly, some vertex in Z_{t+1} is adjacent to both vertices in Z_t . These vertices extend the copies of P_{t-1} to a cycle with $2t$ vertices. \square

Our next lemma shows how to obtain monochromatic P_4 -blowups of long paths.

Lemma 2.4.4. *Let H_0 be a 36-regular (X, Y) -bigraph with girth more than $2t$, and let H be a $K_{3,3}^-$ -blowup of H_0 . If H is $\{\text{red, blue}\}$ -edge-colored, then H contains a monochromatic P_4 -blowup of a path in H_0 on t vertices.*

Proof. We use the $\{\text{red, blue}\}$ -edge-coloring of H to construct an 18-edge-coloring of H_0 . For each vertex u in H_0 , fix an arbitrary indexing $\{u^1, u^2, u^3\}$ of the cluster in H corresponding to u . Each edge uv in H_0 corresponds to a copy of $K_{3,3}^-$ with bipartition (U, V) in H . By Lemma 2.4.1, a monochromatic copy of P_4 occurs. We color uv with a 3-tuple. The first coordinate records the color in which the copy of P_4 occurs (red or blue). The second and third coordinates record the subsets U_0 of U and V_0 of V contained in the monochromatic P_4 . Without loss of generality, assume that $u \in X$ and $v \in Y$. We use the second coordinate to

record the set U_0 and we use the third coordinate to record the set V_0 . There are two choices for the first coordinate and three choices for the second and third coordinates, yielding 18 colors in total.

Since H_0 has girth larger than $2t$, Theorem 2.2.10 implies that $H_0 \xrightarrow{18} P_t$, and a monochromatic copy of P_t in H_0 lifts to a monochromatic P_4 -blowup of P_t in H . \square

Theorem 2.4.5. $R_\Delta(C_{2n}) \leq 96$.

Proof. Let H_0 be a 36-regular bipartite graph with girth more than $2(n+1)$. By Lemma 2.4.2, there is a 96-regular $K_{3,3}^-$ -blowup H of H_0 . We claim that $H \rightarrow C_{2n}$. Consider a {red, blue}-edge-coloring of H . By Lemma 2.4.4, H contains a monochromatic P_4 -blowup of P_{n+1} . By Lemma 2.4.3, it follows that H contains a monochromatic copy of C_{2n} . \square

Odd Cycles

Here, we prove a general bound of $R_\Delta(C_n) \leq 3458$. We need a simple proposition, again about $K_{3,3}^-$ and P_4 .

Proposition 2.4.6. *If H is obtained from $K_{3,3}^-$ by deleting one vertex from each part of the bipartition, then P_4 is a subgraph of H .*

Proof. Deleting one vertex from each part yields $K_{2,2}$ or $K_{2,2}^-$. Each contains P_4 as a subgraph. \square

To complete our proof for odd cycles, we need two lemmas that adapt Lemma 2.4.3 to produce odd cycles.

Lemma 2.4.7. *Let H be a P_4 -blowup of a path $\langle z_1, \dots, z_t \rangle$, where each vertex z_j expands to z_j^1 and z_j^2 in H . If $z_1^1 z_1^2$ is an edge in H , then H contains cycles of every odd length from 3 to $2t - 1$.*

Proof. Let $Z_j = \{z_j^1, z_j^2\}$. Note that there is a copy of P_4 joining Z_j and Z_{j+1} , and hence they are joined by a matching of size two. Choose $k \in [t - 1]$. Linking these matchings from Z_1 through Z_k yields two vertex-disjoint copies of P_k . Because there is a copy of P_4 joining Z_k and Z_{k+1} , some vertex in Z_{k+1} is adjacent to both vertices in Z_k , linking the two copies of P_k to form a copy of P_{2k+1} with both endpoints in Z_1 . The edge $z_1^1 z_1^2$ completes a copy of C_{2k+1} . \square

Lemma 2.4.8. *Let H be a P_4 -blowup of a path $\langle z_1, \dots, z_t \rangle$, where each vertex z_j expands to a set Z_j of two vertices in H . If H contains an edge joining Z_1 and Z_3 , then H contains cycles of every odd length from 7 to $2t - 3$.*

Proof. Note that there is a copy of P_4 joining Z_j and Z_{j+1} , and hence they are joined by a matching of size two. Linking these matchings from Z_1 through Z_t yields two vertex-disjoint copies of P_t . Index the elements of each Z_j as $\{z_j^1, z_j^2\}$ so that $\langle z_1^1, \dots, z_t^1 \rangle$ and $\langle z_1^2, \dots, z_t^2 \rangle$ are paths in H .

Let l be an odd integer with $7 \leq l \leq 2t - 3$; we show that H contains an l -cycle. For $3 \leq k \leq t - 1$, some vertex in Z_{k+1} is adjacent to both vertices in Z_k and completes a path Q_k of length $2k - 4$ with $\langle z_3^1, \dots, z_k^1 \rangle$ and $\langle z_3^2, \dots, z_k^2 \rangle$.

We consider two cases. First, suppose that the edge joining Z_1 and Z_3 has endpoints with different superscripts; we may assume that it is $z_1^1 z_3^2$. Let $k = (l + 1)/2$, and note that $4 \leq k \leq t - 1$. Now the path $\langle z_3^1, z_2^1, z_1^1, z_3^2 \rangle$ has length 3 and connects the endpoints of Q_k to complete a cycle of length $2k - 1$.

Otherwise, we may assume that the edge is $z_1^1 z_3^1$. Because there is a copy of P_4 joining Z_1 and Z_2 , either $z_1^1 z_2^2$ or $z_2^1 z_1^2$ is an edge in H ; we consider both subcases. Suppose that $z_1^1 z_2^2$ is an edge. With $k = (l + 1)/2$, the path $\langle z_3^1, z_1^1, z_2^2, z_3^2 \rangle$ has length 3 and connects the endpoints of Q_k to complete a cycle of length $2k - 1$. Otherwise, $z_2^1 z_1^2$ is an edge. Let $k = (l - 1)/2$, and note that $3 \leq k \leq t - 2$. The path $\langle z_3^1, z_1^1, z_2^1, z_2^2, z_3^2 \rangle$ has length 5 and connects the endpoints of Q_k to complete a cycle of length $2k + 1$. \square

Theorem 2.4.9. $R_\Delta(C_n) \leq 3458$.

Proof. If n is even, the result follows from Theorem 2.4.5. If $n \in \{3, 5\}$, then n is small enough so that a complete graph with sufficiently small maximum degree is a Ramsey host for C_n . We assume that n is odd and at least 7.

Let H_0 be a 36-regular (X_0, Y_0) -bigraph with girth at least $2t$, where t is sufficiently large in terms of n . Let H_1 be the graph on $V(H_0)$ where a pair of vertices is adjacent if they are at distance at most 2 in H_0 ; this graph is called the *square* of H_0 and is denoted by H_0^2 . Observe that H_1 contains H_0 as a subgraph. Because H_0 is bipartite, the extra edges in H_1 have both endpoints in the same part of the bipartition (X_0, Y_0) of H_0 . Note that each vertex in H_1 is incident to 36 edges that cross the bipartition and to exactly $36 \cdot 35$ edges inside its partite set because H_0 has large girth. Therefore, H_1 is 1296-regular. By Lemma 2.4.2, there is a 3456-regular $K_{3,3}^-$ -blowup H_2 of H_1 . We construct a 3458-regular graph H from H_2 by adding a triangle inside each cluster in H_2 .

Let X be the set of vertices in H that are members of clusters in H_2 corresponding to vertices in X_0 , and let Y be the set of vertices in H that are members of clusters in H_2 corresponding to vertices in Y_0 . Note that X and Y form a partition of $V(H)$, and the subgraph of H given by the cut $[X, Y]$ is a $K_{3,3}^-$ -blowup of H_0 .

We claim that $H \rightarrow C_n$. Consider a {red, blue}-edge-coloring of H . Apply Lemma 2.4.4 to the cut $[X, Y]$ to obtain a monochromatic P_4 -blowup of a path $\langle z_1, \dots, z_t \rangle$ in H_0 that alternates between X_0 and Y_0 . Without loss of generality, we assume that the P_4 -blowup

occurs in red. For each z_j , let Z_j be the corresponding cluster of three vertices in H , let z_j^1 and z_j^2 be the corresponding vertices in the P_4 -blowup, and let $Z'_j = \{z_j^1, z_j^2\}$.

Because each Z_j induces a complete graph, $z_j^1 z_j^2$ is an edge in H . We call these *internal edges*. Also, because z_j and z_{j+2} are at distance 2 in H_0 , it follows that $z_j z_{j+2}$ is an edge in H_1 and hence the corresponding clusters Z_j and Z_{j+2} are joined by a copy of $K_{3,3}^-$ in H . By Proposition 2.4.6, there is a (not necessarily monochromatic) copy of P_4 joining Z'_j and Z'_{j+2} . We call the edges in these copies of P_4 *skip edges*.

If there is a red internal edge in Z'_j , then we apply Lemma 2.4.7 to whichever of $\langle z_j, \dots, z_t \rangle$ and $\langle z_j, \dots, z_1 \rangle$ is longer to obtain a red copy of C_n . If there is a red skip edge joining Z'_j and Z'_{j+2} , then we apply Lemma 2.4.8 to whichever of $\langle z_j, \dots, z_t \rangle$ and $\langle z_{j+2}, \dots, z_1 \rangle$ is longer to obtain a red copy of C_n .

Finally, if each internal edge and each skip edge is blue, then H contains a blue P_4 -blowup of the path $\langle z_1, z_3, z_5, \dots, z_n \rangle$. Because internal edges are blue, Lemma 2.4.7 gives a blue copy of C_n . \square

Theorem 2.4.9 and Corollary 2.3.3 yield the following result.

Theorem 2.4.10. *The family of cycles is weakly R_Δ -bounded.*

We presently lack a proof that the family of cycles is R_Δ -bounded, but we believe it to be true.

Conjecture 2.4.11. The family of cycles is R_Δ -bounded.

Chapter 3

Online Degree Ramsey Theory

This chapter is based on joint work with J. Butterfield, T. Grauman, B. Kinnersley, C. Stocker, and D. B. West that appears in [22].

3.1 Introduction

Interesting combinatorial questions can arise when two players compete to achieve opposing goals. In this chapter, we combine the notion of degree Ramsey numbers from Chapter 2 with a game version of Ramsey theory. This model was introduced by Beck [11] and was explored in several contexts by Grytczuk, Hałuszczak, and Kierstead [53]. We describe it using two colors, but the game extends naturally to s colors.

Two players, Builder and Painter, play a game with a target graph G . During each round, Builder presents a new edge uv to Painter (u and/or v may be a vertex not yet used). Painter must color uv red or blue, and Painter's choice is permanent. Builder wins if a monochromatic copy of G arises. Painter wins if Builder cannot force this.

When Builder's moves are unrestricted, Ramsey's Theorem implies that Builder wins by presenting a large complete graph. As in parameter Ramsey theory, we may want to require Builder to keep the presented graph within a restricted class \mathcal{H} . This defines the *online Ramsey game* (G, \mathcal{H}) . Given G and \mathcal{H} , the question is which player has a winning strategy.

We say that (G, \mathcal{H}) is played *on* \mathcal{H} . Grytczuk, Hałuszczak, and Kierstead [53] showed that Builder wins on the class of k -colorable graphs when G is k -colorable. Also, Builder wins on the class of forests when G is a forest. They also showed for $G = K_3$ that Painter wins on outerplanar graphs but Builder wins on planar 2-degenerate graphs. On planar graphs, Builder wins when G is a cycle or is a 4-cycle plus one chord (a slight extension is that Builder can force any fixed cycle plus chords at any one vertex). They conjectured that on planar graphs, Builder wins if and only if G is outerplanar; this remains wide open.

For any graph parameter ρ , we define the *online ρ -Ramsey number* \mathring{R}_ρ of G to be the least k such that Builder has a winning strategy to force G when playing on the family $\{H: \rho(H) \leq k\}$. Because Builder has more power when Painter must color edges as they are presented, always $\mathring{R}_\rho(G; s) \leq R_\rho(G; s)$. The online ρ -Ramsey number can be much smaller than the corresponding ρ -Ramsey number. For example, the main result of [53] is

that $\mathring{R}_\chi(G) = \chi(G)$ for every graph. This implies that $\mathring{R}_\chi(K_n) = n$, but a theorem in [21] implies that $R_\chi(K_n) = R(K_n)$.

The notation $\tilde{r}(G)$ has been used for the online size Ramsey number. Grytczuk, Kierstead, and Prałat [54] proved that $\tilde{r}(P_n) \leq 4n - 7$ for $n \geq 2$ (and found the exact values for $n \leq 6$; see [92, 94, 93] for additional exact results on paths). They also proved $\tilde{r}(G) \geq \frac{1}{2}b(\Delta(G) - 1) + m$ when G has m edges and vertex cover number b . Using the latter, they proved that the maximum of $\tilde{r}(G)$ over trees with m edges is $\Theta(m^2)$.

Kierstead and Konjevod [66] studied an extension of online Ramsey games to s -uniform hypergraphs. A variant of online Ramsey games in which Builder is replaced with a sequence of random edges is studied in [47].

Online Ramsey theory on directed graphs has also been studied, even though the scope of non-trivial questions is limited in this context. If the target graph contains a directed cycle, then Painter easily avoids producing a copy of the target graph by fixing an arbitrary ordering of the vertices (new vertices can be inserted anywhere), using red for edges whose heads have a lower index than their tails, and using blue for other edges.

Nonetheless, when the target graph is acyclic, interesting questions arise. Ferrara and Tennenhouse considered a modification to the online Ramsey game where Builder presents an undirected edge, and Painter chooses the orientation. Ferrara and Tennenhouse [44] observed that Builder can force a directed path on n vertices by playing only $O(n \log n)$ edges. No superlinear lower bound on the number of edges Builder needs is known.

Here, we study the *online degree Ramsey number* $\mathring{R}_\Delta(G)$. Let \mathcal{S}_k be the class of graphs with maximum degree at most k ; $\mathring{R}_\Delta(G)$ is the least k such that Builder wins (G, \mathcal{S}_k) . Table 3.1 compares 2-color results for online degree Ramsey number of trees and cycles with the corresponding results for degree Ramsey number. The results on \mathring{R}_Δ for stars and large odd cycles appear here and in [22]; we provide citations for other results. The *double-star* $S_{a,b}$ is the tree with $a + b$ vertices having adjacent vertices of degrees a and b .

The value of $\mathring{R}_\Delta(G)$ is not known for any connected graph G that is not a tree or a cycle. The smallest such graphs are the graphs C_4^+ and $K_{1,3}^+$ obtained by adding one edge to the 4-cycle or the claw. It is proved in [22] that $\mathring{R}_\Delta(C_4^+) \in \{5, 6, 7\}$ and $\mathring{R}_\Delta(K_{1,3}^+) \in \{4, 5\}$; we present a proof of the former in Section 3.2.

In most of our results, the lower bounds rely on “greedy” strategies for Painter, in which Painter makes an edge red if and only if it keeps the red graph within a specified class, such as \mathcal{S}_k or the class of linear forests. Greedy Painters of both types are used to prove that $\mathring{R}_\Delta(G) \leq 3$ if and only if each component of G is a path or each component is a subgraph of the claw $K_{1,3}$ (Theorem 3.3.2). Another wrinkle in this proof is that when Builder wins (G_1, \mathcal{H}) and (G_2, \mathcal{H}) , it does not follow that Builder wins $(G_1 + G_2, \mathcal{H})$, where $G_1 + G_2$ is the disjoint union of G_1 and G_2 . For example, $\mathring{R}_\Delta(P_4) = \mathring{R}_\Delta(K_{1,3}) = 3$, but Builder cannot force $P_4 + K_{1,3}$ in \mathcal{S}_3 (see Theorem 3.3.2).

| G | $\mathring{R}_\Delta(G)$ | $R_\Delta(G)$ |
|-----------------------|---|--|
| Path P_n | 3, for $n \geq 4$ [22] | 4, for $n \geq 7$; [109] and [4] combined |
| Star $K_{1,m}$ | m | $\begin{cases} 2m - 1; & m \text{ odd} \\ 2m - 2; & m \text{ even} \end{cases}$ [21] |
| Double-star $S_{a,b}$ | $a + b - 1$ [22] | $\begin{cases} 2b - 1; & a = b \text{ odd} \\ 2 \max\{a, b\} - 2; & \text{otherwise} \end{cases}$ [67] |
| Tree G | $\leq 2\Delta(G) - 1$ (sharp) [22] | $\leq 4\Delta(G) - 4$ [60] |
| Even cycle C_{2k} | 4 [22] | $\begin{cases} 5; & k = 2 \\ \leq 96; & k \geq 2 \end{cases}$ [67] [60] |
| Odd cycle C_{2k+1} | $\begin{cases} \leq 5; & \text{always} \\ 4; & k \geq 344 \end{cases}$ [22] | $\begin{cases} \geq 5; & \\ \leq 3458; & \end{cases}$ [67] [60] |

Table 3.1: Results for online degree and degree Ramsey numbers.

Upper bounds require strategies for Builder. The upper bounds in Table 3.1 for paths, stars, and triangles follow by direct arguments involving induction and/or the pigeonhole principle. To obtain the more difficult upper bounds, we simplify the search for a Builder strategy. We prove (Theorem 3.5.6) that Builder may assume that Painter plays “consistently”, meaning that the color Painter assigns to an edge depends only on the components of the current edge-colored graph containing its endpoints. This reduction applies to the Ramsey game (G, \mathcal{H}) whenever \mathcal{H} is monotone and additive (closed under taking subgraphs or disjoint unions).

The main open question in online degree Ramsey theory is the natural online analogue of the big question in degree Ramsey theory.

Question 3.1.1. Is there a function f such that $\mathring{R}_\Delta(G) \leq f(k)$ when $\Delta(G) \leq k$?

The only progress toward this question is that $\mathring{R}_\Delta(G) \leq 6$ when $\Delta(G) \leq 2$ (see [22]). The question remains wide open even for graphs with maximum degree at most 3.

3.2 The greedy Painter, stars, and kites

In this section, we begin our study of online degree Ramsey numbers by presenting bounds for some simple graphs. Our first proposition is particularly simple and establishes the online degree Ramsey number of stars.

Proposition 3.2.1. $\mathring{R}_\Delta(K_{1,m}) = m$.

Proof. Because $K_{1,m}$ contains a vertex of degree m , Builder needs at least degree m in the host graph.

For the upper bound, we use induction on m . For $m > 1$, Builder first plays on \mathcal{S}_{m-1} to force $2m - 1$ disjoint monochromatic copies of $K_{1,m-1}$. By the pigeonhole principle, at least m copies share a color. Builder then plays a star $K_{1,m}$ whose leaves are the centers of those copies. The resulting graph is in \mathcal{S}_m , and $K_{1,m}$ is forced. \square

Upper bounds on $\mathring{R}_\Delta(G)$ arise from strategies for Builder. Lower bounds arise from strategies for Painter. In general, it is more difficult to obtain lower bounds. Nevertheless, a simple greedy strategy provides tight lower bounds in some cases.

Definition 3.2.2. Let \mathcal{F} be a family of graphs. The *greedy \mathcal{F} -Painter* colors each new edge red if the resulting red graph lies in \mathcal{F} ; otherwise, the edge is colored blue.

When $k = \Delta(G) - 1$, the greedy \mathcal{S}_k -Painter establishes a useful general lower bound.

Theorem 3.2.3. For every graph G , $\mathring{R}_\Delta(G) \geq \Delta(G) - 1 + \max_{uv \in E(G)} \min\{d(u), d(v)\}$.

Proof. Let $k = \Delta(G) - 1$, and let $t = \max_{uv \in E(G)} \min\{d(u), d(v)\}$. The greedy \mathcal{S}_k -Painter never makes a red G , because no vertex ever receives $\Delta(G)$ incident red edges. To force a blue G , Builder must force a blue copy of an edge xy such that $\min\{d_G(x), d_G(y)\} = t$. Making this edge blue requires k red edges already at at least one endpoint. In the blue G , each endpoint has at least d edges. Hence at x or y at least $k + t$ edges have been played. \square

Theorem 3.2.3 yields the lower bounds for trees in Table 3.1. The greedy \mathcal{S}_k -Painter also easily yields the lower bound in [54] for the online size Ramsey number ($\tilde{r}(G) \geq \frac{1}{2}bk + m$ when G has m edges, maximum degree $k + 1$, and vertex cover number b). Theorem 3.2.3 also provides a tight lower bound for the online degree Ramsey number of paths. We include its short proof for completeness.

Proposition 3.2.4 (Butterfield *et al.* [22]). $\mathring{R}_\Delta(P_n) = 3$.

Proof. The lower bound is a special case of Theorem 3.2.3. For the upper bound, we prove by induction on n that Builder playing on \mathcal{S}_3 can force a monochromatic path with at least n vertices such that no other edges have been played at the endpoints; the single edge has this property for $n = 2$. For $n \geq 3$, Builder first plays on \mathcal{S}_3 to force $n - 2$ such paths of the same color having at least $n - 1$ vertices each; let “red” be this color. Builder then plays a path Q with n vertices using one endpoint of each of these paths plus two new vertices as the endpoints of Q . If any edge of Q is red, then the desired path arises in red; if they are all blue, then Q becomes the desired path. \square

We conclude this section with bounds on the online degree Ramsey number of the kite C_4^+ which is obtained from C_4 by adding an edge.

Theorem 3.2.5. $5 \leq \mathring{R}_\Delta(C_4^+) \leq 7$.

Proof. Let $G = C_4^+$. By Theorem 3.2.3, $\mathring{R}_\Delta(G) \geq 5$. For the upper bound, Builder first obtains a monochromatic copy of $K_{1,4}$, say in red, in which each vertex has degree at most 4 (see Proposition 3.2.1). Let w be its center and x_1, x_2, x_3, x_4 be its leaves.

Using new vertices, Builder next presents $K_{1,7}$ with center u ; let E be its set of edges. If Painter colors five edges in E blue, then Builder presents K_5 on those five neighbors of u ; let E' be its edge set. If Painter gives blue to two incident edges in E' , then a blue kite arises. Otherwise, at most two edges in E' are blue and a red kite arises, since E' now contains a copy of K_4 in which at most one edge is not red. Hence Painter must make three edges in E red; let the neighbors of u in this red claw be v_1, v_2 , and v_3 .

Next, Builder presents the claw with center w and leaves v_1, v_2, v_3 . Now w has total degree 7 and each v_i has total degree 2. If this claw is all blue, then Builder wins by presenting the triangle on $\{v_1, v_2, v_3\}$. Hence we may assume that wv_1 is red.

Next, Builder presents the edges $\{v_1x_i : 1 \leq i \leq 4\}$ giving degree 6 to v_1 and degree (at most) 5 to each x_i . If Painter makes two of these edges red, then they form a red kite with w . Otherwise, we have a red claw and a blue claw with three common leaves among $\{x_1, x_2, x_3, x_4\}$. Builder now wins by presenting a triangle on those three vertices (giving them degree 7). \square

In fact, Builder can force C_4^+ on \mathcal{S}_5 against the \mathcal{S}_2 -greedy Painter (start with 4 copies of P_3 , which Painter colors red, and then play C_4^+ on the internal vertices of the paths). To improve the lower bound, we must find a better Painter strategy.

3.3 Graphs with small online degree Ramsey number

When Builder is limited to host graphs that have very small maximum degree, it is possible to analyze the online Ramsey game exactly. First, note that $\mathring{R}_\Delta(G) \leq 2$ is easily characterized, and it is trivial that $\mathring{R}_\Delta(G) \leq 1$ if and only if G is a matching.

Proposition 3.3.1. $\mathring{R}_\Delta(G) \leq 2$ if and only if each component of G is a subgraph of P_3 .

Proof. If $\mathring{R}_\Delta(G) \leq 2$, then $\Delta(G) \leq 2$. By Theorem 3.2.3, the vertices with degree 2 are nonadjacent. Hence each component of G is a subgraph of P_3 . Conversely, Builder forces any such graph by presenting enough disjoint triangles. \square

For C_3 , Theorem 3.2.3 yields only $\mathring{R}_\Delta(C_3) \geq 3$, but $\mathring{R}_\Delta(C_3) = 4$ (see [22] for the upper bound). To improve the lower bound, we use a different greedy Painter to characterize the graphs that Builder can force when playing on \mathcal{S}_3 ; the triangle is not among them.

Theorem 3.3.2. $\mathring{R}_\Delta(G) \leq 3$ if and only if each component of G is a path or each component of G is a subgraph of the claw $K_{1,3}$.

Proof. By Proposition 3.2.1 and Proposition 3.2.4, Builder can force the claw or any path on \mathcal{S}_3 . By the pigeonhole principle, Builder can thus force any disjoint union of subgraphs of $K_{1,3}$. Also Builder can force a path long enough to contain any specified disjoint union of paths.

For necessity, suppose that $\mathring{R}_\Delta(G) \leq 3$. Hence Builder can force a monochromatic copy of G in \mathcal{S}_3 . Consider a greedy \mathcal{L} -Painter, where \mathcal{L} is the family of linear forests (disjoint unions of paths). If G appears in red, then each component of G is a path.

Suppose that G appears in blue. A vertex v with degree 3 in G has three incident blue edges and hence no incident red edges, since Builder is playing in \mathcal{S}_3 . For the greedy \mathcal{L} -Painter to make these edges blue, each neighbor already has two incident red edges. Hence a blue claw must be a full component of the blue graph.

Next suppose that the blue G contains P_5 centered at a vertex v ; we have shown that the maximum degree in such a blue component is 2. Let u_1 and u_2 be the neighbors of v on the path. Since the red degree at each of u_1, v, u_2 is at most 1, the condition that prevented $u_i v$ from being red must be that there is already a red path from u_i to v . Since v already has three incident edges, the paths from u_1 and u_2 must merge before reaching v , but this contradicts the red graph being a linear forest.

It remains to show that Builder cannot force a graph containing both P_4 and $K_{1,3}$ in \mathcal{S}_3 . Since such a graph G has maximum degree at least 3 and has an edge with both endpoints having degree at least 2, Theorem 3.2.3 implies that $\mathring{R}_\Delta(G) \geq 4$. \square

Theorem 3.3.2 shows that Builder can force P_4 or $K_{1,3}$ in \mathcal{S}_3 but not their disjoint union. That is, the family of weighted graphs that Builder can force on the class of all graphs is not closed under disjoint union. It would be interesting to characterize the graphs with $\mathring{R}_\Delta(G) \leq 4$.

3.4 Weighted graphs

Strategies for Builder playing on \mathcal{S}_k often involve keeping track of how many edges have been played incident to each vertex. The argument we gave for paths (Proposition 3.2.4) had this flavor; to facilitate the induction we needed to maintain degree 1 at the leaves of the monochromatic path, but we could allow degree 3 at internal vertices. We can view the allowed degree at each vertex as a “capacity”.

Definition 3.4.1. A *c-weighted graph* is a graph G equipped with a nonnegative integer *capacity function* c on $V(G)$. A *copy* of a c -weighted graph G exists in a graph H if G embeds as a subgraph of H via an injection f such that $d_H(f(v)) \leq c(v)$ for all $v \in V(G)$.

When the capacity function is constant, say $c(v) = k$ for all $v \in V(G)$, we simply refer to the c -weighted graph G as a k -weighted graph. The statement that Builder wins (G, \mathcal{S}_k)

is equivalent to the statement that Builder can force the k -weighted graph G when playing on the unrestricted family of all graphs. Vertices that acquire more than k incident edges are forbidden from the desired monochromatic copy of G when Builder is restricted to \mathcal{S}_k .

It is easy to characterize the weighted stars that can be forced.

Proposition 3.4.2. *Builder can force a weighted star with m edges if and only if the center has capacity at least $2m - 1$ or if each vertex has capacity at least m .*

Proof. For sufficiency, the pigeonhole principle suffices when the center has capacity $2m - 1$, and $\mathring{R}_\Delta(K_{1,m}) \leq m$ (Proposition 3.2.1) suffices when each vertex has capacity at least m .

For necessity, consider a strategy for Builder to force a weighted star with m edges against the greedy \mathcal{S}_{m-1} -Painter. As in the proof of Theorem 3.2.3, Builder must force a blue $K_{1,m}$, and for each edge of this star one of the endpoints has $m - 1$ incident red edges. If the center of the blue $K_{1,m}$ has $m - 1$ incident red edges, then its degree is at least $2m - 1$. Otherwise, each leaf has one incident blue edge and $m - 1$ incident red edges, and hence each vertex has degree at least m . \square

In the proof of Proposition 3.2.4, we see that for each n_0 , Builder can force a path on n_0 or more vertices whose endpoints have degree 1 in the host graph. When Builder wishes to force cycles, it is convenient to obtain paths of specified lengths whose endpoints have small degree. Our next theorem shows that this is generally not possible. Painter can avoid providing Builder with a monochromatic path on 4 vertices whose endpoints have degree 1.

Theorem 3.4.3. *Builder cannot force a weighted 4-vertex path whose endpoints have unit capacity.*

Proof. We allow any degree on the central vertices; only the endpoints are restricted. We provide a strategy for Painter to avoid this weighted subgraph. Refer to vertices of degree 1 in the played graph H as *leaves* and vertices of degree at least 2 as *non-leaves*.

Painter maintains a partition of the nonisolated vertices of H into five sets. Sets S, R', B' will partition the leaves, and the non-leaves will lie in R or B . Let S consist of the vertices incident to isolated edges (which Painter makes all red). Other leaves lie in R' if the incident edge is red and in B' if it is blue. The non-leaves are partitioned into R and B so that edges with both endpoints in R are red, edges with both endpoints in B are blue, the neighbor of each vertex in R' is in B , and the neighbor of each vertex in B' is in R . Edges joining R to B may have either color.

If Painter can maintain this partition, then no monochromatic 4-vertex path arises with endpoints of degree 1. The color would force the endpoints to be both in R' or both in B' . In the former case, the central vertices both would be in B , but then the edge joining them would be blue. The latter case is symmetric.

When the game begins, the sets are empty and the partition is valid. We show how Painter responds to each move by Builder. Whenever a vertex of R' receives another incident edge, it moves to R ; since it already has an incident red edge to a vertex of B , this change causes no problem. Similarly, a vertex of B' receiving another incident edge moves to B . Once a vertex enters R or B , it does not move again.

Now consider edges that introduce a new vertex w with neighbor v . If v is new, then vw becomes red and both vertices enter S . If $v \in S \cup R' \cup B'$, then let u be the other neighbor of v . If $v \in S$, then vw is made red, v moves to B , and both w and u move to R' . If $v \in R'$, then vw becomes blue and w enters B' (v moves to R). Similarly, if $v \in B'$, then vw becomes red and w enters R' (v moves to B). If $v \in R$, then vw becomes blue and w enters B' . If $v \in B$, then vw becomes red and w enters R' .

In the remaining cases, neither endpoint is new. Again we add vw ; suppose that $w \in S$, and let u be the other neighbor of w . Wherever v is, vw is made blue, w moves to B , and u moves to R' . If $v \in S$, then v moves to B and the other neighbor of v moves to R' . If $v \in R' \cup B'$, then v is “promoted” to $R \cup B$ as described above. No change is made for v if $v \in R \cup B$. In all cases, making vw blue causes no problem since $w \in B$.

Hence we may assume that both endpoints of the new edge lie in $R' \cup B' \cup R \cup B$. An endpoint lying in $R' \cup B'$ is promoted to $R \cup B$ by removing the “prime”; as remarked earlier, this causes no trouble for its old incident edge. After making the promotion(s), it suffices to make the color of the new edge agree with the set label on at least one of its endpoints. \square

3.5 The consistent Painter

Strategies for Builder often involve pigeonholing arguments. Pigeonholing also applies to 2-edge-colored graphs. When Builder presents m edges, Painter can produce at most 2^m distinguishable 2-edge-colored graphs. By presenting isomorphic copies of the graph formed by these m edges, Builder can force many copies of some single pattern. Nevertheless, when strategies become more complicated and repeated copies of larger patterns are needed, citing the pigeonholing argument becomes unwieldy. Arguments simplify if Builder can assume that Painter plays “consistently”.

Definition 3.5.1. A Painter strategy is *consistent* if the color Painter chooses for an edge uv depends only on the 2-edge-colored component(s) containing u and v when uv arrives.

For example, a consistent Painter always colors an isolated triangle in the same way. If there are nonisomorphic ways to order the edges of a graph (such as K_4), then a consistent Painter may produce different colorings depending the order in which the edges arrive.

Our aim is to reduce the problem of proving that Builder wins to proving that Builder wins against consistent Painters. The argument can be given for the s -color model, but we state it only for two colors, red and blue. We need several technical notions.

Definition 3.5.2. Given a monotone additive family \mathcal{H} , an \mathcal{H} -strategy specifies a color for each pair (H, e) such that H is a 2-edge-coloring of a graph in \mathcal{H} and e is an edge not in H (either or both endpoints of e may be new vertices). An \mathcal{H} -list is an ordering of the edges of some graph in \mathcal{H} ; every initial segment of an \mathcal{H} -list forms a graph in \mathcal{H} . For each \mathcal{H} -list E and each \mathcal{H} -strategy \mathcal{A} , let $\mathcal{A}(E)$ denote the edge-colored graph that results when Builder presents E to \mathcal{A} . An edge-colored graph F contains another such graph F' if there is an injection of $V(F')$ into $V(F)$ that preserves edges and preserves their colors.

To reduce the Builder problem to winning against consistent Painters, we will show that for every Painter strategy there is a consistent strategy that does at least as well for Painter. That is, when \mathcal{A} is an \mathcal{H} -strategy, there is a consistent \mathcal{H} -strategy \mathcal{A}' such that any 2-edge-colored graph Builder can force against \mathcal{A}' can also be forced against \mathcal{A} . A special set of 2-edge-colored graphs will enable us to produce \mathcal{A}' .

Definition 3.5.3. A uv -augmentation of a 2-edge-colored graph H with nonadjacent vertices u and v is obtained by adding uv to H with color red or blue. Let \mathcal{H} be a monotone additive family. A class \mathcal{C} of connected 2-edge-colored graphs is \mathcal{H} -coherent if it contains K_1 and satisfies the following *augmentation property*: If H is a 2-edge-colored copy of a graph H' in \mathcal{H} , and H' has nonadjacent vertices u and v such that $H' + uv$ is a connected graph in \mathcal{H} and the component(s) of H are in \mathcal{C} , then \mathcal{C} contains a uv -augmentation of H .

An \mathcal{H} -coherent class \mathcal{C} yields a consistent \mathcal{H} -strategy \mathcal{A}' as follows. When an edge uv is added to the current 2-colored graph H , \mathcal{A}' consults \mathcal{C} to find which color on uv yields a uv -augmentation in \mathcal{C} for the component(s) of H containing the endpoints of the added edge. When both colors yield uv -augmentations, \mathcal{A}' always chooses the same one, say red.

Definition 3.5.4. Let \mathbf{C} be the class of connected 2-edge-colored (unlabeled) graphs; every 2-edge-coloring of a graph in \mathcal{H} is a multiset of elements of \mathbf{C} having finitely many distinguishable components, each with finite multiplicity. Given an \mathcal{H} -strategy \mathcal{A} , a 2-edge-colored graph H is \mathcal{A} -realizable if for some \mathcal{H} -list E , the outcome $\mathcal{A}(E)$ contains H . A family $\mathcal{C} \subseteq \mathbf{C}$ is \mathcal{A} -plentiful if, for every finite subset $C \subseteq \mathcal{C}$ and every positive integer n , the 2-edge-colored graph consisting of n components isomorphic to each element of C is \mathcal{A} -realizable.

In order to be \mathcal{H} -coherent for some monotone additive family \mathcal{H} , a family \mathcal{C} contained in \mathbf{C} must somehow be “large enough”. In order to be \mathcal{A} -plentiful for some \mathcal{H} -strategy \mathcal{A} , the family \mathcal{C} must somehow be “small enough”. We seek a family achieving both properties. We use Zorn’s Lemma in the following form: if every chain in a partial order P has an upper bound, then P has a maximal element.

Lemma 3.5.5. *If \mathcal{H} is a monotone additive family of graphs, and \mathcal{A} is an \mathcal{H} -strategy, then some family \mathcal{C} is both \mathcal{H} -coherent and \mathcal{A} -plentiful.*

Proof. Note first that $\{K_1\}$ is \mathcal{A} -plentiful. Also, if $\mathcal{C}_1, \mathcal{C}_2, \dots$ are \mathcal{A} -plentiful families with $\mathcal{C}_1 \subseteq \mathcal{C}_2 \subseteq \dots$, then the union of these families is also \mathcal{A} -plentiful, because the definition of \mathcal{A} -plentiful requires \mathcal{A} -realizability only of repeated copies of *finite* subsets, and each finite subset appears in some \mathcal{C}_j . It follows from Zorn's Lemma that there is a maximal \mathcal{A} -plentiful family \mathcal{C} containing K_1 . We claim that \mathcal{C} is \mathcal{H} -coherent.

By construction, $K_1 \in \mathcal{C}$. Consider a fixed 2-edge-colored graph H in \mathcal{C} , and let H' be its underlying uncolored graph in \mathcal{H} . Let u and v be nonadjacent vertices in H' such that $H'_u v \in \mathcal{H}$. Let H_1 and H_2 be the possible uv -augmentations of H (using red or blue on uv). If neither lies in \mathcal{C} , then by the maximality of \mathcal{C} , both $\mathcal{C} \cup \{H_1\}$ and $\mathcal{C} \cup \{H_2\}$ are not \mathcal{A} -plentiful. Hence there are positive integers t_1 and t_2 and finite sets $C_1, C_2 \subseteq \mathcal{C}$ such that the 2-edge-colored graphs $t_1(C_1 \cup \{H_1\})$ and $t_2(C_2 \cup \{H_2\})$ are not \mathcal{A} -realizable, where for $C \subseteq \mathcal{C}$ we use qC to denote the 2-edge-colored graph with q copies of each element of C as components.

Let $D = C_1 \cup C_2 \cup \{H\}$. Since D is a finite subset of \mathcal{C} , and \mathcal{C} is \mathcal{A} -plentiful, $2(t_1 + t_2 - 1)D$ is \mathcal{A} -realizable via some \mathcal{H} -list E . When E is presented, $\mathcal{A}(E)$ contains at least $2(t_1 + t_2 - 1)$ disjoint copies of H .

Since \mathcal{H} is additive, the list E' formed by adding to E the copies of uv in $2(t_1 + t_2 - 1)$ components isomorphic to H' is an \mathcal{H} -list; Builder may legally present these edges after E . Consider the first $t_1 + t_2 - 1$ of the added edges. Either \mathcal{A} colors at least t_1 of them red, or \mathcal{A} colors at least t_2 of them blue. In the first case, $t_1(C_1 \cup \{H_1\})$ is \mathcal{A} -realizable: we have obtained t_1 copies of H_1 , and at least $t_1 + t_2 - 1$ copies of H remain (needed if $H \in C_1$). In the second case, $t_2(C_2 \cup \{H_2\})$ is similarly \mathcal{A} -realizable. The contradiction implies that \mathcal{C} contains a uv -augmentation of H .

Essentially the same argument shows that \mathcal{C} contains a uv -augmentation of H when H consists of two 2-edge-colored components, each containing one of u and v . \square

Theorem 3.5.6. *If \mathcal{H} is a monotone additive family of graphs, and \mathcal{A} is an \mathcal{H} -strategy for Painter, then there is a consistent \mathcal{H} -strategy \mathcal{A}' such that for every \mathcal{H} -list E' , there is an \mathcal{H} -list E such that $\mathcal{A}(E) \supseteq \mathcal{A}'(E')$. That is, Builder can force against \mathcal{A} any monochromatic target that Builder can force against \mathcal{A}' .*

Proof. By Lemma 3.5.5, there is an \mathcal{H} -coherent, \mathcal{A} -plentiful family $\mathcal{C} \subseteq \mathbf{C}$. As described after Definition 3.5.3, from the \mathcal{H} -coherence of \mathcal{C} we define a consistent \mathcal{H} -strategy \mathcal{A}' .

Whenever \mathcal{A}' is given a new edge uv , the definition of \mathcal{C} being \mathcal{H} -coherent implies that the uv -augmentation chosen by \mathcal{A}' for the component being formed is in \mathcal{C} . Thus every component of $\mathcal{A}'(E')$ is in \mathcal{C} . Since \mathcal{C} is \mathcal{A} -plentiful, it now follows that $\mathcal{A}'(E') \subseteq \mathcal{A}(E)$ for some \mathcal{H} -list E . \square

To show that Builder wins (G, \mathcal{H}) , it now suffices to show that Builder can force a monochromatic G against any consistent \mathcal{H} -strategy for Painter. In particular, if some \mathcal{H} -

list results in a particular 2-edge-colored component, then Builder can recreate another copy of that component by playing an isomorphic list of edges on a new set of vertices.

As mentioned earlier, the argument applies to any monotone additive family \mathcal{H} . For example, to prove sufficiency in the conjecture in [53], one need only show that Builder can force any outerplanar graph when playing on planar graphs against a consistent Painter. In [22], the consistent Painter is invoked to simplify the Builder's strategy for obtaining monochromatic trees.

3.6 Large odd cycles

In this section, we examine the online degree Ramsey number of cycles. From Theorem 3.3.2, it follows that $\mathring{R}_\Delta(C_n) \geq 4$ for each n . In [22], we prove that $\mathring{R}_\Delta(C_n) = 4$ when n is even, large, or $n = 3$. Here, we present the proof that $\mathring{R}_\Delta(C_n) = 4$ when n is large and odd. We begin with two lemmas that describe how Builder can extend given strategies to force larger structures.

Lemma 3.6.1 (Butterfield *et al.* [22]). *Against a consistent Painter, let F be a weighted graph Builder can force in red, with $u \in V(F)$ having capacity c . Form F' from $F + F$ by adding an edge joining the two copies of u and changing the capacity at its endpoints to $c + 2$. If n is odd, then Builder can force a red F' or a blue t -weighted n -cycle, where $t = c + 2$.*

Proof. Builder forces n copies of F in red, respecting capacities. Builder then plays an n -cycle on the copies of u . If these edges are all blue, then they form a blue t -weighted n -cycle. Otherwise, a red F' arises. \square

Lemma 3.6.2. *Against a consistent Painter, let F be a weighted graph Builder can force in red, with $u \in V(F)$ having capacity c . Let F' be the weighted graph obtained from F by changing the capacity at u to $c + 2$ and adding a new vertex v with capacity 2 adjacent only to u . Let $t = \max\{c + 2, 4\}$. If n is odd, then Builder can force a red F' or a monochromatic t -weighted n -cycle.*

Proof. Builder forces $n(n - 1)/2$ red copies of F against this Painter. Next, Builder plays an $n(n - 1)$ -cycle, alternating between copies of u and new vertices. Using red on any such edge produces a red copy of F' ; otherwise, there is a blue $n(n - 1)$ -cycle in which alternate vertices have degree 2. This cycle decomposes into n paths P_1, \dots, P_n , each consisting of $n - 1$ consecutive edges. Since each path has even length, we may assume that the endpoints of each path have degree 2.

Next, for each path P_j , Builder plays the edge joining its two endpoints. Using blue for any such edge creates a blue n -cycle (respecting capacities). Otherwise, these edges form a red n -cycle in which every vertex has degree 4. \square

Lemma 3.6.3. *Every tree T has a vertex w such that each component of $T - w$ contains at most half the leaves of T . Moreover, if T is not a path, then some such vertex has degree at least 3.*

Proof. Construct a directed graph D on $V(T)$ as follows. If uv is an edge in T and the component of $T - u$ containing v has more than half the leaves of T , then we include an edge from u to v in D . Since D cannot have both uv and vu , it is acyclic. For each vertex w having outdegree zero in D , each component of $T - w$ has at most half the leaves in T .

If w has degree 2 in T , then let P be a maximal path in T whose internal vertices all have degree 2 in T . If T is not a path, then at least one endpoint of P has degree at least 3. Let w' be such an endpoint of P . The components of $T - w'$ yield a partition of $V(T)$ that refines the partition given by components of $T - w$, so we can use w' instead of w . \square

Our next lemma shows that if Builder can force monochromatic trees with many leaves and small diameter, then Builder can force monochromatic odd cycles. We will need Dirac's Theorem [28], which states that every n -vertex graph with minimum degree at least $n/2$ has a spanning cycle.

Lemma 3.6.4. *Let T be a weighted tree whose leaves have capacity 2 and non-leaves have capacity 4. Let $d = \text{diam}(T)$, and let l be the number of leaves in T . Let G be a 4-weighted n -cycle with n odd. If $2d + 1 \leq n \leq l$ and Builder can force T , then Builder can force G .*

Proof. If $l > n$, then we iteratively delete leaves to obtain a subtree \hat{T} of T with n leaves; since $\text{diam}(\hat{T}) \leq \text{diam}(T)$, we have $2\text{diam}(\hat{T}) + 1 \leq n$. If Builder can force T , then Builder can force \hat{T} . Hence we may assume that $l = n$.

Since $n \geq 3$, T is not a path. Hence there is a vertex w in T with degree at least 3 such that every component of $T - w$ has at most half of the leaves, by Lemma 3.6.3. Since $\text{diam}(T) \leq (n - 1)/2$, each leaf has distance at most $(n - 1)/2$ from w . Construct a tree T' from T as follows. For every leaf u of T having distance less than $(n - 1)/2$ from w , append a path at u to reach a new leaf with distance $(n - 1)/2$ from w . The resulting tree T' has the same number of leaves as T , and every leaf has distance $(n - 1)/2$ from w . Weight T' by giving capacity 2 to each leaf and capacity 4 to each non-leaf.

Against a consistent Painter, we may assume that Builder can force T in red. Note that T' arises from T by repeatedly increasing the capacity of a leaf to 4, appending an edge there, and giving the new leaf capacity 2. By repeated application of Lemma 3.6.2, Builder can force a red T' or a monochromatic G .

If Painter avoids the monochromatic G , then a red T' results. By construction, each leaf of T' has distance $(n - 1)/2$ from w , so the distance in T' between leaves belonging to different components of $T' - w$ is $n - 1$. There are $d_T(w)$ such components, each with at most $n/2$ leaves. Create an auxiliary graph F whose vertices are the n leaves of T' ; two such

vertices are adjacent in F if they lie in different components of $T' - w$. Since each component has at most $n/2$ leaves, Dirac's Theorem [28] implies that F has a spanning cycle.

Builder now plays the edges of the spanning cycle in F . Using red for one of these edges completes a red n -cycle in which every vertex has degree at most 4. Otherwise, there is a 4-weighted n -cycle in blue. \square

Next, we show how Builder can use Lemma 3.6.1 to force trees with many leaves and small diameter, permitting the application of Lemma 3.6.4. We use $\text{dist}_G(x, y)$ to denote the distance between vertices x and y in a graph G .

Lemma 3.6.5. *Let G be a 4-weighted cycle. For each $r \geq 0$, there is a weighted tree T_r with leaves of capacity 2 and non-leaves of capacity 4 such that Builder can force a monochromatic graph in $\{T_r, G\}$ and*

1. T_r has $2(4^r + 1)$ leaves,
2. $\text{diam}(T_r) \leq 11 \cdot 2^r - 8$, and
3. T_r has disjoint pairs $\{x, y\}$ and $\{x', y'\}$ of leaves such that $\text{dist}_{T_r}(x, y) = \text{dist}_{T_r}(x', y') = 2$.

Proof. We use induction on r . For $r = 0$, let T_0 be the 6-vertex double-star $S_{3,3}$, with capacity 2 on leaves and capacity 4 on non-leaves. Against a consistent Painter, we may assume that every isolated triangle yields a red 2-weighted P_3 . It follows from Lemma 3.6.1 that Builder can force a red T_0 when Painter avoids G . Also, T_0 satisfies the other specified properties.

For $r \geq 1$, we construct T_r from T_{r-1} . Let x and y be leaves in T_{r-1} with $\text{dist}_{T_{r-1}}(x, y) = 2$. Form a tree T from two copies of T_{r-1} by adding an edge joining the two copies of x . By Lemma 3.6.1, Builder can force a red T . In T , let y^* be one of the two copies of y . Form T_r from two copies of T by adding an edge joining the two copies of y^* . By Lemma 3.6.1, Builder can force a red T_r . The construction of T_1 from T_0 is shown in the figure below.

With each of the two steps in the construction of T_r from T_{r-1} , we doubled the number of leaves and then killed two leaves. Hence if T_{r-1} has l leaves, then T_r has $4l - 6$ leaves.

Let $d = \text{diam}(T_{r-1})$. Since every vertex of T_{r-1} has distance at most d from x , every vertex in T has distance at most $d + 1$ from each copy of x . Since y^* has distance 2 in T from one copy of x , the distance in T_r from each vertex to one specified copy of y^* is at most $d + 4$. By the triangle inequality, $\text{diam}(T_r) \leq 2d + 8$. Hence $\text{diam}(T_r) \leq 2[11 \cdot 2^{r-1} - 8] + 8 = 11 \cdot 2^r - 8$.

Finally, note that each copy of T_{r-1} has two disjoint pairs of leaves at distance 2. The construction of T destroys one pair from each copy, so T has two such pairs. These pairs in both copies of T survive in T_r . \square

We now have the tools to prove our main result about odd cycles.

Theorem 3.6.6. *If n is odd and $337 \leq n \leq 514$ or $n \geq 689$, then $\mathring{R}_\Delta(C_n) = 4$.*

Proof. By Theorem 3.3.2, $\mathring{R}_\Delta(C_n) > 3$.

For any r , Lemma 3.6.5 implies that Builder can force a graph in $\{T_r, G\}$, where G is the 4-weighted C_n . If $2(11 \cdot 2^r - 8) + 1 \leq n \leq 2(4^r + 1)$, then by Lemma 3.6.4 Builder can force G . The theorem now follows by showing that

$$\bigcup_{r \geq 0} \{2(11 \cdot 2^r - 8) + 1, \dots, 2(4^r + 1)\} = \{337, \dots, 514\} \cup \{689, \dots\}.$$

The interval of suitable values of n is empty when $r \leq 3$. For $r = 4$, the interval is $\{337, \dots, 514\}$. The interval for $r = 5$ begins at 689. For $r \geq 5$, the end of the interval for r is after the beginning of the interval for $r + 1$, so all larger values are covered. \square

Theorem 3.6.6 and results in [22] imply that the online degree Ramsey number of every cycle is 4, with the possible exception of a finite number of odd cycles that have online degree Ramsey number 5. Recently, Rolnick [97] proved that $\mathring{R}_\Delta(C_n) = 4$ for each n .

Chapter 4

Subtrees with Few Labeled Paths

In this chapter, we prove several quantitative Ramseyan results involving ternary complete trees with $\{0, 1\}$ -labeled edges where we attempt to find a complete binary subtree with as few labels as possible along its paths. One of these is used to answer a question of Simpson's in computability theory.

This chapter is based on joint work with R. G. Downey, N. Greenberg, and C. G. Jockusch that appears in [30].

4.1 Introduction

There have been many fruitful interactions between combinatorics and computability theory. Examples include new combinatorial proofs of classical results such as Miletì's proof of the canonical Ramsey theorem [79], Montalbán's newly devised invariants for infinite linear orderings [81], Kierstead's online version of Dilworth's Theorem [65], and Füredi *et al.* on inverting the difference operator [49]. This work is another example of such an interaction.

We study edge-labelings of rooted trees. A tree is *ternary* if each non-leaf has 3 children and *binary* if each non-leaf has 2 children. A tree is *complete* if all leaves are at the same distance from the root, and the *depth* of a complete tree is the distance between the root and a leaf. The *level* of an edge is the depth of the endpoint farthest from the root. If T is a complete ternary tree of depth n , we define $\mathcal{B}(T)$ to be the set of all binary subtrees of T that are complete with depth n . A tree T is *edge-labeled* if each edge in T is assigned a label from the set $\{0, 1\}$. We define \mathcal{T}_n to be the set of all ternary, complete, edge-labeled trees of depth n .

If $T \in \mathcal{T}_n$, r is the root of T , and σ is a leaf in T , then reading the elements along the path from r to σ in T gives a *path-label* $x \in \{0, 1\}^n$, and we say that σ has path-label x . We define $L(T)$ to be the set of all path-labels in T . Given $T \in \mathcal{T}_n$, we wish to find a subtree $S \in \mathcal{B}(T)$ that minimizes $|L(S)|$. For each $T \in \mathcal{T}_n$, let $f(T) = \min\{|L(S)| : S \in \mathcal{B}(T)\}$, and for each n , let $f(n) = \max\{f(T) : T \in \mathcal{T}_n\}$.

The combinatorial thrust of this chapter is to study the behavior of $f(n)$ as n grows. In Section 4.2, we show that $\lim_{n \rightarrow \infty} (f(n))^{1/n}$ exists; our bounds on $f(n)$ imply that this limit has a value between $2^{1/\lg 3} \approx 1.548$ and 2. In Section 4.3, we show that if $c < \sqrt{\lg(4/3)} \approx$

0.644, then there is a constant γ such that $f(n) \leq \gamma 2^{n-c\sqrt{n}}$. Consequently, the ratio $f(n)/2^n$ tends to zero as n grows. This result has the following Ramsey interpretation: for large n , every edge-labeled complete ternary tree of depth n admits a complete binary subtree of depth n whose path-labels constitute an arbitrarily small fraction of the space of all possible path-labels. In Section 4.4, we prove that $f(n) \geq 2^{(n-3)/\lg 3}$. Our techniques lead to a solution of a problem in computability theory and effective randomness.

In his survey paper [101] on mass problems and randomness, Simpson asked whether for all $k \geq 3$, the Medvedev degree of DNR_k bounds the Medvedev degree of every Π_1^0 class of positive measure. We give precise definitions in Section 4.5, but the gist of the question concerns comparing the computational difficulty of diagonalization with a constant bound with that of constructing a set which is effectively random. The full background and motivation for this question, which we answer in the negative in this chapter, can be found in Section 4.5.

4.2 Some facts and a question about f

We begin by collecting a few simple facts about $f(n)$. The following recursive bounds on $f(n)$ are instructive.

Proposition 4.2.1.

1. If n is a positive integer, then $f(n+1) \leq 2f(n)$.
2. If r and s are positive integers, then $f(r+s) \geq f(r)f(s)$.

Proof. To prove (1), let $T \in \mathcal{T}_{n+1}$ be a tree with root r , and let T_0 and T_1 be subtrees of T rooted at two children of r . Since $T_0, T_1 \in \mathcal{T}_n$, by induction each T_i has a binary subtree $S_i \in \mathcal{B}(T_i)$ containing at most $f(n)$ path-labels, and combining these subtrees with the root of T yields a binary subtree of T with at most $2f(n)$ path-labels.

To prove (2), let $R \in \mathcal{T}_r$ be a tree in which each $R' \in \mathcal{B}(R)$ contains at least $f(r)$ path-labels, and let $S \in \mathcal{T}_s$ be a tree in which each $S' \in \mathcal{B}(S)$ contains at least $f(s)$ path-labels. Obtain $T \in \mathcal{T}_{r+s}$ by attaching a copy of S at each leaf in R . Each binary subtree of T contains at least $f(r)f(s)$ labels. \square

It is instructive to consider the behavior of f when n is small. It is clear that $f(1) = 1$. It is also straightforward to establish the value of $f(2)$.

Proposition 4.2.2. $f(2) = 2$.

Proof. By Proposition 4.2.1, it follows that $f(2) \leq 2f(1) = 2$. For the lower bound, see Figure 4.1. \square

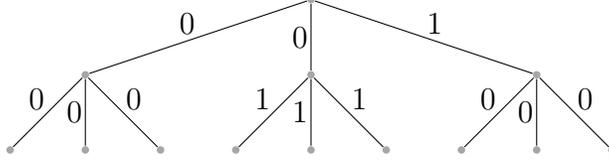


Figure 4.1: An example of a tree $T \in \mathcal{T}_2$ with $f(T) = 2$.

The following proposition about trees T of depth 2 that achieve equality in $f(T) \leq 2$ is useful in establishing more values of $f(n)$.

Proposition 4.2.3. *If $T \in \mathcal{T}_2$ and $f(T) = 2$, then there exists $X \subseteq \{0, 1\}^2$ with $|X| = 3$ such that for each $Y \subseteq X$ with $|Y| = 2$, there is a binary subtree $S \in \mathcal{B}(T)$ with $L(S) = Y$.*

Proof. Let $T \in \mathcal{T}_2$ with $f(T) = 2$ and root r . Let T_1, T_2 , and T_3 be the subtrees of T rooted at the children of r , and let x_j be a path-label in $\{0, 1\}^2$ such that at least two leaves in T_j have path-label x_j . Note that the x_j are distinct since $f(T) = 2$. The proposition follows with $X = \{x_1, x_2, x_3\}$. \square

Given two trees T_1 and T_2 , it is also useful to find binary subtrees $S_1 \in \mathcal{B}(T_1)$ and $S_2 \in \mathcal{B}(T_2)$ such that $L(S_1) \cup L(S_2)$ is small.

Proposition 4.2.4. *If $T_1, T_2 \in \mathcal{T}_2$, then either $\min\{f(T_1), f(T_2)\} = 1$ or there are binary subtrees $S_j \in \mathcal{B}(T_j)$ such that $|L(S_1) \cup L(S_2)| \leq 2$.*

Proof. We may assume that $f(T_1) = f(T_2) = 2$. Apply Proposition 4.2.3 to T_1 and again to T_2 to obtain X_1 and X_2 respectively. Because $|X_1| = |X_2| = 3$ and $X_1, X_2 \subseteq \{0, 1\}^2$, we have that $|X_1 \cap X_2| \geq 2$. Choose $Y \subseteq X_1 \cap X_2$ so that $|Y| = 2$. Proposition 4.2.3 implies that there exist $S_j \in \mathcal{B}(T_j)$ with $L(S_j) = Y$. \square

Corollary 4.2.5. *If $T_1, T_2 \in \mathcal{T}_2$, then there are binary subtrees $S_j \in \mathcal{B}(T_j)$ such that $|L(S_1) \cup L(S_2)| \leq 3$.*

Proof. By Proposition 4.2.2, $f(T_1) \leq 2$ and $f(T_2) \leq 2$. If $\min\{f(T_1), f(T_2)\} = 1$, then there are binary subtrees S_j with $|L(S_1)| + |L(S_2)| \leq 3$, which suffices. Otherwise Proposition 4.2.4 implies that there are binary subtrees S_j with $|L(S_1) \cup L(S_2)| \leq 2$. \square

Proposition 4.2.6. $f(3) = 3$.

Proof. Let $T \in \mathcal{T}_3$ with root r . Let v_1, v_2, v_3 be the children of r ; we may assume that rv_1 and rv_2 have the same edge-label. Let T_1 and T_2 be the subtrees of T of depth 2 rooted at v_1 and v_2 respectively. Corollary 4.2.5 implies that there are binary subtrees $S_j \in \mathcal{B}(T_j)$ such that $|L(S_1) \cup L(S_2)| \leq 3$. Combining S_1 and S_2 with r , we obtain a binary subtree S of T with $|L(S)| \leq 3$. It follows that $f(3) \leq 3$. For the lower bound, see Figure 4.2. \square

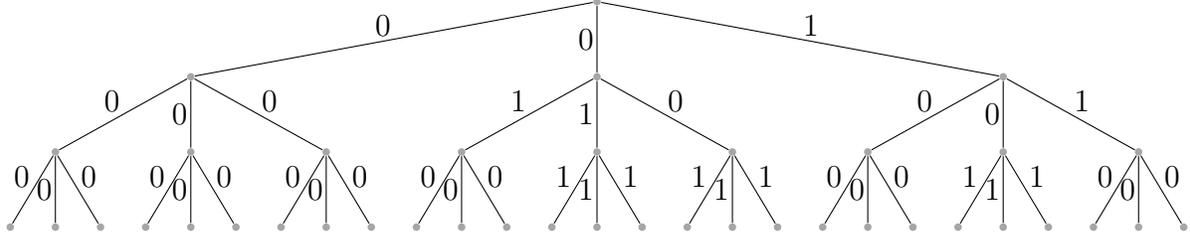


Figure 4.2: An example of a tree $T \in \mathcal{T}_3$ with $f(T) = 3$.

Next, we examine the structure of trees $T \in \mathcal{T}_3$ that achieve equality in $f(T) \leq 3$.

Proposition 4.2.7. *Let T be a tree in \mathcal{T}_3 with root r and $f(T) = 3$, and let v_1, v_2, v_3 be the children of r . The edges rv_1, rv_2 , and rv_3 do not all have the same label.*

Proof. Suppose for a contradiction that rv_1, rv_2 , and rv_3 all have edge-label 0. Let T_j be the subtree of T of depth 2 rooted at v_j . Each T_j has a binary subtree S_j in which the edges incident to the root of T_j have the same edge-label. At least two of S_1, S_2 , and S_3 have the same edge-label on all 4 edges at level 1. Combining these with r yields a subtree $S \in \mathcal{B}(T)$ in which all path-labels share the first 2 coordinates. It follows that $|L(S)| \leq 2$, contradicting that $f(T) = 3$. \square

Proposition 4.2.8. *Let T be a tree in \mathcal{T}_3 with root r and $f(T) = 3$. Let v_1, v_2, v_3 be the children of r , indexed so that rv_1 and rv_2 have the same edge-label, and rv_3 has the opposite edge-label. For $j \in \{1, 2, 3\}$, let T_j be the subtree of T rooted at v_j . We have $\min\{f(T_1), f(T_2)\} = 1$ and $\max\{f(T_1), f(T_2)\} = f(T_3) = 2$.*

Proof. If $f(T_1) = f(T_2) = 2$, then Proposition 4.2.4 implies that there are binary subtrees $S_j \in \mathcal{B}(T_j)$ such that $|L(S_1) \cup L(S_2)| = 2$. Combining these with r would yield a binary subtree $S \in \mathcal{B}(T)$ with $|L(S)| \leq 2$. It follows that $\min\{f(T_1), f(T_2)\} = 1$. We may assume that $f(T_1) = 1$.

It follows that $f(T_2) = f(T_3) = 2$. Indeed, if $f(T_j) = 1$ for some $j \in \{2, 3\}$, then a binary subtree of T_j with one path-label and a binary subtree of T_1 with one path-label would combine with r to produce a binary subtree $S \in \mathcal{B}(T)$ with $|L(S)| \leq 2$. \square

Lemma 4.2.9. *Let $T_1, T_2 \in \mathcal{T}_3$. There are binary subtrees $S_j \in \mathcal{B}(T_j)$ such that $|L(S_1) \cup L(S_2)| \leq 4$.*

Proof. We may assume that $f(T_1) = 3$, since the Proposition readily follows when both $f(T_1) \leq 2$ and $f(T_2) \leq 2$. For $j \in \{1, 2\}$, let r_j be the root of T_j , and let $v_{j,1}, v_{j,2}, v_{j,3}$ be the children of r_j . By Proposition 4.2.7, we may index $v_{1,1}, v_{1,2}$, and $v_{1,3}$ so that $r_1v_{1,1}$ and $r_1v_{1,2}$ have the same edge-label, which we may assume is 0, and $r_1v_{1,3}$ has the opposite edge-label 1. By Proposition 4.2.8, we may assume that $f(T_{1,1}) = 1$ and $f(T_{1,2}) = f(T_{1,3}) = 2$.

Consider the labels on edges incident to r_2 . If at least two of these are 0, then there are subtrees $S_j \in \mathcal{B}(T_j)$ such that all 4 edges at level 1 have label 0, which implies that $|L(S_1) \cup L(S_2)| \leq 4$. Hence we may assume that $r_2v_{2,1}$ and $r_2v_{2,2}$ have edge-label 1.

Apply Corollary 4.2.5 to obtain binary subtrees $S_{2,j} \in \mathcal{B}(T_{2,j})$ with $|L(S_{2,1}) \cup L(S_{2,2})| \leq 3$. Note that there is a path-label $x \in \{0,1\}^2$ such that $x \notin L(S_{2,1}) \cup L(S_{2,2})$. Because $f(T_{1,3}) = 2$, Proposition 4.2.3 implies that there exists $S_{1,3} \in \mathcal{B}(T_{1,3})$ such that $x \notin L(S_{1,3})$. It follows that $|L(S_{1,3}) \cup L(S_{2,1}) \cup L(S_{2,2})| \leq 3$.

Choose $S_{1,1} \in \mathcal{B}(T_{1,1})$ so that $|L(S_{1,1})| = 1$. Combine $S_{1,1}$ and $S_{1,3}$ with r_1 to form S_1 , and combine $S_{2,1}$ and $S_{2,2}$ with r_2 to form S_2 . Because $r_1v_{1,1}$ has edge-label 0 and $r_1v_{1,3}$, $r_2v_{2,1}$, and $r_2v_{2,2}$ have edge-label 1, we have that $|L(S_1) \cup L(S_2)| = |L(S_{1,1})| + |L(S_{1,3}) \cup L(S_{2,1}) \cup L(S_{2,2})| \leq 4$. \square

Lemma 4.2.10. $f(4) = 4$.

Proof. Let $T \in \mathcal{T}_4$ with root r . Let v_1, v_2, v_3 be the children of r ; we may assume that rv_1 and rv_2 have the same edge-label. Let T_1 and T_2 be the subtrees of T of depth 3 rooted at v_1 and v_2 respectively. Lemma 4.2.9 implies that there are binary subtrees $S_j \in \mathcal{B}(T_j)$ such that $|L(S_1) \cup L(S_2)| \leq 4$. Combining S_1 and S_2 with r , we obtain a binary subtree S of T with $|L(S)| \leq 4$. It follows that $f(3) \leq 3$. By Proposition 4.2.1 and Proposition 4.2.2, we have that $f(4) \geq f(2)f(2) = 4$. \square

Summarizing our lemmas, we obtain the following.

Theorem 4.2.11. *If $n \leq 4$, then $f(n) = n$. Also, $6 \leq f(5) \leq 8$.*

Proof. Proposition 4.2.2, Proposition 4.2.6, and Lemma 4.2.10 imply the first statement. For the second, note that Proposition 4.2.1 implies that $f(5) \geq f(3)f(2) = 6$ and $f(5) \leq 2f(4) = 8$. \square

The exact value of $f(n)$ is unknown when $n \geq 5$. Proposition 4.2.1 has further consequences. First, the upper bound $f(n+1) \leq 2f(n)$ shows that $f(n)/2^n$ is a non-increasing sequence, and because $f(n) \geq 0$, it follows that $\lim_{n \rightarrow \infty} f(n)/2^n$ exists. Indeed, we shall see that this limit is zero. Another consequence of Proposition 4.2.1 is that $\lim_{n \rightarrow \infty} (f(n))^{1/n}$ exists.

Proposition 4.2.12. *If $a_n = (f(n))^{1/n}$ and $\beta = \sup \{a_n\}$, then $\lim a_n = \beta$.*

Proof. Note that $\beta \leq 2$ since $f(n) \leq 2^n$. Fix $\varepsilon > 0$ and choose m so that $a_m \geq \beta - \varepsilon/2$. Let n be large, and divide n by m to get a quotient q and remainder r . Iteratively applying Proposition 4.2.1, we have that

$$f(n) \geq (f(m))^q f(r) = (a_m)^{mq} f(r) \geq (a_m)^{n-r}.$$

Hence, we have that $a_n \geq (a_m)^{1-r/n}$. Because $(a_m)^{1-r/n} \rightarrow a_m$ as $n \rightarrow \infty$, it follows that there exists n_0 such that $n \geq n_0$ implies that $a_n \geq a_m - \varepsilon/2$. Therefore, for each $n \geq n_0$, we have that $\beta - \varepsilon \leq a_n \leq \beta$.

This result can also be proved by noting that $g(n) = \log f(n)$ is a superadditive function, i.e. $g(m+n) \geq g(m) + g(n)$ for all positive integers m, n . Also $g(n)/n$ is bounded. It then follows from a result known as Fekete's Lemma (see [91], #98, page 23, solution on page 198) that $\lim_n g(n)/n$ exists and equals $\sup g(n)/n$. Restating this in terms of f completes the proof. \square

It follows from Proposition 4.2.12 and $f(3) = 3$ that

$$\lim_{n \rightarrow \infty} (f(n))^{1/n} \geq (f(3))^{1/3} = \sqrt[3]{3} \geq 1.442.$$

We shall see in Corollary 4.4.4 that

$$\lim_{n \rightarrow \infty} (f(n))^{1/n} \geq 2^{\frac{1}{\lg 3}} \geq 1.548.$$

On the other hand, the best known upper bound is the trivial bound

$$\lim_{n \rightarrow \infty} (f(n))^{1/n} \leq 2.$$

This leads to the main open problem regarding bounds on $f(n)$.

Question 4.2.13. What is $\lim_{n \rightarrow \infty} (f(n))^{1/n}$?

4.3 An upper bound on $f(n)$

We begin collecting results needed to establish our upper bound on $f(n)$. The following proposition is central to the task at hand. It is implicit in the proof of Theorem 6 of [61], which is a sort of forerunner of our Theorem 4.5.2. It was also stated explicitly by Robert Goldblatt in [50] (bottom of page 561) where it was applied to solve a problem in modal logic. We include its short proof for completeness.

Proposition 4.3.1. *Let T be a complete ternary tree of depth n . If each leaf in T is colored red or blue, then there exists $S \in \mathcal{B}(T)$ such that all leaves in S share a common color.*

Proof. Let r be the root of T and let T_1, T_2 , and T_3 be the subtrees of T rooted at the children of r . By induction, each tree T_j contains a subtree $S_j \in \mathcal{B}(T_j)$ in which all leaves share a common color. By the pigeonhole principle, at least two of the S_j contain leaves of the same color. Combining these with the root of T , we obtain $S \in \mathcal{B}(T)$ as required. \square

Proposition 4.3.1 has a very useful consequence. When $X \subseteq \{0, 1\}^n$, we let \overline{X} be the complementary set $\{0, 1\}^n \setminus X$.

Corollary 4.3.2. *If $T \in \mathcal{T}_n$ and X is a subset of $\{0, 1\}^n$, then there exists $S \in \mathcal{B}(T)$ such that either $L(S) \subseteq X$ or $L(S) \subseteq \overline{X}$.*

Proof. Label a leaf σ in T red if the path-label from the root of T to σ is in X and blue otherwise. By Proposition 4.3.1, there exists $S \in \mathcal{B}(T)$ such that all leaves share a common color. If this color is red, then $L(S) \subseteq X$. Otherwise, $L(S) \subseteq \overline{X}$. \square

Our strategy to bound $f(n)$ from above is as follows. We prove our bound by induction. To prove that $f(n)$ is small, we use that $f(m)$ is small for a carefully chosen number m that is less than n .

Consider $T \in \mathcal{T}_n$. We first find a complete binary subtree S' of depth m such that $|L(S')| \leq f(m)$. For each leaf σ in S' , let T_σ be the subtree of T rooted at σ . Note that each T_σ is a member of \mathcal{T}_{n-m} . To extend S' to a complete binary subtree of depth n , we wish to find a family of complete binary subtrees $S_\sigma \in \mathcal{B}(T_\sigma)$ such that $|\bigcup_\sigma L(S_\sigma)|$ is as small as possible.

The key for this process is arguing that given a family of edge-labeled ternary trees, we can find binary subtrees of each such that the total number of path-labels used in all of the binary subtrees is small. Corollary 4.3.2 gives some control over the path-labels that appear in the binary subtrees. In order to find the binary subtrees, we apply Corollary 4.3.2 numerous times with different subsets of $\{0, 1\}^n$. We are particularly interested in applying Corollary 4.3.2 to families of subsets of $\{0, 1\}^n$ with a certain structure.

Definition 4.3.3. Let $\alpha \in [0, 1]$, and let Υ be a ground set. Two partitions $\{X, \overline{X}\}$ and $\{Y, \overline{Y}\}$ of Υ are α -orthogonal if all four of the cross intersections ($X \cap Y$, $X \cap \overline{Y}$, $\overline{X} \cap Y$, and $\overline{X} \cap \overline{Y}$) have size at least $\alpha|\Upsilon|/4$. A family of partitions \mathcal{X} is α -orthogonal if each pair of distinct partitions in \mathcal{X} is α -orthogonal.

While we construct large α -orthogonal families for an arbitrary ground set Υ , we apply our construction in the case $\Upsilon = \{0, 1\}^n$. Note that if $\alpha < 1$ and X and Y are chosen independently and uniformly at random from all subsets of a large ground set, then $\{X, \overline{X}\}$ and $\{Y, \overline{Y}\}$ are α -orthogonal with high probability. This suggests a natural way of constructing large α -orthogonal families. We use Chernoff's inequality.

Theorem (Chernoff's Inequality (See [82, Theorem 4.2].)). *Let Z_1, Z_2, \dots, Z_t be mutually independent random indicator variables where $Z_l = 1$ with probability p_l and $Z_l = 0$ with probability $1 - p_l$, let $Z = \sum_{l=1}^t Z_l$, and let $\mu = \mathbf{E}[Z] = \sum_{l=1}^t p_l$. If $0 \leq \delta \leq 1$, then $\Pr[Z < (1 - \delta)\mu] < e^{-\mu\delta^2/2}$.*

Lemma 4.3.4. *Let $\alpha \in (0, 1)$, and Υ be a ground set of size t . There exists a family of pairwise α -orthogonal partitions of Υ of size at least*

$$\left\lfloor \frac{\sqrt{2}}{2} \left(e^{\frac{(1-\alpha)^2}{16}} \right)^t \right\rfloor.$$

Proof. Let $r = \left\lfloor \frac{\sqrt{2}}{2} e^{\frac{(1-\alpha)^2}{16}} t \right\rfloor$. For each $1 \leq j \leq r$, choose a subset $X_j \subseteq \Upsilon$ uniformly and independently at random. We claim that with positive probability, $\{X_i, \overline{X}_i\}$ and $\{X_j, \overline{X}_j\}$ are α -orthogonal when $i \neq j$. In particular, this implies that with positive probability, the partitions are all distinct and that $\{\{X_j, \overline{X}_j\} : 1 \leq j \leq r\}$ is an α -orthogonal family of size r , which implies that some such family exists. Let $\mathcal{X} = \{\{X_j, \overline{X}_j\} : 1 \leq j \leq r\}$.

For each pair $\{i, j\}$ with $1 \leq i < j \leq r$, let A_{ij} be the event that one of the four cross intersections between $\{X_i, \overline{X}_i\}$ and $\{X_j, \overline{X}_j\}$ has size less than $\alpha \frac{t}{4}$, and let $A = \bigcup_{ij} A_{ij}$, so that A is the event that \mathcal{X} is not an α -orthogonal family. We show that $\Pr[A] < 1$.

Of course $\Pr[A] \leq \sum_{ij} \Pr[A_{ij}]$. Similarly, we have that $\Pr[A_{ij}] \leq 4p$, where p is the probability that $|X \cap Y| < \alpha \frac{t}{4}$ where $X \subseteq \Upsilon$ and $Y \subseteq \Upsilon$ are chosen uniformly and independently at random. For each $x \in \Upsilon$, let Z_x be the random indicator variable for the event that $x \in X \cap Y$, and let $Z = \sum_x Z_x$, so that $Z = |X \cap Y|$. Note that the Z_x are mutually independent random indicator variables and $Z_x = 1$ with probability $1/4$. By Chernoff's inequality,

$$p = \Pr \left[Z < \alpha \frac{t}{4} \right] < e^{-(t/4)(1-\alpha)^2/2} = e^{-\frac{(1-\alpha)^2}{8}t}.$$

It follows that

$$\Pr[A] \leq \sum_{ij} \Pr[A_{ij}] \leq \sum_{ij} 4p = 4 \binom{r}{2} p < 2r^2 e^{-\frac{(1-\alpha)^2}{8}t} \leq 1$$

and hence $\Pr[A] < 1$ as required. \square

It is possible to construct larger α -orthogonal families using more sophisticated probabilistic tools, such as the Lovász Local Lemma, but this does not give a substantial improvement to our bounds on $f(n)$.

Lemma 4.3.5. *Let $\varepsilon > 0$ and let $k = \lg(\varepsilon^{-2} \ln 4)$. If $T_1, \dots, T_r \in \mathcal{T}_n$ and $n \geq \lg(r) + k$, then there are binary subtrees $S_j \in \mathcal{B}(T_j)$ such that*

$$\left| \bigcup_j L(S_j) \right| \leq \left(\frac{3}{4} + \varepsilon \right) 2^n.$$

Proof. Let $\alpha = 1 - 4\varepsilon$, so $\alpha < 1$. By Lemma 4.3.4, there is a family \mathcal{X} of pairwise α -orthogonal partitions of $\{0, 1\}^n$ of size

$$\begin{aligned}
\left\lfloor 2^{-1/2} e^{\varepsilon^2 2^n} \right\rfloor &\geq \left\lfloor 2^{-1/2} e^{\varepsilon^2 (r\varepsilon^{-2} \ln 4)} \right\rfloor \\
&= \left\lfloor 2^{-1/2} e^{r \ln 4} \right\rfloor \\
&= \left\lfloor 2^{2r-1/2} \right\rfloor \\
&> 2^r
\end{aligned}$$

Fix an arbitrary linear ordering on $\{0, 1\}^n$. For each $\{X, \bar{X}\} \in \mathcal{X}$, we apply Corollary 4.3.2 to each of the trees T_1, \dots, T_r . Let $D_{\{X, \bar{X}\}}$ be the subset of $\{T_1, \dots, T_r\}$ consisting of those trees T for which Corollary 4.3.2 produces a binary subtree $S \in \mathcal{B}(T)$ where $L(S) \subseteq \min\{X, \bar{X}\}$, where the minimization is with respect to the chosen ordering on $\{0, 1\}^n$.

Because $|\mathcal{X}| > 2^r$, there exist distinct partitions $\{X, \bar{X}\}$ and $\{Y, \bar{Y}\}$ in \mathcal{X} with $D_{\{X, \bar{X}\}} = D_{\{Y, \bar{Y}\}}$. Let $D = D_{\{X, \bar{X}\}}$. For each T_j , we select $S_j \in \mathcal{B}(T_j)$ as follows. If $T_j \in D$, then we may choose $S_j \in \mathcal{B}(T_j)$ such that $L(S_j) \subseteq \min\{X, \bar{X}\}$. Alternately, if $T_j \notin D$, then we may choose $S_j \in \mathcal{B}(T_j)$ such that $L(S_j) \subseteq \max\{Y, \bar{Y}\}$.

Note that none of the S_j contains a path-label in $Z = \max\{X, \bar{X}\} \cap \min\{Y, \bar{Y}\}$. Moreover, because \mathcal{X} is α -orthogonal, we have that $|Z| \geq (\alpha/4)2^n$. It follows that

$$\left| \bigcup_j L(S_j) \right| \leq 2^n - (\alpha/4)2^n \leq \left(\frac{3}{4} + \varepsilon \right) 2^n.$$

□

We remark that the hypothesis $n \geq \lg(r) + k$ cannot be relaxed beyond reducing k . Indeed, suppose that $r = 2^n$ and index the ternary trees by vectors in $\{0, 1\}^n$. If each T_x is edge-labeled so that $L(T_x) = \{x\}$, then $L(S) = \{x\}$ for each $S \in \mathcal{B}(T_x)$. Consequently, regardless of which subtrees are chosen, $\bigcup_x L(S_x) = \{0, 1\}^n$.

Our main result (Theorem 4.3.7 below) asserts that for sufficiently small constants $c > 0$ the function f is $O(2^{n-c\sqrt{n}})$. We now briefly outline the proof of this result. Using induction, assume the result for some $m < n$ (which we will now need to pick carefully, given n), and given T , pick a complete binary subtree S' of depth m such that $|L(S')|$ is bounded by $\gamma 2^{m-c\sqrt{m}}$ for an appropriate choice of the constant γ .

Now we have two kinds of path-labels in S' : those that occur often (in the proof, at least $2^{c\sqrt{m}}$ many times), and those that do not. If a path-label x appears often, it doesn't matter how we choose to extend S' at leaves σ with path-label x , because the total number of path-labels for all leaves of S extending any such σ will be limited. And if a label x does not appear often, then we can apply Lemma 4.3.5 to obtain trees S_σ extending all the leaves of S' which are labeled by x , with a bounded total number of labels.

Our next lemma is technical and determines how we choose the depth m of subtree on which we apply induction. Because we apply Lemma 4.3.5 to a collection of 2^m trees of depth $n - m$, we need $n - m$ to be large. On the other hand, we will replace \sqrt{m} with \sqrt{n} in some of our bounds, so we want \sqrt{m} to be close to \sqrt{n} .

Lemma 4.3.6. *Let $c > 0$ and $k > 0$. If n is a sufficiently large integer, then there exists an integer m with $1 \leq m < n$ such that $n - m \geq c\sqrt{m} + k$ and $\sqrt{n} - \sqrt{m} \leq c$.*

Proof. For positive $x \in \mathbb{R}$, let $y(x) = x - c\sqrt{x} - (k + 1)$. We have $y(x) \rightarrow \infty$, so for large enough real x , we can let $h(x) = \sqrt{x} - \sqrt{y(x)}$. Algebraic manipulation yields $h(x) = \frac{n(x)}{d(x)}$, where $n(x) = c + (k + 1)/\sqrt{x}$ and $d(x) = (1 + \sqrt{1 - c/\sqrt{x} + (k + 1)/x})$, for all sufficiently large x . Hence $h(x) \rightarrow c/2$ as x grows. Therefore $h(x) \leq c$ when x is sufficiently large. Let n be large enough so that $h(n) \leq c$ and $y(n) > 0$, and let $m = \lceil y(n) \rceil$. Note that $m < n$ since $y(n) < n - 1$. Because $m - 1 < y(n) \leq m$, we have that

$$n - m = n - (m - 1) - 1 \geq n - y(n) - 1 = c\sqrt{n} + k \geq c\sqrt{m} + k.$$

Similarly, we have $\sqrt{n} - \sqrt{m} \leq \sqrt{n} - \sqrt{y(n)} = h(n) \leq c$. Finally, note that because $y(n) > 0$, we have that $1 \leq m$. \square

Theorem 4.3.7. *If $0 \leq c < \sqrt{\lg(4/3)} \approx 0.644$, then there is a constant γ such that $f(n) \leq \gamma 2^{n - c\sqrt{n}}$.*

Proof. Because $2^{c^2} < 4/3$, we may choose $\delta \in (3/4, 1/2^{c^2})$. Let $\varepsilon = \delta - 3/4$, let $k = \lg(\varepsilon^{-2} \ln 4)$ as in Lemma 4.3.5, and let n_0 be large enough so that for all $n \geq n_0$ there is some m as in Lemma 4.3.6. Note that because $\delta 2^{c^2} < 1$, we may choose γ to be large enough so that $(1 + \gamma\delta)2^{c^2} \leq \gamma$ holds and $f(n) \leq \gamma 2^{n - c\sqrt{n}}$ holds for all $n < n_0$. We prove that the bound holds for all n by induction.

Let $n \geq n_0$, apply Lemma 4.3.6 to obtain m , and consider $T \in \mathcal{T}_n$ with root r . Let T' be the complete ternary subtree of T rooted at r with depth m . By induction, there exists a complete $S' \in \mathcal{B}(T')$ with $|L(S')| \leq \gamma 2^{m - c\sqrt{m}}$.

For each $x \in \{0, 1\}^m$, let A_x be the set of leaves of S' with path-label x . We say that $x \in \{0, 1\}^m$ is *frequent* if $|A_x| \geq 2^{c\sqrt{m}}$, and we say that x is *infrequent* otherwise. Let α be the number of frequent labels, and let β be the number of infrequent labels.

For each leaf σ of T' , let T_σ be the complete ternary subtree of T rooted at σ of depth $n - m$. For each leaf σ in S' , we extend S' at σ by selecting some $S_\sigma \in \mathcal{B}(T_\sigma)$. The choice for S_σ depends on whether the path-label of σ in S' is frequent or not.

If x is frequent, then for each $\sigma \in A_x$, we choose $S_\sigma \in \mathcal{B}(T_\sigma)$ arbitrarily. Otherwise, suppose that x is infrequent, and let $\sigma_1, \dots, \sigma_r$ be the leaves in S' with path-label x . Because x is infrequent, we have $r \leq 2^{c\sqrt{m}}$. Moreover, each T_σ has depth $n - m$ and $n - m \geq c\sqrt{m} + k \geq \lg(r) + k$. Therefore Lemma 4.3.5 implies that there exist $S_\sigma \in \mathcal{B}(T_\sigma)$ such that

$|\bigcup_{\sigma \in A_x} L(S_\sigma)| \leq \delta 2^{n-m}$. Gluing together all the trees S_σ yields $S \in \mathcal{B}(T)$. We bound $|L(S)|$ as follows.

First, we bound the number of path-labels in $L(S)$ that extend frequent path-labels. Note that by the definition of ‘‘frequent,’’ $\alpha 2^{c\sqrt{m}} \leq 2^m$. If $x \in \{0, 1\}^m$, then the total number of path-labels in $L(S)$ which extend x is at most 2^{n-m} . Hence the total number of path-labels in $L(S)$ which extend a frequent path-label is at most $\alpha 2^{n-m} \leq 2^{n-c\sqrt{m}}$.

Next, we bound the number of path-labels in $L(S)$ that extend infrequent path-labels. If x is not frequent, then the number of path-labels in $L(S)$ that extend x is at most $\delta 2^{n-m}$. Note that $\beta \leq |L(S')| \leq \gamma 2^{m-c\sqrt{m}}$. Hence the number of path-labels in $L(S)$ that extend an infrequent path-label is at most $\beta \delta 2^{n-m} \leq \gamma 2^{m-c\sqrt{m}} \delta 2^{n-m} = \gamma \delta 2^{n-c\sqrt{m}}$.

Adding these two bounds, we have that

$$\begin{aligned} |L(S)| &\leq (1 + \gamma \delta) 2^{n-c\sqrt{m}} \\ &= (1 + \gamma \delta) 2^{c(\sqrt{n}-\sqrt{m})} 2^{n-c\sqrt{n}} \\ &\leq (1 + \gamma \delta) 2^{c^2} 2^{n-c\sqrt{n}} \\ &\leq \gamma 2^{n-c\sqrt{n}} \end{aligned}$$

as required. □

4.4 A lower bound on $f(n)$

Our strategy for bounding $f(n)$ from below is to construct edge-labeled ternary trees in which each path-label occurs along a limited number of paths, and then extend these trees slightly.

Lemma 4.4.1. *Define a sequence $\{a_m\}$ of integers via $a_0 = 1$ and $a_m = \lceil 3a_{m-1}/2 \rceil$ for $m \geq 1$. For each m , there exists $T_m \in \mathcal{T}_m$ such that for each $x \in \{0, 1\}^m$, the set A_x of all leaves in T_m with path-label x satisfies $|A_x| \leq a_m$.*

Proof. By induction on m . If $m = 0$, the statement holds trivially. For $m \geq 1$, the inductive hypothesis implies that there is $T_{m-1} \in \mathcal{T}_{m-1}$ in which each path label occurs at most a_{m-1} times. We extend T_{m-1} to a complete ternary tree of depth m as follows. Consider a path-label $x \in \{0, 1\}^{m-1}$. At each vertex u in A_x , add three children v_1, v_2, v_3 adjacent to u . Of the $3|A_x|$ new edges, arbitrarily label $\lceil 3|A_x|/2 \rceil$ with label 0 and label the others with label 1. Repeating for each $x \in \{0, 1\}^{m-1}$ yields T_m . □

It is straightforward to argue by induction that $(3/2)^m \leq a_m \leq 2(3/2)^m - 1$. Solving the recurrence exactly has received some study. Odlyzko and Wilf showed that $a_m = \lfloor K(3/2)^m \rfloor$ where $K \approx 1.6222$ [86]; see also [45]. The sequence appears in the On-Line Encyclopedia of

Integer Sequences with sequence identifier A061419. Our application requires only the easy upper bound $a_m \leq 2(3/2)^m$. By extending the trees provided in Lemma 4.4.1, we obtain a lower bound on $f(n)$.

Lemma 4.4.2. *If $m \geq 0$ and $s = \lceil \lg 2(3/2)^m \rceil$, then $f(m + s) \geq 2^m$.*

Proof. Obtain T_m as in Lemma 4.4.1, and let $s = \lceil \lg 2(3/2)^m \rceil$. We obtain a tree $T \in \mathcal{T}_{m+s}$ by extending T_m as follows. Fix some $x \in \{0, 1\}^m$, and let A_x be the set of all leaves in T_m with path-label x . Because $|A_x| \leq 2(3/2)^m$, we may choose distinct labels $\theta(\sigma) \in \{0, 1\}^s$ for each $\sigma \in A_x$. Extend T_m at σ by attaching the tree $T_\sigma \in \mathcal{T}_s$ with $L(T_\sigma) = \{\theta(\sigma)\}$. Following the same extension procedure for each label in $\{0, 1\}^m$ yields T .

Consider $S \in \mathcal{B}(T)$ and let $\sigma_1, \dots, \sigma_r$ be the vertices of S at depth m . For each σ_j , let τ_j be a leaf in S that is a descendant of σ_j . Because no two distinct leaves τ_i, τ_j share a common path-label, we have that $|L(S)| \geq r = 2^m$, as required. \square

Lemma 4.4.2 yields a lower bound on $f(n)$ only when n is of a special form; however, we claim that for each n , either n or $n - 1$ is of a form to which Lemma 4.4.2 applies. Let $b_m = m + \lceil \lg 2(3/2)^m \rceil$, and note that for $m \geq 1$, we have that

$$b_m - b_{m-1} = 1 + \lceil \lg 2(3/2)^m \rceil - \lceil \lg 2(3/2)^{m-1} \rceil < 2 + \lg 3/2 < 3.$$

Because $b_m - b_{m-1}$ is an integer, we have that $b_m - b_{m-1} \leq 2$. We obtain the following general lower bound.

Theorem 4.4.3. *For each n , we have $f(n) \geq 2^{\frac{n-3}{\lg 3}} \geq (0.269) \cdot (1.548)^n$.*

Proof. Let m be an integer such that either n or $n - 1$ is equal to $m + \lceil \lg 2(3/2)^m \rceil$. Lemma 4.4.2 implies that $f(n) \geq f(m + \lceil \lg 2(3/2)^m \rceil) \geq 2^m$. Note that

$$n - 1 \leq m + \lceil \lg 2(3/2)^m \rceil \leq m + (\lg 2(3/2)^m) + 1 = (\lg 3)m + 2$$

and therefore $m \geq (n - 3)/\lg 3$. \square

Corollary 4.4.4. *We have that $\lim_{n \rightarrow \infty} (f(n))^{1/n} \geq 2^{\frac{1}{\lg 3}} \geq 1.548$.*

A generalization of f

Our techniques partially extend to a natural generalization of f , where the branching factor the host tree is an integer q and we seek subtrees where each non-leaf has p children. A tree is q -ary if every non-leaf has q children. Let $\mathcal{T}_n^{(q,t)}$ be the set of all complete $\{0, \dots, t - 1\}$ -edge-labeled q -ary trees of depth n . For $2 \leq t \leq p < q$ and $T \in \mathcal{T}_n^{(q,t)}$, consider the problem of finding a p -ary subtree of T with few path-labels; we have focused on the case $(t, p, q) =$

(2, 2, 3). When $T \in \mathcal{T}_n^{(q,t)}$ and S is a complete p -ary subtree of depth n , we write $S \sqsubset_p T$. Let $f(T; p) = \min\{|L(S)| : S \sqsubset_p T\}$, and define $f(n; t, p, q) = \max\{f(T; p) : T \in \mathcal{T}_n^{(q,t)}\}$. While the behavior of $f(n; t, p, q)/t^n$ remains largely unexplored, we are able to prove the following.

Theorem 4.4.5.

$$\lim_{n \rightarrow \infty} \frac{f(n; t, p, q)}{t^n} = \begin{cases} 0 & \text{if } p < \frac{1}{2}q + 1 \\ 1 & \text{if } p \geq \frac{t-1}{t}q + 1 \end{cases}$$

Proof (sketch). Suppose that $p < q/2 + 1$, or equivalently, $2p - 1 \leq q$. In this case, a modified version of Proposition 4.3.1 holds. If T is a complete q -ary tree of depth n whose leaves are all colored red or blue, then there is a p -ary subtree S with $S \sqsubset_p T$ such that the leaves of S are monochromatic. We then extend Lemma 4.3.5 in the natural way, so that if $T_1, \dots, T_r \in \mathcal{T}_n^{(q,t)}$ and n is sufficiently large in terms of r , then there are p -ary subtrees S_j of T_j such that $\bigcup_j L(S_j)$ contains at most a constant fraction of the total space of path-labels $\{0, \dots, t-1\}^n$. We find a p -ary subtree with few path-labels by iteratively applying this lemma at the leaves of the current p -ary subtree. Each application results in a constant factor reduction of the fraction of used path-labels.

If $p \geq \frac{t-1}{t}q + 1$, then in fact $f(n; t, p, q) = t^n$. Let $k = \lfloor q/t \rfloor$, and let $T \in \mathcal{T}_n^{(q,t)}$ be a tree in which each of the t edge-labels is used on at least k of the q descendant edges of each non-leaf vertex. Because p is an integer, we have $p \geq \lceil \frac{t-1}{t}q + 1 \rceil = \lceil \frac{t-1}{t}q \rceil + 1$, and therefore $p > \lceil \frac{t-1}{t}q \rceil = \lceil q - q/t \rceil = q - k$. Let u be a non-leaf in T and let $j \in \{0, \dots, t-1\}$. Since $p + k > q$, every choice of p children at u includes at least one child v such that uv has edge-label j . It follows that $L(S) = \{0, \dots, t-1\}^n$ for each S with $S \sqsubset_p T$. \square

Theorem 4.4.5 determines the limiting behavior of $f(n; t, p, q)/t^n$ as $n \rightarrow \infty$ when $t = 2$. The first unknown case is $(t, p, q) = (3, 3, 4)$.

4.5 An application to computability theory

In this section, we present a result in computability theory. Computability theory traditionally borrows much of its notation from set theory. When A and B are sets, we use B^A to denote the set of functions from A to B . Let $\omega = \{0, 1, \dots\}$. The set ω^ω is called *Baire space*.

A fundamental goal of computability theory is to understand which functions in Baire space are computable. Informally, a function $f: \omega \rightarrow \omega$ is *computable* if there is a procedure which, when given a number $n \in \omega$ as input, outputs $f(n)$ after a finite number of steps. A *partial function from A to B* is a function from a subset A' of A to B ; we write $\phi: A \rightarrow B$ when ϕ is a partial function from A to B . When $n \in A'$, we say that $\phi(n)$ is *defined* or

converges, and we write $\phi(n) \downarrow$. When $n \in A$ but $n \notin A'$, we say that $\phi(n)$ is *undefined* or *diverges*, and we write $\phi(n) \uparrow$. When $A' = A$, we say that ϕ is *total*. A partial function $\phi: \omega \rightarrow \omega$ is *partial computable* if there is a procedure which, when given a number $n \in \omega$ as input, outputs $\phi(n)$ after a finite number of steps if $\phi(n) \downarrow$ and does not terminate if $\phi(n) \uparrow$.

There are several different mathematical models of procedures, such as Turing Machines and functions in lambda calculus. The definitions of “computable” and “partial computable” are robust in that they do not depend on the model chosen. We use the Turing Machine model. Here, we describe the Turing Machine model informally; for a formal definition, see [103, 102]. A *Turing Machine* M consists of a finite set of states, a transition function, and a one-way infinite tape for data storage. The tape is divided into cells c_1, c_2, \dots and each cell is blank or contains a 0 or a 1. A tape head which is positioned over one of the cells on the tape, which is called the current cell. One of the states is designated as the “start state”, and another is designated the “halt state”. Based upon the current state and the contents of the current cell, the transition function specifies the next state, potentially overwrites the contents of current cell, and optionally moves the tape head left or right.

To run M on an input $n \in \omega$, the number n is first written on the tape in binary, starting at the left end (other cells are initially blank), and the current state of M is set to the start state. If M enters the halt state after a finite number of steps, then we write $M(n) \downarrow$, we say that $M(n)$ *converges* or *halts*, and we use $M(n)$ to denote the contents of the tape. If M never enters the halt state, we write $M(n) \uparrow$ and we say that $M(n)$ *diverges*.

One crucial property of Turing Machines is that each admits a finite description. Consequently, the set of all Turing Machines is countable. Let M_1, M_2, \dots be a reasonable¹ list of Turing Machines. Each Turing Machine M_e computes a computable partial function φ_e , and conversely every computable partial function is computed by some Turing Machine. Hence, the set of partial computable functions is a countable subset of ω^ω .

One common variant of the Turing Machine is the oracle Turing Machine. An *oracle Turing Machine* has all the components of a Turing Machine, plus an additional one-way infinite tape called the *oracle tape* and an *oracle query state*. The oracle tape has its own head which operates independently of the data tape head. An oracle Turing Machine M is run with respect to an oracle $f \in \omega^\omega$. When the oracle Turing Machine enters the oracle query state, the contents of the oracle tape are interpreted as the binary representation of a number n and are replaced with the binary representation of $f(n)$. When M is run with respect to f on an input n , we extend the notions of divergence, convergence, and output using the notation $M^f(n)$ in place of $M(n)$. The *Turing functional* Φ corresponding to oracle Turing Machine M is the function from ω^ω to the set of partial functions from ω to ω where $\Phi(f)$ is the partial function computed by M^f . Hence, the notation $\Phi(f)(n)$ is the same as

¹Here, “reasonable” means that there is a computable way of passing from the combinatorial description of M_e to e and conversely. It is helpful to think of e as being the “source code” for the Turing Machine M_e .

$M^f(n)$ when Φ is the Turing functional corresponding to M .

Our application requires a generalization to partial edge-labelings of the infinite ternary tree. A *ternary sequence* is a finite sequence of 0's, 1's, and 2's. The *full ternary tree* is the collection of all ternary sequences, ordered by sequence extension. This partial ordering can also be viewed as a (connected, acyclic) graph where two sequences are joined by an edge if one is an immediate extension of the other, that is, the one extends the other by one digit. The empty sequence is the root of the tree. The set of vertices at depth k is $\{0, 1, 2\}^k$.

We consider partial edge-labelings of the full ternary tree. Recall that if $\sigma\tau$ is an edge in the full ternary tree and τ extends σ by one character, then the level of $\sigma\tau$ is the depth of τ . Hence, edges incident to the root are at level 1. Let U be an infinite set of positive integers, which will indicate a set of levels of the full ternary tree; let u_1, u_2, u_3, \dots be an increasing enumeration of the elements of U . A *U -edge-labeling* of the full ternary tree is an assignment of a label in $\{0, 1\}$ to each edge at every level in U . As before, reading the labels along edges in the path from the root to a vertex σ gives a path-label, and reading the labels along the edges of an infinite path starting at the root gives a path-label in $\{0, 1\}^\omega$, where $\omega = \{0, 1, 2, \dots\}$. A binary subtree S of the full ternary tree is *complete* if it is nonempty and has no leaves. (Note that we are considering subtrees in the graph-theoretic sense. In particular, such an S is “2-bushy” in the sense that every node at depth n has two children at depth $n + 1$.) For such a subtree S , let $L(S)$ be the set of path-labels of paths through S , as before. Also as before, our object is to find such an S with $L(S)$ “small.” However, it is easily seen that it is not possible in general to choose such an S with $L(S)$ countable. Instead, we ensure that $L(S)$ has measure 0 in the usual fair-coin measure on 2^ω . This amounts to choosing S so that $\lim_n |L(S_{u_n})|/2^n = 0$, where S_k is the set of nodes and vertices of S with depth at most k .

A set $U \subseteq \mathbb{N}$ is *computable* if there is an algorithm which, given $n \in \mathbb{N}$, decides if $n \in U$. A *U -edge-labeling* is *computable* if there is an algorithm which, given an edge $\sigma\tau$ in the full ternary tree, outputs the label on $\sigma\tau$. A full binary subtree S of the full ternary tree is *computable* if there is an algorithm that, given a ternary sequence σ , decides if σ is a vertex in S . For our application to computability theory, we also need the proof to be effective in the sense that we can choose S to be computable if T and U are computable.

Theorem 4.5.1. *Let U be an infinite set of positive integers, and let T be a U -edge-labeling of the full ternary tree. Then there is a complete binary subtree S of T such that $L(S)$ has measure 0 as a subset of $\{0, 1\}^\omega$. Furthermore, if U and T are computable, we may require S to be computable.*

Proof. We prove the computable version of the result, and of course the other version follows by the same argument, omitting all mention of computability. Let r be the root of T . We obtain S by computing a sequence S_0, S_1, \dots of finite, complete binary subtrees rooted at r . Each tree S_j is a proper subtree of S_{j+1} and $L(S_j)$ has size at most $(3/4)^j \cdot 2^n$, where n is

the length of the path-labels in S_j . We set $S = \bigcup_j S_j$. Note that S is a full binary subtree of T and $L(S)$ has measure 0. Moreover, S is computable; to see if σ is a vertex in S , simply compute S_j for large enough j so that path-labels in S_j have length at least as long as σ and test if σ is in S_j .

Let S_0 be the binary subtree of depth 0 rooted at r . Given S_j , we show how to compute S_{j+1} . We obtain S_{j+1} by gluing trees of the same depth to the leaves of S_j . These trees are obtained from a modified version of Lemma 4.3.5. This modified version is explained next.

The argument of Lemma 4.3.5 easily extends to the partial edge-labeling case when the n of the lemma is replaced by the length of the path-labels in the given partially edge-labeled trees. In fact, the argument becomes easier because we are no longer trying to establish a delicate upper bound on the number of labels. The key to applying the pigeon-hole principle in the proof of Lemma 4.3.5 is that $|\mathcal{X}| > 2^n$, where \mathcal{X} is a family of pairwise α -orthogonal partitions of $\{0, 1\}^n$. If we now set $\alpha = 1$ and require $n > 2^r$, we now achieve this easily by taking $\mathcal{X} = \{\{X_i, \overline{X_i}\} : 1 \leq i \leq n\}$, where X_i is the set of binary words of length n with a 1 in the i th bit, so that \mathcal{X} is a 1-orthogonal family of partitions of $\{0, 1\}^n$. (This avoids the use of Chernoff's Inequality to construct a large α -orthogonal family, at the cost of making n much larger than in the original version of Lemma 4.3.5.) Since now $\alpha = 1$, the first inequality in the final line of the proof of Lemma 4.3.5 yields $|\bigcup_j L(S_j)| \leq (3/4)2^n$.

Let m be the length of the path-labels in S_j , and let A be the set of all leaves in S_j . Let $n = 2^{|A|} + 1$. For each σ in A , let T_σ be the complete ternary subtree rooted at σ whose leaves have depth u_{m+n} in T . Note that by construction, the path-labels in T_σ all have length n .

Therefore, by the modified version of Lemma 4.3.5 discussed above, for each $\sigma \in A$, there exists a complete binary subtree S_σ of T_σ of full depth such that $|\bigcup_{\sigma \in A} L(S_\sigma)| \leq (3/4)2^n$. Because A is finite and there are only a finite number of candidates for each S_σ , we may compute such a collection of subtrees using brute force. Let S_{j+1} be the binary subtree obtained by gluing S_σ at each leaf σ in S_j . Note that S_{j+1} has depth u_{m+n} and the path-labels in S_{j+1} have length $m + n$.

For each $x \in L(S_j)$, there are at most $(3/4)2^n$ path-labels in $L(S_{j+1})$ that extend x . It follows that $|L(S_{j+1})| \leq (3/4)2^n \cdot |L(S_j)| = (3/4)2^n (3/4)^j 2^m = (3/4)^{j+1} 2^{m+n}$, as required. \square

Equip ω with the discrete topology. The product space ω^ω , also known as *Baire space*, is a universal Polish space (a separable, completely metrizable space). Medvedev [76] considered subsets of Baire space to be “mass problems,” where the idea is that the elements of a set \mathcal{A} are the solution of the “problem.” For example, if \mathcal{A} is a singleton $\{f\}$, then the problem \mathcal{A} is the problem of computing f . For another example, if \mathcal{A} consists of all functions whose range is some nonempty set X , then \mathcal{A} is the problem of enumerating the elements of X .

When is one mass problem at least as difficult as another? Medvedev [76] introduced a reducibility on mass problems which is now often called Medvedev reducibility. Namely, \mathcal{B} is

Medvedev reducible to \mathcal{A} , denoted $\mathcal{B} \leq_M \mathcal{A}$, if there is a uniform way to compute a solution for \mathcal{B} given any solution for \mathcal{A} . Formally, this means there is a Turing functional Φ such that $\Phi(f) \in \mathcal{B}$ for all $f \in \mathcal{A}$. In other words, there is a fixed oracle Turing machine which, given any function $f \in \mathcal{A}$ as oracle, computes a function $g \in \mathcal{B}$, which must of course be a total function. Note that Medvedev reducibility extends Turing reducibility in the sense that for $f, g \in \omega^\omega$, g is Turing reducible to f if and only if $\{g\}$ is Medvedev reducible to $\{f\}$.

The relation \leq_M is a pre-partial ordering on Baire space. We call two mass problems *Medvedev equivalent* if each is Medvedev reducible to the other, and Medvedev equivalence is an equivalence relation. The equivalence classes are called *Medvedev degrees*; the collection of degrees is turned into a *degree structure* by adding the induced partial ordering. In fact, this degree structure is a distributive lattice, where the least upper bound is induced by pairwise effective join

$$\mathcal{A} \times \mathcal{B} = \{f \oplus g : f \in \mathcal{A} \ \& \ g \in \mathcal{B}\}^2$$

and greatest lower bound given by effective disjoint union

$$\mathcal{A} \sqcup \mathcal{B} = \{0f : f \in \mathcal{A}\} \cup \{1g : g \in \mathcal{B}\}.$$

The Medvedev degrees have a least element $\mathbf{0}$ which consists of all mass problems that contain a computable function. The greatest element, the degree of the empty set, is usually ignored. The Medvedev degree of a mass problem \mathcal{A} is denoted by $\text{deg}_M(\mathcal{A})$ or sometimes simply \mathbf{a} .

Medvedev reducibility is also known as *strong reducibility*. This is because Muchnik [83] later introduced a weaker version of Medvedev reducibility, the difference being that uniform computation is no longer required: a mass problem \mathcal{B} is Muchnik (or *weakly*) reducible to a mass problem \mathcal{A} if, for each $f \in \mathcal{A}$, there is a Turing functional Φ such that $\Phi(f) \in \mathcal{B}$. Here, the order of quantifiers allows for a different functional Φ for each $f \in \mathcal{A}$, and so the behavior of the reduction is no longer uniform over the functions in \mathcal{A} . The corresponding degree structure is isomorphic to the sublattice of the power set of the Turing degrees, consisting of all the sets of Turing degrees which are closed upwards, i.e. are unions of cones.

Now recall that Baire space ω^ω is actually a topological space with basis $\{O_\tau : \tau \in \omega^{<\omega}\}$, where $\omega^{<\omega} = \bigcup_{n \geq 0} \omega^n$ and the basic open set O_τ is given by $O_\tau = \{f \in \omega^\omega : f \text{ extends } \tau\}$. Because τ is a finite list of numbers, algorithms can output the basic open set O_τ implicitly by referring to τ .

The topological notions of open and closed sets can be refined using computability theory. We say that an open set $\mathcal{O} \subseteq \omega^\omega$ is *effectively open* (or Σ_1^0) if the collection of basic open subsets of \mathcal{O} is computably enumerable. The complement of an effectively open set is *effectively closed*, or Π_1^0 . We note that a set is effectively closed if and only if it is the set of

²Where $(f \oplus g)(2n) = f(n)$ and $(f \oplus g)(2n + 1) = g(n)$.

infinite paths through a computable subtree of $\omega^{<\omega}$.

Of particular interest are effectively closed sets which are *computably bounded*, in other words, are subsets of the subspace

$$\prod_{n \geq 0} \{0, 1, \dots, h(n) - 1\}$$

of ω^ω , for some computable function h . These closed sets are the sets of paths through trees which are finitely branching, and so are compact subsets of ω^ω ; the extra effectiveness condition implies that in some sense these sets are “effectively compact.” Computably bounded Π_1^0 classes have occurred as the set of solutions to problems in logic, combinatorics, and algebra and other areas; for a survey of this aspect of Π_1^0 classes, see [23], for example. If \mathcal{A} is a computably bounded Π_1^0 class, then the Medvedev degree of \mathcal{A} contains a Π_1^0 class \mathcal{B} such that \mathcal{B} is contained in *Cantor space* $\{0, 1\}^\omega$. Hence, we may study the computably bounded Π_1^0 classes via their Cantor space representatives. Much research was devoted to studying the Medvedev and Muchnik degrees of effectively closed subsets of Cantor space; see [101] for more details. Both substructures of the full degree structures are again distributive lattices, as the lattice operations, applied to effectively closed sets, yield effectively closed sets.

Let \mathcal{P}_M be the lattice of Medvedev degrees of effectively closed subsets of Cantor space. In the lattice \mathcal{P}_M there is a greatest element $\mathbf{1}$, which is the degree of the set of all consistent completions of Peano arithmetic, viewed as a mass problem via standard coding. Another Π_1^0 mass problem in $\mathbf{1}$ is the set DNR_2 of $\{0, 1\}$ -valued diagonally nonrecursive functions. A function $f \in \omega^\omega$ is called *diagonally nonrecursive* if for all e we have $f(e) \neq \varphi_e(e)$, where $\{\varphi_e\}_{e \in \omega}$ is an effective list of all partial computable functions. Here $f(e) \neq \varphi_e(e)$ means that either $\varphi_e(e)$ is undefined, or it is defined with a value unequal to $f(e)$.

The set DNR of all diagonally nonrecursive functions is effectively closed, but is not recursively bounded, indeed is not compact. However, we let, for every $k < \omega$, DNR_k be the set of all diagonally nonrecursive functions

$$f: \omega \rightarrow \{0, 1, 2, \dots, k - 1\},$$

i.e., k -valued diagonally nonrecursive functions. These sets are Π_1^0 and are recursively bounded, hence their Medvedev degrees lie in \mathcal{P}_M . As mentioned,

$$\text{deg}_M(\text{DNR}_2) = \mathbf{1}$$

is the greatest degree in \mathcal{P}_M . However, Jockusch [61, Theorem 6] showed that if we let $\mathbf{d}_k = \text{deg}_M(\text{DNR}_k)$, then

$$\mathbf{d}_2 > \mathbf{d}_3 > \mathbf{d}_4 > \dots$$

is a strictly decreasing sequence. This contrasts with Jockusch’s result [61, Theorem 5]) that the Muchnik (weak) degrees of all of the classes DNR_k coincide, i.e., the classes $\text{DNR}_2, \text{DNR}_3, \dots$ are all Muchnik equivalent.

If the classes DNR_k are “high”, or close to $\mathbf{1}$, then classes of positive measure should be considered “fat”, and so “low”, or close to $\mathbf{0}$. The measure we use is the product probability measure, using the fair coin measure on $\{0,1\}$. Simpson [101, Corollary 7.11] showed a non-join result: if $\mathbf{a} \in \mathcal{P}_M$ is the Medvedev degree of a Π_1^0 class of positive measure, $k \geq 2$, and $\mathbf{b} \in \mathcal{P}_M$ is not above \mathbf{d}_k , then the join $\mathbf{a} \vee \mathbf{b}$ is not above \mathbf{d}_k either. Thus Π_1^0 classes of positive measure are so weak, that they cannot help any other Π_1^0 class compute bounded diagonally nonrecursive functions. Simpson asked [101, Remark 7.12] if the reason for this is that all Π_1^0 classes of positive measure are Medvedev reducible to each DNR_k . Here we answer Simpson’s question in the negative:

Theorem 4.5.2. *There is a Π_1^0 class $P \subseteq 2^\omega$ of positive measure which is not Medvedev reducible to DNR_3 .*

In fact, we prove a stronger result. Measure and Π_1^0 classes are closely tied to notions of effective randomness. For more background see [101]; here we just mention that every Π_1^0 class of positive measure contains a tail of every Martin-Löf random set³ (Kučera [74]), and that since the collection of Martin-Löf random sets is a Σ_2^0 set (an effective F_σ set), there are nonempty Π_1^0 classes, necessarily of positive measure, which contain only Martin-Löf random sets. Thus in the Muchnik (weak) degrees, the degree of the set of Martin-Löf random sets is the same as the degree of any Π_1^0 class which contains only Martin-Löf random sets, is the greatest degree of Π_1^0 classes of positive measure. In the Medvedev degrees, the picture is not as tidy; the Medvedev degree \mathbf{r} of the set of Martin-Löf random sets is not in \mathcal{P}_M , and in fact there is no greatest Medvedev degree of Π_1^0 classes of positive measure (Terwijn [108]). However, trivially, if P is a Π_1^0 class which only contains Martin-Löf random sets, then the identity functional witnesses that $\mathbf{r} \leq \text{deg}_M(P)$, and so Theorem 4.5.2 follows from the following theorem, which answers a question raised by J. Miller.

Theorem 4.5.3. $\mathbf{r} \not\leq \mathbf{d}_3$.

Indeed, we prove a little more. Recall that a set $X \in 2^\omega$ is *Kurtz random* (or *weakly 1-random*) if it is not a member of any null Π_1^0 subset of 2^ω . This is a notion of randomness which is much weaker than Martin-Löf randomness. We prove the following, which implies Theorem 4.5.3:

³Recall that a sequence $X \in 2^\omega$ is Martin-Löf random if whenever $\langle U_n \rangle$ is an effective sequence of effectively open subsets of 2^ω such that the measure of each U_n is at most 2^{-n} , we have $X \notin \bigcap_n U_n$. Equivalently, the initial segments of X are incompressible, in the sense that there is a constant c such that for all n , $K(X \upharpoonright n) \geq n - c$; here K denotes prefix-free Kolmogorov complexity. See, for example, [31] for more details on effective randomness.

Theorem 4.5.4. *The set of Kurtz random sets is not Medvedev reducible to DNR_3 .*

We note that in contrast, recently Greenberg and J. Miller [52] showed that the set of reals which have effective Hausdorff dimension 1 is Medvedev reducible to each DNR_k .

Proof. Suppose for a contradiction that Φ is a Turing functional that witnesses that DNR_3 is Medvedev reducible to the Kurtz random sets. Hence, for each $f \in \text{DNR}_3$, we have that $\Phi(f)$ is the characteristic function of a Kurtz random set and is thus *total* (i.e. given an input, returns a value after a finite amount of time). Without loss of generality, we may assume that for all $f \in 3^\omega$, $\Phi(f)$ is total and $\{0, 1\}$ -valued. We do this by replacing Φ by a modified Turing functional which, with oracle $f \in 3^\omega$ and input n , simulates Φ until either $\Phi(f)(n)$ converges or it is discovered by a systematic search that $f(e) = \varphi_e(e)$ for some e , so $f \notin \text{DNR}_3$. One of these events must occur since $\Phi(f)$ is total for all $f \in \text{DNR}_3$. If the former occurs first, the modified functional outputs $\Phi(f)(n)$, and otherwise it outputs 0 (say).

By compactness, and buffering the use of Φ -computations, we can obtain an effective increasing sequence $u_1 < u_2 < u_3 < \dots$ such that for all $n \geq 1$, for all $X \in 3^\omega$, the X -use of computing $\Phi(X) \upharpoonright n$ is exactly u_n . Let $U = \{u_1, u_2, u_3, \dots\}$. Thus Φ yields a computable U -edge labeling of the full ternary tree $3^{<\omega}$: for $\sigma \in \{0, 1, 2\}^{u_n}$, we let $\Phi(\sigma) \in \{0, 1\}^n$ be the result of applying Φ to the oracle σ , and so we label the parent edge incident to σ in the full ternary tree with the last bit of $\Phi(\sigma)$. That is, the path-label of σ is exactly $\Phi(\sigma)$ for all $\sigma \in \bigcup_{n \geq 1} \{0, 1, 2\}^{u_n}$, and hence for all $f \in 3^\omega$, $\Phi(f)$ is the path-label of f .

By Theorem 4.5.1, there is a computable full binary subtree S of $3^{<\omega}$ such that $L(S)$ is null in Cantor space. We show that $L(S)$ is effectively closed by enumerating the basic open subsets of its complement $\{0, 1\}^\omega - L(S)$. For each $n \geq 1$, compute u_n and compute the set A of all vertices in S at depth u_n . Next, for each $\sigma \in A$, compute the path-label of σ . In the enumeration, output all basic open sets O_τ such that $\tau \in \{0, 1\}^n$ is not the path-label of some $\sigma \in A$. It follows that $L(S)$ is a null Π_1^0 class.

Let T be the subtree of $3^{<\omega}$ induced by the vertices τ such that DNR_3 contains an extension of τ . Note that if τ is in T , then because τ has three children in $3^{<\omega}$ and at most one violates the diagonally non-recursive condition, at least two of its children are in T . It follows that T and S contain a common infinite path $f \in \{0, 1, 2\}^\omega$. Because f is an infinite path in T , we have that f is in DNR_3 . Because f is an infinite path in S , we have that $\Phi(f) \in L(S)$, and so $\Phi(f)$ is a member of a null Π_1^0 class, which implies that $\Phi(f)$ is not Kurtz random. Hence $f \in \text{DNR}_3$ but $\Phi(f)$ is not Kurtz random, and so Φ fails to witness that the class of Kurtz random sets is Medvedev-reducible to DNR_3 , as required. \square

Chapter 5

Parity Edge-Colorings of Graphs

A *parity walk* in an edge-coloring of a graph is a walk along which each color is used an even number of times. Let $p(G)$ be the least number of colors in an edge-coloring of G having no parity path (a *parity edge-coloring*). Let $\widehat{p}(G)$ be the least number of colors in an edge-coloring of G having no open parity walk (a *strong parity edge-coloring*). Always $\widehat{p}(G) \geq p(G) \geq \chi'(G)$. We prove that $\widehat{p}(K_n) = 2^{\lceil \lg n \rceil} - 1$ for all n . The optimal strong parity edge-coloring of K_n is unique when n is a power of 2, and the optimal colorings are completely described for all n .

This chapter is based on joint work with D. P. Bunde, D. B. West, and H. Wu [19, 20].

5.1 Introduction

Our work began by studying which graphs embed in the hypercube Q_k , the graph with vertex set $\{0, 1\}^k$ in which vertices are adjacent when they differ in exactly one coordinate. Coloring each edge with the position of the bit in which its endpoints differ yields two necessary conditions for the coloring inherited by a subgraph G :

- 1) every cycle uses each color an even number of times,
- 2) every path uses some color an odd number of times.

Existence of a k -edge-coloring satisfying conditions (1) and (2) is also sufficient for a connected graph G to be a subgraph of Q_k . This characterization of subgraphs of Q_k appeared in 1972 (Havel and Morávek [57]). The problem was studied as early as 1953 (Shapiro [100]).

Let the *usage* of a color on a walk be the parity of the number of times it appears along the walk. A *parity walk* is a walk in which the usage of every color is even. Condition (1) above states that every cycle is a parity walk, and (2) states that no path is a parity walk.

In general, a *parity edge-coloring* is an edge-coloring with no parity path, and a *strong parity edge-coloring (spec)* is an edge-coloring with no open parity walk (that is, every parity walk is closed). Some graphs embed in no hypercube, but giving the edges distinct colors produces a spec for any graph. Hence the *parity edge-chromatic number* $p(G)$ and the *strong parity edge-chromatic number* $\widehat{p}(G)$, defined respectively to be the minimum numbers of colors in a parity edge-coloring of G and in a spec of G , are well defined.

5.2 Trees, hypercubes, paths, and cycles

In this section, we present elementary results about parity edge-colorings. Some of these were obtained previously by authors whose aim was to study related concepts; we include proofs for completeness and to emphasize the connections with parity edge-colorings.

Remark 5.2.1. For every graph G , $\widehat{p}(G) \geq p(G) \geq \chi'(G)$, and the parameters \widehat{p} and p are monotone under the subgraph relation.

Proof. We have $p(G) \geq \chi'(G)$ by considering paths of length 2, and $\widehat{p}(G) \geq p(G)$ since closed walks are not paths. For $H \subseteq G$, a pec or spec of G restricts to such an edge-coloring on H , since every parity walk in the restriction to H is a parity walk in the coloring on G . \square

When G is a forest, every pec is also a spec, so $p(G) = \widehat{p}(G)$. Edge-coloring the hypercube by coordinates shows that $p(Q_k) \leq \widehat{p}(Q_k) \leq k$. Hence $p(G) \leq k$ if $G \subseteq Q_k$. For trees, we prove the converse.

Given a k -edge-coloring f and a walk W , we use $\pi(W)$ to denote the parity vector of W , recording the usage of each color as 0 or 1. When walks W and W' are concatenated, the parity vector of the concatenation is the vector binary sum $\pi(W) + \pi(W')$. The *weight* of a vector is the number of nonzero positions.

Theorem 5.2.2 ([57]). *A tree T embeds in the k -dimensional hypercube Q_k if and only if $p(T) \leq k$.*

Proof. We have observed necessity. Conversely, let f be a parity k -edge-coloring of T (there may be unused colors if $p(T) < k$). Fix a root vertex r in T . Define $\phi: V(T) \rightarrow V(Q_k)$ by setting $\phi(v) = \pi(W)$, where W is the r, v -path in T .

When $uv \in E(T)$, the r, u -path and r, v -path in T differ in one edge, so $\phi(u)$ and $\phi(v)$ are adjacent in Q_k . It remains only to check that ϕ is injective. The parity vector for the u, v -path P in T is $\phi(u) + \phi(v)$, since summing the r, u -path and r, v -path cancels the portion from r to P . Since f is a parity edge-coloring, $\phi(P)$ is nonzero, and hence $\phi(u) \neq \phi(v)$. \square

When k is part of the input, recognizing subgraphs of Q_k is NP-complete [72], and this remains true when the input is restricted to trees [111]. Therefore, computing $p(G)$ or $\widehat{p}(G)$ is NP-hard even when G is a tree. Perhaps there is a polynomial-time algorithm for trees with bounded degree or bounded diameter.

The Havel–Movárek characterization [57] of subgraphs of Q_k follows easily from Theorem 5.2.2 (they also proved statements equivalent to Theorem 5.2.2 and Corollary 5.2.5.) Their proof is essentially the same as ours, but our organization is different in the language of pecs.

Corollary 5.2.3 ([57]). *A graph G is a subgraph of Q_k if and only if G has a parity k -edge-coloring in which every cycle is a parity walk.*

Proof. We have observed necessity. For sufficiency, choose a spanning tree T . Since $p(T) \leq p(G) \leq k$, Theorem 5.2.2 implies that $T \subseteq Q_k$. Map T into Q_k using ϕ as defined in the proof of Theorem 5.2.2. For each $xy \in E(G) - E(T)$, the cycle formed by adding xy to T is given to be a parity walk. Hence the x, y -path in T has parity vector with weight 1. This makes $\phi(x)$ and $\phi(y)$ adjacent in Q_k , as desired. \square

Mitas and Reuter [80] later gave a lengthier proof motivated by studying subdiagrams of the subset lattice. They also characterized the graphs occurring as induced subgraphs of Q_k as those having a k -edge-coloring satisfying properties (1) and (2) and (3), where property (3) essentially states that if the parity vector of a walk W has weight 1, then the endpoints of W are adjacent.

Spanning trees yield a general lower bound on $p(G)$, which holds with equality for paths, even cycles, and connected spanning subgraphs of Q_k .

Corollary 5.2.4. *If G is connected, then $p(G) \geq \lceil \lg n(G) \rceil$.*

Proof. If T is a spanning tree of G , then $p(G) \geq p(T)$. Since T embeds in the hypercube of dimension $p(T)$, we have $n(G) = n(T) \leq 2^{p(T)} \leq 2^{p(G)}$. \square

Corollary 5.2.5. *For all n , $p(P_n) = \widehat{p}(P_n) = \lceil \lg n \rceil$. For even n , $p(C_n) = \widehat{p}(C_n) = \lceil \lg n \rceil$.*

Proof. The lower bounds follow from Corollary 5.2.4. The upper bounds hold because Q_k contains cycles of all even lengths up to 2^k . \square

A result equivalent to $p(P_n) = \widehat{p}(P_n) = \lceil \lg n \rceil$ appears in [57] (without defining either parameter). When n is odd, C_n needs an extra color beyond $\lceil \lg n \rceil$. To prove this, we begin with simple observations about adding an edge.

Lemma 5.2.6. (a) *If e is an edge in a graph G , then $p(G) \leq p(G - e) + 1$.*

(b) *If also $G - e$ is connected, then $\widehat{p}(G) \leq \widehat{p}(G - e) + 1$.*

Proof. (a) Put an optimal parity edge-coloring on $G - e$ and add a new color on e . There is no parity path avoiding e , and any path through e uses the new color exactly once.

(b) Put an optimal spec on $G - e$ and add a new color on e . Let P be a u, v -path in $G - e$, where u and v are the endpoints of e . Suppose that there is an open parity walk W . Note that W traverses e an even number of times, since no other edge has the same color as e . Form W' by replacing each traversal of e by P or its reverse, depending on the direction of traversal of e . Every edge is used with the same parity in W' and W , and the endpoints are unchanged, so W' is an open parity walk in $G - e$. This is a contradiction. \square

Lemma 5.2.6(b) does not hold when $G - e$ is disconnected (see Example 5.2.8).

Theorem 5.2.7. *If n is odd, then $p(C_n) = \widehat{p}(C_n) = \lceil \lg n \rceil + 1$.*

Proof. Lemma 5.2.6(b) yields the upper bound, since $\widehat{p}(P_n) = \lceil \lg n \rceil$.

For the lower bound, we show first that $\widehat{p}(C_n) = p(C_n)$ (this and Lemma 5.2.6(a) yield an alternative proof of the upper bound). Let W be an open walk, and let W' be the subgraph formed by the edges with odd usage in W . The sum of the usage by W of edges incident to a vertex x is odd if and only if x is an endpoint of W . Hence W' has odd degree precisely at the endpoints of W . Within C_n , this requires W' to be a path P joining the endpoints of W . Under a parity edge-coloring f , some color has odd usage along P , and this color has odd usage in W . Hence f has no open parity walk, and every parity edge-coloring is a spec.

It now suffices to show that $\widehat{p}(C_n) \geq p(P_{2n})$. Given a spec f of C_n , we form a parity edge-coloring g of P_{2n} with the same number of colors. Let v_1, \dots, v_n be the vertices of C_n in order, and let $u_1, \dots, u_n, w_1, \dots, w_n$ be the vertices of P_{2n} in order. Define g by letting $g(u_i u_{i+1}) = g(w_i w_{i+1}) = f(v_i v_{i+1})$ for $1 \leq i \leq n - 1$ and letting $g(u_n w_1) = f(v_n v_1)$.

Each path in P_{2n} corresponds to an open walk in C_n or to one trip around the cycle. There is no parity path of the first type, since f is a spec. There is none of the second type, since C_n has odd length. \square

The “unrolling” technique of Theorem 5.2.7 leads to an example G with $\widehat{p}(G) > p(G)$, which easily extends to generate infinite families.

Example 5.2.8. Form a graph G by identifying a vertex of K_3 with an endpoint of P_8 . Since $p(K_3) = p(P_7) = 3$, adding the connecting edge yields $p(G) \leq 4$ (see Lemma 5.2.6(a)).

We claim that $\widehat{p}(G) \geq p(P_{18}) = 5$. We copy a spec f of G onto P_{18} with the path edges doubled. Beginning with the vertex of degree 1 in G , walk down the path, once around the triangle, and back up the path. This walk has length 17; copy the colors of its edges in order to the edges of P_{18} in order to form an edge-coloring g of P_{18} .

Each path in P_{18} corresponds to an open walk in G or a closed walk that traverses the triangle once. There is no parity path of the first type, since f is a spec. There is none of the second type, since such a closed walk has odd length. This proves the claim.

Since $\widehat{p}(K_3) = \widehat{p}(P_7) = 3$, this graph G also shows that adding an edge can change \widehat{p} by more than 1 when G is disconnected.

We know of no bipartite graph G with $\widehat{p}(G) > p(G)$. Nevertheless, it is not true that every optimal parity edge-coloring of a bipartite graph is a spec.

Example 5.2.9. Let G be the graph obtained from C_6 by adding two pendant edges at one vertex. Let W be the spanning walk that starts at one pendant vertex, traverses the cycle, and ends at the other pendant vertex. Let f be the 4-edge-coloring that colors the edges of W in order as a, b, a, c, b, d, c, d . Although f is an optimal parity edge-coloring ($\Delta(G) = 4$), it uses each color twice on the open walk W , so it is not a spec. Changing the edge of color d on the cycle to color a yields a strong parity 4-edge-coloring.

5.3 Complete graphs

When n is a power of 2, we will prove that the complete graph K_n has a unique optimal spec (up to isomorphism), which will help us determine $\widehat{p}(K_n)$ for all n . With a suitable naming of the vertices, we call this edge-coloring of K_n the “canonical” coloring.

Definition 5.3.1. For $A \subseteq \mathbb{F}_2^k$, let $K(A)$ be the complete graph with vertex set A . The *canonical coloring* of $K(A)$ is the edge-coloring f defined by $f(uv) = u + v$, where $u + v$ is binary vector addition. When $n = 2^k$, letting $A = \mathbb{F}_2^k$ yields the canonical coloring of K_n .

Lemma 5.3.2. For $A \subseteq \mathbb{F}_2^k$, the canonical coloring of $K(A)$ is a spec. Consequently, if $n = 2^k$, then $\widehat{p}(K_n) = p(K_n) = \chi'(K_n) = n - 1$.

Proof. If W is an open walk, then its endpoints differ in some bit i . Thus in the canonical coloring the total usage of colors flipping bit i along W is odd, and hence some color has odd usage on W . The canonical coloring of $K(\mathbb{F}_2^k)$ uses $2^k - 1$ colors (the color 0^k is not used). The lower bound follows from $\widehat{p}(G) \geq p(G) \geq \chi'(G) \geq \Delta(G)$. \square

Since every complete graph is a subgraph of the next larger complete graph, we obtain $\widehat{p}(K_n) \leq 2^{\lceil \lg n \rceil} - 1$. In this section, we prove that this upper bound is exact. Our proof is expressed using linear subspaces of binary vector spaces.

Parity edge-coloring relates to a less restrictive problem. A walk of length $2k$ is *repetitive* if the i th and $(k + i)$ th edges have the same color, for $1 \leq i \leq k$. A *Thue coloring* is an edge-coloring with no repetitive path, and the *Thue number* $t(G)$ is the minimum number of colors in a Thue coloring of G . Every parity edge-coloring is a Thue coloring, so $t(G) \leq p(G)$. Alon, Grytczuk, Hałuszczak, and Riordan [5] observed that the canonical coloring yields $t(K_n) \leq 2^{\lceil \lg n \rceil} - 1$. It seems no good lower bounds on $t(K_n)$ are known. Our lower bound on $\widehat{p}(K_n)$ shows that a Thue coloring of K_n better than the canonical coloring must contain an open parity walk.

We use the closure of linear spaces under addition to prove that $\widehat{p}(K_n) \geq 2^{\lceil \lg n \rceil} - 1$. The main idea is to introduce an additional vertex without needing additional colors until a power of 2 is reached. We begin by proving that every optimal spec of K_n is a canonical coloring when n is a power of 2.

Definition 5.3.3. An edge-coloring f of a graph G satisfies the *4-constraint* if whenever $f(uv) = f(xy)$ and $vx \in E(G)$, also $uy \in E(G)$ and $f(uy) = f(vx)$.

Lemma 5.3.4. If f is a parity edge-coloring in which every color class is a perfect matching, then f satisfies the 4-constraint.

Proof. Otherwise, given $f(uv) = f(xy)$, the edge of color $f(vx)$ incident to u forms a parity path of length 4 with uv , vx , and xy . \square

Theorem 5.3.5. *If f is a parity edge-coloring of K_n in which every color class is a perfect matching, then f is a canonical coloring and n is a power of 2.*

Proof. Every edge is a canonically colored K_2 . Let R be a largest vertex set such that $|R|$ is a power of 2 and f restricts to a canonical coloring on R . Define j by $|R| = 2^{j-1}$. Let ϕ be a bijection from R to \mathbb{F}_2^{j-1} under which f is the canonical coloring. We prove the claim by showing that if $|R| < n$, then f canonically colors a subgraph twice as big.

Since f is canonical on R , every color used within R by f is used on a perfect matching of R . The vertices of R have other neighbors in K_n , so there is a color c not used within R . Since c is used on a perfect matching, c matches R to some set U . Let $R' = R \cup U$. Define $\phi': R' \rightarrow \mathbb{F}_2^j$ as follows: for $x \in R$, obtain $\phi'(x)$ by appending 0 to $\phi(x)$; for $x \in U$ obtain $\phi'(x)$ by appending 1 to $\phi(x')$, where x' is the neighbor of x in color c . Within R' , we henceforth refer to the vertices by their names under ϕ' .

By Lemma 5.3.4, the 4-constraint holds for f . The 4-constraint copies the coloring from the edges within R to the edges within U . That is, consider $x', y' \in U$ arising from $x, y \in R$, with $f(xx') = f(yy') = c$. Now $f(x'y') = f(xy) = x + y = x' + y'$, using the 4-constraint, the fact that f is canonical on R , and the definition of ϕ' . Hence f is canonical within U .

Finally, let u be the name of the color on the edge $0^j u$, for $u \in U$. For any $v \in R$, let $w = u + v$; note that $w \in U$. Both $0^j v$ and uw have color v , since f is canonical within R and within U . By applying the 4-constraint to $\{v0^j, 0^j w, uw\}$, we conclude that $f(uv) = f(0^j w) = w$. Since $w = u + v$, this completes the proof that f is canonical on R' . \square

For a proper edge-coloring of K_n , the 4-constraint is equivalent to the property that the six edges on any four vertices receive three colors or six colors. Independently of our proof, Keevash and Sudakov [64] proved that if that property holds for a proper edge-coloring with $n - 1$ colors (one where every color class is a perfect matching), then n is a power of 2. They also observed that the canonical coloring has this property. Their proof appears also in [13].

In connection with the uniqueness result, Mubayi asked whether a stability property holds. That is, when n is a power of 2, does there exist a parity edge-coloring or a spec of K_n that has only $(1 + o(1))n$ colors but is “far” from the canonical coloring?

The main result needs several algebraic observations. Relative to any k -edge-coloring f , the *parity space* L_f is the set of parity vectors of closed walks. We note that L_f is a linear subspace of \mathbb{F}_2^k .

Lemma 5.3.6. *If f is an edge-coloring of a connected graph G , then L_f is a binary vector space.*

Proof. Since $L_f \subseteq \mathbb{F}_2^k$, it suffices to show that L is closed under addition. Given a u, u -walk W and a v, v -walk W' , let P be a u, v -path in G , and let \bar{P} be its reverse. Following W, P, W', \bar{P} in succession yields a u, u -walk with parity vector $\pi(W) + \pi(W')$. \square

For a vector space L , let $w(L)$ be the least number of nonzero coordinates of a nonzero vector in L (set $w(L) = \infty$ if $L = \{0\}$). For an edge-coloring f of K_n , $w(L_f)$ determines whether f is a spec. Let the *weight* of a vector in \mathbb{F}_2^k be the number of nonzero entries.

Lemma 5.3.7. *If an edge-coloring f of a graph G is a spec, then $w(L_f) \geq 2$. The converse holds when $G = K_n$.*

Proof. If the parity vector of a closed walk W has weight 1, then one color has odd usage in W (say on edge e). Now $W - e$ is an open parity walk, and f is not a spec.

If f is not a spec, then $\pi(W') = 0$ for some open walk W' . In K_n , the ends of W' are adjacent, and adding that edge yields a closed walk whose parity vector has weight 1. \square

In fact, the minimum weight of a non-zero vector in the parity space of an optimal spec is at least three.

Lemma 5.3.8. *If f is an optimal spec of K_n , then $w(L_f) \geq 3$.*

Proof. If W is a closed walk having odd usage for colors a and b only, then form f' by merging the colors a and b into a single color a' . We use Lemma 5.3.7 repeatedly. Since f is optimal, there is a closed walk Z on which f' has odd usage for only one color c . If $c = a'$, then f has odd usage on Z for only a or b ; hence $c \neq a'$. Now since f has odd usage for at least two colors on Z , both a and b must also have odd usage on Z . Now in L_f we have $\pi(W) + \pi(Z)$ with weight 1. \square

Lemma 5.3.9. *For any colors a and b in an optimal spec f of K_n , there is some closed walk W on which the colors having odd usage are a , b , and one other.*

Proof. We use Lemma 5.3.7 repeatedly. Since f is optimal, merging the colors a and b into a single color a' yields an edge-coloring f' that is not a spec. Hence under f' there is a closed walk W on which f' has odd usage for only one color c . Also $c \neq a'$, since otherwise f has odd usage on W for only a or b . With $c \neq a'$ and the fact that f has odd usage for at least two colors on W , both a and b also have odd usage on W , and W is the desired walk. \square

The same idea as in Lemma 5.3.9 shows that $w(L_f) \geq 3$ when f is an optimal spec of K_n , but we do not need this observation. We note, however, that the condition $w(L_f) \geq 3$ is the condition for L_f to be the set of codewords for a 1-error-correcting code. Indeed, when $n = 2^k$ and f is the canonical coloring, L_f is a perfect 1-error-correcting code of length $n - 1$.

A *dominating vertex* in a graph is a vertex adjacent to all others. We use $d_H(v)$ and $N_H(v)$ to denote the degree and neighborhood of a vertex v in a graph H .

Lemma 5.3.10. *If f is an edge-coloring of a graph G with a dominating vertex v , then L_f is the span of the parity vectors of triangles containing v .*

Proof. By definition, the span is contained in L_f . Conversely, consider any $\pi(W) \in L_f$. Let S be the set of edges with odd usage in W , and let H be the spanning subgraph of G with edge set S . Since the total usage at each vertex of W is even, H is an even subgraph of G . Hence H decomposes into cycles, which are closed walks, and $\pi(W)$ is the sum of the parity vectors of these cycles.

It therefore suffices to show that S is the set of edges that appear in an odd number of the triangles formed by v with edges of $H - v$. Each edge of $H - v$ is in one such triangle, so we need consider only edges involving v . An edge vw lies in an odd number of these triangles if and only if $d_{H-v}(w)$ is odd, which occurs if and only if $w \in N_H(v)$, since $d_H(w)$ is even. By definition, $vw \in E(H)$ if and only if vw has odd usage in W and hence lies in S . \square

Lemma 5.3.11. *If f is an optimal spec of K_n that uses some color fewer than $n/2$ times, then f extends to a spec of K_{n+1} using the same colors.*

Proof. View K_{n+1} as arising from K_n by adding a vertex u . Let a be a color used fewer than $n/2$ times by f , and let v be a vertex of K_n at which a does not appear.

We use f to define f' on $E(K_{n+1})$. Let f' agree with f on $E(K_n)$, and let $f'(uv) = a$. To define f' on each remaining edge uw , first let $b = f(vw)$. By Lemma 5.3.9, there is a closed walk W with odd usage precisely for a and b and some third color c under f . Let $f'(uw) = c$.

Note that f' uses the same colors as f . It remains only to show that f' is a spec. To do this we prove that $w(L_{f'}) \geq 2$, by showing that $L_{f'} \subseteq L_f$. By Lemma 5.3.10, it suffices to show that $\pi(T) \in L_f$ when T is a triangle in K_{n+1} containing v .

Triangles not containing u lie in the original graph and have parity vectors in L_f . Hence we consider the triangle T formed by $\{u, v, w\}$. Now $\pi(T) = \pi(W) \in L_f$, where W is the walk used to specify $f'(uw)$. \square

Theorem 5.3.12. $\widehat{p}(K_n) = 2^{\lceil \lg n \rceil} - 1$.

Proof. If some color class in an optimal spec is not a perfect matching, then $\widehat{p}(K_n) = \widehat{p}(K_{n+1})$, by Lemma 5.3.11. This vertex absorption cannot stop before the number of vertices reaches a power of 2, because when every color class is a perfect matching the coloring is canonical, by Theorem 5.3.5. It cannot continue past $2^{\lceil \lg n \rceil}$ vertices, since the maximum degree is a lower bound on χ' and \widehat{p} . Hence $\widehat{p}(K_n) = \widehat{p}(K_{2^{\lceil \lg n \rceil}}) = 2^{\lceil \lg n \rceil} - 1$. \square

Corollary 5.3.13. *If f is an optimal spec of K_n , then f is obtained by deleting vertices from the canonical coloring of $K_{2^{\lceil \lg n \rceil}}$.*

Proof. By Lemma 5.3.11, we may extend f to an optimal spec f' of $K_{2^{\lceil \lg n \rceil}}$; by Theorem 5.3.5, f' is the canonical coloring. \square

One may ask whether every edge-coloring of K_n that satisfies the 4-constraint is a spec or a parity edge-coloring. Examples show that the answer is no. Similarly, not every parity edge-coloring of K_n is a spec. Nevertheless, it may be that every *optimal* parity edge-coloring is a spec. We offer the following conjecture, which in [19] we proved for $n \leq 16$.

Conjecture 5.3.14. $p(K_n) = \widehat{p}(K_n)$ for every positive integer n .

5.4 Complete bipartite graphs

To further motivate our focus on complete graphs, we show that our main result strengthens a special case of Yuzvinsky's Theorem on sums of binary vectors. To state it, we need the Hopf–Stiefel function from the theory of quadratic forms.

Definition 5.4.1 (Hopf [59], Stiefel [105]). For positive integers r and s , define $r \circ s$ to be the least integer n such that $(x + y)^n$ is in the ideal of $\mathbb{F}_2[x, y]$ generated by x^r and y^s .

In non-algebraic language, the definition has the following equivalent phrasing: $r \circ s$ is the least n such that $\binom{n}{k}$ is even for each k with $n - s < k < r$. The condition becomes vacuous if $n \geq r + s - 1$, so trivially $r \circ s \leq r + s - 1$.

Theorem 5.4.2 (Yuzvinsky [113]). For $A, B \subseteq \mathbb{F}_2^k$, let $C = \{a + b : a \in A, b \in B\}$. If $|A| = r$ and $|B| = s$, then $|C| \geq r \circ s$.

Generalizations and alternative proofs of Yuzvinsky's Theorem appear in [2], [14], [32]. The theorem is related to our results via a simple formula for the Hopf–Stiefel function recently proved by Plagne [89]. Subsequently, Károlyi [63] gave a short inductive proof. See [33] for a thorough survey of alternative formulas, related results, and generalizations.

Theorem 5.4.3 (Plagne [89], Károlyi [63]). $r \circ s = \min_{k \in \mathbb{N}} \{2^k (\lceil \frac{r}{2^k} \rceil + \lceil \frac{s}{2^k} \rceil - 1)\}$.

When $A = B$ and both have size r , the minimization yields $r \circ r = 2^{\lceil \lg r \rceil}$. Yuzvinsky's Theorem for this case says that every canonical coloring of K_r uses at least $2^{\lceil \lg r \rceil} - 1$ colors. Our result shows that in the more general family of strong parity edge-colorings, it remains true that at least $2^{\lceil \lg r \rceil} - 1$ colors are needed.

The canonical coloring extends to complete bipartite graphs in a natural way: if $A, B \subseteq \mathbb{F}_2^k$ and $K(A, B)$ is the complete bipartite graph with partite sets A and B , then the edge-coloring defined by $f(ab) = a + b$ is a spec. The bound in Yuzvinsky's Theorem is always tight (see [32]); that is, for $r, s \leq 2^k$ there exist $A, B \subseteq \mathbb{F}_2^k$ with $|A| = r$, $|B| = s$, and $|C| = r \circ s$. Consequently, $\widehat{p}(K_{r,s}) \leq r \circ s$. We conjecture that equality holds. A direct proof in the graph-theoretic setting would strengthen all cases of Yuzvinsky's Theorem.

Conjecture 5.4.4. $\widehat{p}(K_{r,s}) = r \circ s$.

Yuzvinsky's Theorem as stated describes a bipartite situation, with the application to complete graphs as a special case. This relationship extends to specs, which means that proving the special case of Conjecture 5.4.4 for $r = s = n$ also implies that on K_n . That implication uses the following proposition.

Proposition 5.4.5. $\widehat{p}(K_{n,n}) \leq \widehat{p}(K_n) + 1$.

Proof. Let f be a spec of K_n with vertex set u_1, \dots, u_n . Given $K_{n,n}$ with partite sets v_1, \dots, v_n and w_1, \dots, w_n , let $f'(v_i w_j) = f(u_i u_j)$ when $i \neq j$, and give a single new color to all $v_i w_i$ with $1 \leq i \leq n$. A parity walk W' under f' starts and ends in the same partite set. Let W be the walk obtained by mapping it back to K_n , which collapses v_i and w_i into u_i , for each i . The edges that had the new color disappear; this number of edges is even, since W' was a parity walk. Hence W is a parity walk under f .

Since f is a spec, W is a closed walk in K_n . Hence W' starts and ends at vertices in the same partite set that have the same index. Since $K_{n,n}$ has only one vertex with each index in each partite set, W' is closed. Hence f' is a spec of $K_{n,n}$. \square

Proposition 5.4.5 and canonical colorings of K_n imply that $\widehat{p}(K_{n,n}) \leq 2^{\lceil \lg n \rceil}$ for all n , and Conjecture 5.4.4 requires that equality holds. Toward this special case of the conjecture, we offer the following.

Proposition 5.4.6. *If some optimal spec of $K_{n,n}$ uses a color on at least $n - 1$ edges, then $\widehat{p}(K_{n,n}) = \widehat{p}(K_n) + 1 = 2^{\lceil \lg n \rceil}$. If a color is used $n - r$ times, then $\widehat{p}(K_{n,n}) \geq 2^{\lceil \lg n \rceil} - \binom{r}{2}$.*

Proof. We prove the general statement. Let f be such a spec, and let c be such a color. Let U be one partite set, with vertices u_1, \dots, u_n . Whenever color class c is incident to at least one of distinct vertices $u_i, u_j \in U$, let $P_{i,j}$ be a u_i, u_j -path of length 2 in which one edge has color c under f . Choose these so that $P_{j,i}$ is the reverse of $P_{i,j}$. When c appears at neither u_i nor u_j , leave $P_{i,j}$ undefined.

Let G be the graph obtained from K_n with vertex set v_1, \dots, v_n by deleting the edges $v_i v_j$ such that $P_{i,j}$ is undefined; there are $\binom{r}{2}$ such edges. Define a coloring f' on G by letting $f'(v_i v_j)$ be the color other than c on $P_{i,j}$.

We claim that f' is a spec. Given a parity walk W' under f' , define a walk W in $K_{n,n}$ as follows. For each edge $v_i v_j$ in W' , follow $P_{i,j}$. By construction, the usage in W of each color other than c is even. Hence also the usage of c is even. Hence W is a parity walk under f and therefore is closed. Since W starts and ends at the same vertex $u_i \in U$, also W' starts and ends at the same vertex v_i .

We have proved that every parity walk under f' is closed, so f' is a spec. Hence f' has at least $\widehat{p}(G)$ colors, and f has at least one more. By Lemma 5.2.6(b) and Theorem 5.3.12, $\widehat{p}(G) \geq 2^{\lceil \lg n \rceil} - 1 - \binom{r}{2}$, which completes the proof of the lower bound. \square

Corollary 5.4.7. $\widehat{p}(K_{n,n}) \geq \max_r \min\{2^{\lceil \lg n \rceil} - \binom{r}{2}, \frac{n^2}{n-r-1}\}$.

Proof. If $E(K_{n,n})$ has a spec with s colors, where $s < 2^{\lceil \lg n \rceil} - \binom{r}{2}$, then by Proposition 5.4.6 no color can be used at least $n - r$ times, and hence $n^2/s \leq n - r - 1$. Thus $\widehat{p}(K_{n,n}) \geq \min\{2^{\lceil \lg n \rceil} - \binom{r}{2}, n^2/(n - r - 1)\}$. \square

With $r = 1$, we conclude that $\widehat{p}(K_{n,n}) \geq 2^k$ when $n > 2^k - 3 - 4/(n - 2)$, since then $n^2/(n - 2) > 2^k - 1$. Thus $\widehat{p}(K_{5,5}) = 8$, and $\widehat{p}(K_{n,n}) = 16$ for $13 \leq n \leq 16$. Using $r = 2$, we obtain $14 \leq \widehat{p}(K_{9,9}) \leq 16$. Our last result obtains $\widehat{p}(K_{2,n})$ exactly.

Theorem 5.4.8. $\widehat{p}(K_{2,n}) = p(K_{2,n}) = 2 \lceil n/2 \rceil$.

Proof. The upper bound follows from $\widehat{p}(K_{2,n}) \leq 2 \circ n = 2 \lceil n/2 \rceil$. For the lower bound, since $\Delta(K_{2,n}) = n$ for $n \geq 2$, it suffices to show that n must be even when f is a parity edge-coloring of $K_{2,n}$ with n colors. Let $\{x, x'\}$ be the partite set of size 2. Each color appears at both x and x' . If color a appears on xy and $x'y'$, then $f(xy) = f(x'y')$, since otherwise the colors a and $f(xy)$ form a parity path of length 4.

Hence y and y' have the same pair of incident colors. Making this argument for each color partitions the vertices in the partite set of size n into pairs. Hence n is even. \square

5.5 Open problems

Many interesting questions remain about parity edge-coloring and strong parity edge-coloring. We have already mentioned several and collect them here with additional questions.

In [19], we prove the first conjecture for $n \leq 16$. Here and in [19], we prove various special cases of the second conjecture, which yield further special cases of the first.

Conjecture 5.5.1. $p(K_n) = 2^{\lceil \lg n \rceil} - 1$ for all n .

Conjecture 5.5.2. $p(K_{n,n}) = \widehat{p}(K_{n,n}) = 2^{\lceil \lg n \rceil}$ for all n .

For complete bipartite graphs in general, the full story would be given by proving Conjecture 5.5.3, which we restate here for completeness.

Conjecture 5.5.3. $\widehat{p}(K_{r,s}) = r \circ s$ for all r and s .

We have exhibited families of graphs G such that $\widehat{p}(G) > p(G)$ (see Example 5.2.8), but the difference is only 1, and the graphs we obtained all contain odd cycles.

Question 5.5.4. What is the maximum of $\widehat{p}(G)$ when $p(G) = k$?

Conjecture 5.5.5. $p(G) = \widehat{p}(G)$ for every bipartite graph G .

The lower bound in Corollary 5.2.4 naturally leads us to ask which graphs achieve equality. Every spanning subgraph of a hypercube satisfies $p(G) = \lg n(G)$; is the converse true?

Question 5.5.6. Which connected graphs G satisfy $p(G) = \lceil \lg n(G) \rceil$? Which satisfy $\widehat{p}(G) = \lceil \lg n(G) \rceil$?

Motivated by the uniqueness of the optimal spec of K_{2^k} , Dhruv Mubayi suggested studying the “stability” of the result.

Question 5.5.7. Does there exist a parity edge-coloring of K_{2^k} with $(1 + o(1))2^k$ colors that is “far” from the canonical coloring?

Since the factors can be treated independently in constructing a spec, \widehat{p} is subadditive under Cartesian product. Note that $\widehat{p}(P_2 \square P_2) = 2 = \widehat{p}(P_2) + \widehat{p}(P_2)$.

Question 5.5.8. For what graphs G and H does equality hold in $\widehat{p}(G \square H) \leq \widehat{p}(G) + \widehat{p}(H)$? What can be said about $p(G \square H)$ in terms of $p(G)$ and $p(H)$?

Finally, the definitions of parity edge-coloring and spec extend naturally to directed graphs: the parity condition is the same but is required only for directed paths or walks. Hence $p(D) \leq p(G)$ and $\widehat{p}(D) \leq \widehat{p}(G)$ when D is an orientation of G .

For a directed path \vec{P}_m , the constraints are the same as for an undirected path. More generally, if D is an acyclic digraph, and m is the maximum number of vertices in a path in D , then $p(D) = \widehat{p}(D) = \lceil \lg m \rceil$. The lower bound is from any longest path.

For the upper bound, give each vertex x a label $l(x)$ that is the maximum number of vertices in a path ending at x (sources have label 0). Write each label as a binary $\lceil \lg m \rceil$ -tuple. By construction, $l(v) > l(u)$ whenever uv is an edge. To form a spec of D , use a color c_i on edge uv if the i th bit is the first bit where $l(u)$ and $l(v)$ differ. All walks are paths. Any x, y -path has odd usage of c_i , where the i th is the first bit where $l(x)$ and $l(y)$ differ, since no edge along the path can change an earlier bit.

Thus the parameters equal $\lceil \lg n \rceil$ for the n -vertex transitive tournament, which contains \vec{P}_n . This suggests our final question.

Question 5.5.9. What is the maximum of $p(T)$ or $\widehat{p}(T)$ when T is an n -vertex tournament?

Chapter 6

The Chromatic Number of Circle Graphs

The *intersection graph* of a family of sets has a vertex for each set in the family, with vertices adjacent if and only if the corresponding sets intersect. A *circle graph* is the intersection graph of a family of chords of a circle. Given a circle graph G , a family of chords of a circle whose intersection graph is G is a *circle representation* of G . In this chapter, we study the chromatic number of circle graphs that are K_4 -free. This chapter is based on joint work with A. V. Kostochka that appears in [71].

6.1 Introduction

Recall that the *chromatic number* of a graph G , denoted $\chi(G)$, is the minimum size of a partition of $V(G)$ into independent sets. A *clique* is a set of pairwise adjacent vertices, and the *clique number* of G , denoted $\omega(G)$, is the maximum size of a clique in G .

Vertices in a clique must receive distinct colors, so $\chi(G) \geq \omega(G)$ for every graph G . In general, $\chi(G)$ cannot be bounded above by any function of $\omega(G)$. Indeed, there are triangle-free graphs with arbitrarily large chromatic number [84].

When graphs have additional structure, it may be possible to bound the chromatic number in terms of the clique number. A family of graphs \mathcal{G} is *χ -bounded* if there is a function f such that $\chi(G) \leq f(\omega(G))$ for each $G \in \mathcal{G}$. Some families of intersection graphs of geometric objects have been shown to be χ -bounded (see [68]). In particular, the family of circle graphs is χ -bounded. Gyárfás [55, 56] proved that $\chi(G) \leq k^2 2^k (2^k - 2)$ when G is a circle graph with clique number k . Kostochka [70] subsequently improved the bound to $\chi(G) \leq k(k+2)2^k$. For large k , the best known upper bound is due to Kostochka and Kratochvíl [69], who proved that $\chi(G) \leq 50 \cdot 2^k - 32k - 64$. This latter bound holds even in the larger family of intersection graphs of polygons inscribed on a circle. (Circle graphs arise by expanding one endpoint of each chord in a circle representation to form a thin triangle.) Conversely, Kostochka [70, 68] constructed circle graphs with clique number k and chromatic number at least $0.5k(\ln k - 2)$. The exponential gap has remained open for 25 years.

Exact results are known only for circle graphs with clique number at most 2. Karapetyan [62] showed that $\chi(G) \leq 8$ when G is a triangle-free circle graph. Kostochka [70] showed that $\chi(G) \leq 5$, and Ageev [1] constructed a triangle-free circle graph with chromatic

number 5. When G is a circle graph with $\omega(G) = 3$, the general bound in [70] implies that $\chi(G) \leq 120$. In this chapter, we improve the bound to $\chi(G) \leq 44$.

Two sets *overlap* if they have non-empty intersection and neither is contained in the other. For a family X of closed intervals on the real line, the *overlap graph* of X , denoted $G(X)$, is the graph with vertex set X in which x and y are adjacent if and only if they overlap. The family of circle graphs is the same as the family of overlap graphs. Given a family of chords on a circle, cutting the circle at a point and unrolling gives the corresponding overlap representation. We study circle graphs via their overlap representations.

We may assume that intervals in X have distinct endpoints. Indeed, let a be a real number, index the intervals with right endpoint a as x_1, \dots, x_s so that $l(x_1) < \dots < l(x_s)$, and index the intervals with left endpoint a as y_1, \dots, y_t so that $r(y_1) < \dots < r(y_t)$. Perturbing the endpoints at a within a small range does not change the overlap relation between any pair of intervals, unless both intervals in the pair had an endpoint at a . If the perturbation is performed so that $l(y_t) < \dots < l(y_1) < r(x_s) < \dots < r(x_1)$, then the overlap relation of all pairs is preserved.

Definition 6.1.1. An interval $[a, b]$ is a *left-neighbor* of $[c, d]$ if $a < c < b < d$. We use $L_X(u)$ to denote the set of all left-neighbors of an interval u in a family X , or simply $L(u)$ when X is clear from context. Similarly, $[a, b]$ is a *right-neighbor* of $[c, d]$ if $c < a < d < b$, and $R_X(u)$ denotes the set of all right-neighbors of u . We also define the *closed* left and right neighborhoods via $\bar{L}_X(u) = L_X(u) \cup \{u\}$ and $\bar{R}_X(u) = R_X(u) \cup \{u\}$. For each interval u , we use $l(u)$ to denote the left endpoint of u and $r(u)$ to denote the right endpoint of u .

The *inclusion order* is defined by containment. The *endpoint order* is defined by putting $x \leq y$ if and only if $l(x) \leq l(y)$ and $r(x) \leq r(y)$. Note that $x \leq y$ in the endpoint order if and only if x comes before y in both the left-endpoint order and the right-endpoint order. Note that any two distinct intervals are comparable in exactly one of the inclusion order and the endpoint order.

6.2 Clean circle graphs

Kostochka [70] constructed circle graphs having clique number k and chromatic number as large as $\Omega(k \log k)$. In this section, we show that if a circle graph G satisfies particular additional properties, then $\chi(G) \leq 2k - 1$.

Definition 6.2.1. If S is a set of intervals, then the *center* of S is the intersection of the intervals in S . A family of intervals X is *clean* if no interval is contained in the intersection of two overlapping intervals in X . A circle graph is *clean* if it is the overlap graph of a clean family of intervals.

A set S of vertices in a graph G is a *cutset* if $G - S$ is disconnected. When S is a cutset in G , the graphs induced by the union of S and the vertices of a component of $G - S$ are *S-lobes*. To color G , it suffices to color the S -lobes so that the colorings agree on S . To ensure that S is colored in the same way in all S -lobes, our inductive hypothesis prescribes the way in which S is colored.

Definition 6.2.2. A subset A of a poset P is a *down-set* if $y \in A$ whenever $y \leq x$ for some $x \in A$. For an element $z \in P$, we use $D[z]$ to denote the down-set $\{y \in P: y \leq z\}$ and $D(z)$ to denote the down-set $\{y \in P: y < z\}$. The *height* of an element $x \in X$ is equal to the size of a maximum chain in $D[x]$; note that minimal elements have height 1. When X is a family of intervals, we define $h_X(x)$ (or simply $h(x)$ when X is clear from context) to be the height of x in the endpoint order on X . The *canonical coloring* of a family X of intervals assigns $h(x)$ to each interval $x \in X$. A coloring f of a family X of intervals is *canonical*, and we say that f is *canonical* on X , if the color classes of f form the same partition of X as the color classes of the canonical coloring.

Note that the canonical coloring is a proper coloring; if x and y overlap, then they are comparable in the endpoint order, and therefore $h(x) \neq h(y)$. In what follows, we develop the tools needed to construct proper colorings of clean circle graphs that are canonical on every right neighborhood.

Proposition 6.2.3. *In a clean family of intervals, let x be an interval with $h(x) \geq 2$. If y is chosen from $D(x)$ to maximize $l(y)$, then $h(x) = h(y) + 1$.*

Proof. Let $k = h(x)$; we use induction on k . When $k = 2$, the statement is trivial. Suppose $k \geq 2$. Since $h(y) < h(x)$, it suffices to show that $h(y) \geq h(x) - 1$. Since $h(x) = k$, there is a chain z_1, \dots, z_k with $z_k = x$. We may assume that $y \neq z_{k-1}$, so the choice of y yields $l(z_{k-1}) < l(y)$. Therefore $l(z_{k-2}) < l(z_{k-1}) < l(y)$. Consider the order of $r(y)$ and $r(z_{k-2})$. If $r(y) < r(z_{k-2})$, then y is contained in $z_{k-1} \cap z_{k-2}$, contradicting that the family is clean. Otherwise, $r(y) > r(z_{k-2})$; now $y > z_{k-2}$ and $h(y) \geq k - 1$. \square

Remark 6.2.4. Proposition 6.2.3 requires that the family of intervals is clean.

Proposition 6.2.5. *Let X be a clean family, and let x be an interval in X that contains another interval in X . If $Y = X - \{x\}$, then $h_Y(u) = h_X(u)$ for all $u \in Y$.*

Proof. Let z be an interval in X that is contained in x . Because X is clean, $y < x$ implies $y < z$, and $y > x$ implies $y > z$. Therefore, if C is a chain containing x in the endpoint order on X , then substituting z for x in C yields another chain of the same size. Hence $h_Y(u) \geq h_X(u)$ for all $u \in Y$, and the other inequality holds since $Y \subseteq X$. \square

Proposition 6.2.6. *If X is a family of intervals that share a common point a , and the overlap graph $G(X)$ has clique number k , then the canonical coloring on X uses exactly k colors.*

Proof. Because the canonical coloring is proper, it uses at least k colors. For the other direction, if the canonical coloring uses r colors, then there is a chain C of size r in the endpoint order on X . It follows that C is an independent set in the inclusion order on X . Hence, no two intervals in C are related by containment. However, C is pairwise intersecting because every member of X contains a . It follows that the intervals in C pairwise overlap, and so $k \geq r$. \square

Proposition 6.2.7. *If f is canonical on X , and Y is a down-set of X in the endpoint order, then f is canonical on Y .*

Proof. Because Y is a down-set in X , we have $h_X(x) = h_Y(x)$ for each $x \in Y$. \square

Let X be a family of intervals and let $u \in X$ be an interval that is not inclusion-minimal, where $u = [a, b]$. We define the *subordinate* of u to be the interval with the rightmost right endpoint among all intervals contained in u . Let v be the subordinate of u , where $v = [c, d]$, and define the *modified subordinate* to be the interval v' , where $v' = [c, b]$. The *right-push* operation on u produces the families Y and Y' and a map $\phi : Y \rightarrow Y'$, where $Y = X - u$, $Y' = Y - v + v'$, and $\phi(x) = x$ for $x \neq v$ and $\phi(v) = v'$.

Lemma 6.2.8. *Let X be a family of intervals, let $u \in X$ be an interval that is not inclusion-minimal, and let Y, Y' , and $\phi : Y \rightarrow Y'$ be produced by the right-push operation on u . If X is clean, then the following hold:*

1. *The map ϕ preserves the order of the left endpoints and right endpoints. That is, $l(x) < l(y)$ if and only if $l(\phi(x)) \leq l(\phi(y))$ for each $x, y \in Y$. Similarly, $r(x) < r(y)$ if and only if $r(\phi(x)) < r(\phi(y))$.*
2. *Y and Y' are clean.*
3. *The clique numbers of Y and Y' are both at most the clique number of X .*
4. *For each $w \in Y$ with $w \neq v$, we have $\phi(\overline{R}_Y(w)) = \overline{R}_{Y'}(\phi(w))$.*
5. *$\phi(\overline{R}_Y(v)) \subseteq \overline{R}_{Y'}(\phi(v))$ and $\phi(\overline{R}_Y(v))$ is a lower subset of $\overline{R}_{Y'}(\phi(v))$ in the endpoint order.*

Proof. Define a, b, c, d so that $u = [a, b]$ and $v = [c, d]$.

1. Note that no interval in X has its right endpoint strictly between d and b . Indeed, if there were such an interval w , then either w is contained in u , in which case v is not the subordinate of w , or $\{w, u\}$ is a 2-clique whose center contains v , contradicting that X is clean. In passing from Y to Y' , the right endpoint of v is moved from d to b to form a new interval v' . Doing so preserves the order of the left endpoints and the right endpoints.
2. Because $Y \subseteq X$ and X is clean, we have that Y is clean. Note that x is contained in the center of a 2-clique with $\{y, z\}$ with $y \leq z$ if and only if $l(y) < l(z) < l(x)$ and $r(x) < r(y) < r(z)$. Hence, the property of being clean is determined by the order of the left endpoints and the order of right endpoints. Because ϕ preserves these orders, Y' is also clean.
3. Let k be the clique number of X . Because $Y \subseteq X$, the clique number of Y is at most k . Let $\{x_1, \dots, x_t\}$ be a clique S in Y' with $x_1 < \dots < x_t$, and note that $l(x_1) < \dots < l(x_t) < r(x_1) < \dots < r(x_t)$. Suppose for a contradiction that $t > k$. We have that $x_j = v'$ for some j , or else S is a clique in Y . If $j > 1$, then x_{j-1} cannot have its right endpoint between d and b . Because $d < b$ and $r(x_j) = b$, it follows that $r(x_{j-1}) < d < r(x_j)$. But $d = r(v)$, and so $S - v' + v$ is a clique of size t in Y , a contradiction. Hence it must be that $j = 1$. Recalling that $r(x_1) = r(u) = d$, we have that $l(u) < l(x_2) < \dots < l(x_t) < r(u) < r(x_2) < \dots < r(x_t)$, which implies that $S - v' + u$ is a clique of size t in X , another contradiction.
4. If $x \in \overline{R}_Y(w)$, then passing from x to $\phi(x)$ leaves the left endpoint fixed and possibly increases the right endpoint. Because $w \neq v$ and $\phi(w) = w$, it follows that $\phi(x) \in \overline{R}_{Y'}(\phi(w))$ and so $\phi(\overline{R}_Y(w)) \subseteq \overline{R}_{Y'}(\phi(w))$. Conversely, if $\phi(x) \in \overline{R}_{Y'}(\phi(w))$, then passing from $\phi(x)$ to x leaves the left endpoint fixed and possibly decreases the right endpoint. However, right endpoint must remain above the right endpoint of $\phi(w)$, and so $x \in \overline{R}_Y(w)$. It follows that $\overline{R}_{Y'}(\phi(w)) \subseteq \phi(\overline{R}_Y(w))$.
5. Passing from v to v' increases the right endpoint of v , but in doing so, the right endpoint never crosses the right endpoint of another interval. Hence, each right-neighbor of v in Y is a right-neighbor of v' in Y' , and therefore $\phi(\overline{R}_Y(v)) \subseteq \overline{R}_{Y'}(\phi(v))$. Suppose that $\phi(x), \phi(y) \in \overline{R}_{Y'}(v')$, $\phi(x) \leq \phi(y)$, and $y \in \overline{R}_Y(v)$. It follows that $l(v) = l(v') < l(x) < l(y) < r(v) < r(v') < r(x) < r(y)$, and hence $x \in \overline{R}_Y(v)$ also.

□

Note that because the endpoint order on X only depends on the order of the left endpoints and the order of the right endpoints, a consequence of Lemma 6.2.8 is that ϕ is a poset isomorphism from Y to Y' under the endpoint order.

Definition 6.2.9. A coloring f of a family of intervals X is *good* if, for each $w \in X$, f is canonical on $\overline{R}(w)$.

Note that if f is a good coloring of X , then it follows that f is a proper coloring. While some families of intervals do not admit good colorings with any number of colors, clean families have good colorings.

Proposition 6.2.10. *Let X be a clean family of intervals and let $u \in X$ be a non-minimal element in the inclusion order. If v is chosen from $\{w \in X : w \subseteq u\}$ to minimize the left endpoint, then $h_X(u) = h_X(v)$.*

Proof. We argue that $w < u$ if and only if $w < v$. If $w < u$, then also $w < v$ or else $\{w, u\}$ is a 2-clique with v in the center, contradicting that X is clean. Conversely, if $w < v$, then the extremality of v implies that $w < u$. \square

Lemma 6.2.11. *If X is a clean family and f is the canonical coloring on X , then f is good.*

Proof. Let $z \in X$ and let $S = \overline{R}_X(z)$, and let h_S (resp h_X) be the height function on the endpoint order on S (resp X). We show that for each $u, v \in S$, it holds that $h_S(u) = h_S(v)$ if and only if $h_X(u) = h_X(v)$. For each $k \geq 0$, let $T_k = \{w \in S : h_S(w) = k\}$. Because all elements in T_k have the same height, they are not comparable in the endpoint order, and therefore T is a chain in the inclusion-order. Index the elements of T as u_1, \dots, u_n so that $u_1 \subsetneq u_2 \subsetneq \dots \subsetneq u_n$, and fix $j < n$. We claim there are no intervals in X whose left endpoint is between $l(u_j)$ and $l(u_{j+1})$. Indeed, if there are such intervals, then let v be one that minimizes the left endpoint. Note that $v \subsetneq u_{j+1}$, or else $\{u_{j+1}, v\}$ is a 2-clique with u_j in the center. Also $v \notin S$, or else applying Proposition 6.2.10 to u_{j+1} and v in the family S would give that $h_S(v) = h_S(u_{j+1}) = k$, and hence $v \in T_k$, a contradiction because no interval in T_k has left endpoint between the left endpoints of u_j and u_{j+1} . But now $v \notin S$ implies that v is in the center of the 2-clique $\{z, u_{j+1}\}$, a contradiction. A final application of Proposition 6.2.10 to u_{j+1} and u_j in X gives that $h_X(u_{j+1}) = h_X(u_j)$. It follows that all intervals in T_k have the same height in X .

For the converse, suppose that T_k and $T_{k'}$ with $k < k'$ have the property that all elements in $T_k \cup T_{k'}$ have the same height in the endpoint order on X . Fix $u \in T_{k'}$. Because $h_S(u) = k'$ and $k < k'$, there is an interval $v \in S$ with $v < u$ and $h_S(v) = k$. It follows that $v \in T_k$. But now v and u are comparable in the endpoint order, so they cannot have the same height in X . \square

Lemma 6.2.12. *Let X be a clean family of intervals, let u be a non-minimal element in the inclusion order on X , and obtain Y, Y' , and ϕ from the right-push operation on u . If g' is a good coloring of Y' , then $g' \circ \phi$ is a good coloring of Y .*

Proof. Consider $w \in Y$. Because g' is good on Y' , we have that g' is canonical on $\overline{R}_{Y'}(\phi(w))$. By Lemma 6.2.8, we have that $\phi(\overline{R}_Y(w))$ is a down set of $\overline{R}_{Y'}(\phi(w))$ in the endpoint order (even equality holds when $w \neq v$). By Proposition 6.2.7, we have that g' is canonical on $\phi(\overline{R}_Y(w))$. But $\phi : Y \rightarrow Y'$ is an isomorphism of the endpoint orders on Y and Y' , so $g' \circ \phi$ is canonical on $\overline{R}_Y(w)$. \square

If $a \in \mathbb{R}$, then X^a denotes the subfamily of X consisting of all intervals that contain a in their interior, $X^{>a}$ denotes the subfamily of X consisting of all intervals that are entirely to the right of a , and $X^{<a}$ denotes the subfamily of X consisting of all intervals that are entirely to the left of a .

Proposition 6.2.13. *Let f be a good coloring of X , let α and β be colors, let a be a point on the real line, and suppose that $f(u) \notin \{\alpha, \beta\}$ for each $u \in X^a$. If f' is the coloring of X obtained from f by interchanging α and β on the intervals in $X^{>a}$, then f' is also good.*

Proof. Let $w \in X$ and define c, d so that $w = [c, d]$. If $d > a$, then every interval in $\overline{R}_X(w)$ with a color in $\{\alpha, \beta\}$ is in $X^{>a}$, and so the change in colors does not alter the partition on $\overline{R}_X(w)$ given by the color classes of f . Similarly, if $d < a$, then every interval in $\overline{R}_X(w)$ with a color in $\{\alpha, \beta\}$ is in $X^{<a}$, and so none of these intervals change colors. If $d = a$, then increase a by a small amount and apply the proposition again. \square

Theorem 6.2.14. *If X is a clean family of intervals with clique number $k \geq 1$, then there is a good coloring f of X using at most $2k - 1$ colors.*

Proof. By induction on $|X|$; we may assume $|X| \geq 1$ and $k \geq 2$. Let x be the interval in X which minimizes $l(x)$. If $R(x) = \emptyset$, then x has no neighbors. Let $Y = X - x$, apply induction to Y to obtain good coloring g of Y , and extend g to a coloring f of X by assigning an arbitrarily chosen color to x . Clearly, f is canonical on each right-neighborhood.

Therefore, we may assume that x has right-neighbors. Choose $y \in R(x)$ to minimize $l(y)$, and define a and b so that $y = [a, b]$. Let $Y_1 = \{z \in X : l(z) \leq b\}$ and $Y_2 = \{z \in X : r(z) \geq b\}$. Note that $x \notin Y_2$ and therefore $Y_2 \subsetneq X$. If also $Y_1 \subsetneq X$, then we may apply induction to Y_1 and Y_2 to obtain respective good colorings g_1 and g_2 . Note that $Y_1 \cap Y_2 = \{z \in X : l(z) \leq b \leq r(z)\}$, and because y is inclusion-maximal, $Y_1 \cap Y_2 = \overline{R}_X(y)$. Consequently, all right-neighbors of y survive in Y_1 and Y_2 , and hence $\overline{R}_X(y) = \overline{R}_{Y_1}(y) = \overline{R}_{Y_2}(y)$, which implies that g_1 and g_2 are canonical on $Y_1 \cap Y_2$. Hence, after relabeling the color names, we obtain a coloring g of X that is a common extension of g_1 and g_2 . Clearly, g uses at most $2k - 1$ colors; it remains to show that g is canonical on each right-neighborhood. Consider $u \in X$. If $r(u) \leq b$, then $\overline{R}_X(u) \subseteq Y_1$ and so $\overline{R}_X(u) = \overline{R}_{Y_1}(u)$, which implies that g is canonical on $\overline{R}_X(u)$. Otherwise, $\overline{R}_X(u) \subseteq Y_2$, and so $\overline{R}_X(u) = \overline{R}_{Y_2}(u)$, which again implies that g is canonical on $\overline{R}_X(u)$.

Hence, we may assume $X = Y_1$. Next, we consider the case that x is not inclusion-minimal. Let v be the subordinate of x , let v' be the modified subordinate of x , and obtain Y, Y' , and ϕ from the right-push operation on x . By Lemma 6.2.8, we have that Y and Y' are clean with clique number at most k . By induction and Lemma 6.2.12, obtain good colorings g' of Y' and $g_0 = g' \circ \phi$ of Y using at most $2k - 1$ colors. Extend g_0 to a coloring g of X by defining $g(w) = g_0(w)$ for $w \neq x$ and $g(x) = g_0(v) = g'(v')$. Clearly, g uses at most $2k - 1$ colors. We claim that g is a good coloring. First, note that because x minimizes $l(x)$, we have that $x \in \overline{R}_X(w)$ implies that $w = x$. Therefore g inherits the canonical coloring of g_0 on $\overline{R}_X(w)$ whenever $w \neq x$. Finally, note that because X is clean, we have that $R_X(x) = R_{Y'}(v')$ and hence g inherits the canonical coloring on $\overline{R}_X(x)$ from the canonical coloring of g' on $\overline{R}_{Y'}(v')$.

Hence, we may assume that x is inclusion-minimal; it follows that $y \in \overline{R}_X(w)$ implies that $w \in \{x, y\}$. Next, we consider the case that y is not inclusion-minimal. Let v be the subordinate of y , let v' be the modified subordinate, and obtain Y, Y' and ϕ from the push operation. By Lemma 6.2.8, we have that Y and Y' are clean with clique number at most k . By induction and Lemma 6.2.12, obtain good colorings g' of Y' and $g_0 = g' \circ \phi$ of Y using at most $2k - 1$ colors. We use g_0 to construct a good coloring of X . Because $Y = X - x$, to extend a good coloring of Y to a good coloring of X , we must assign a color to y so that the coloring remains canonical on each closed right-neighborhood. Because y is only in the closed right-neighborhood of x and y , we need only verify that the coloring is canonical on $\overline{R}_X(x)$ and $\overline{R}_X(y)$.

We consider two subcases. First, suppose that y is inclusion-minimal in $\overline{R}_X(x)$. Because y is chosen from $\overline{R}_X(x)$ to minimize $l(y)$, it follows that $x < y < z$ for every $z \in \overline{R}_X(x) - \{x, y\}$. With $Z_1 = \overline{R}_X(x)$ and $Z_2 = \overline{R}_Y(x) = \overline{R}_X(x) - \{y\}$, this implies that two elements have the same height in Z_2 if and only if they have the same height in Z_1 , and y is the only element of height 1 in Z_1 . Consequently, an extension of g_0 to X is canonical on $\overline{R}_X(x)$ if and only if it assigns y a color that is not used on any other interval in $\overline{R}_X(x)$. Similarly, $y < z$ for each $z \in \overline{R}_X(y) - \{y\}$ and hence an extension of g_0 to X is canonical on $\overline{R}_X(y)$ if and only if y is assigned a color that is not used on any other interval in $\overline{R}_X(y)$. Because g_0 is canonical on $\overline{R}_Y(x)$ and the clique number of $\overline{R}_Y(x)$ is at most $k - 1$ (indeed, every maximal clique in $\overline{R}_X(x)$ contains y), it follows that g_0 uses at most $k - 1$ colors on $\overline{R}_Y(x)$. Also, g' uses at most k colors on $\overline{R}_{Y'}(v')$, and hence g_0 uses at most $k - 1$ colors on $R_X(y)$ (indeed, $g'(v')$ is used on $v' \in \overline{R}_{Y'}(v')$ but is not used on any interval in $R_X(y)$). Because $2k - 1$ colors are available and at most $2k - 2$ provide conflicts, one color remains available for assignment to y .

The second subcase is that y is not inclusion-minimal in $\overline{R}_X(x)$. Let z be the interval that minimizes $l(z)$ among all intervals in $\overline{R}_X(x)$ that are contained in y . Note that z is also the interval that minimizes $l(z)$ among all that are contained in y . Let $\alpha = g_0(z)$. By

Proposition 6.2.10, the height of y and the height of z are the same in all subsets of X containing z and y . By induction, we have that g_0 is canonical on $\overline{R}_X(x) - y$. Applying Proposition 6.2.5 to $\overline{R}_X(x)$, an extension of g_0 to X is canonical on $\overline{R}_X(x)$ if and only if y is assigned color α . Also, an extension of g_0 to X is canonical on $\overline{R}_X(y)$ if and only if y is assigned a color different from every other interval in $\overline{R}_X(y)$. If α is not used on $R_X(y)$, then we may assign α to y . Otherwise, we first modify g_0 before extending to X . Note that z is inclusion-maximal in Y , and let a be a point slightly to the right of $r(z)$. Because z is inclusion-maximal in Y , every interval in Y that contains a is in $R_Y(z)$. Let A be the set of colors that g_0 uses on intervals containing a . Because g_0 is canonical on $\overline{R}_Y(z)$, at most k colors are used on these intervals; because g_0 uses α on $z \in \overline{R}_Y(z)$, we have $\alpha \notin A$ and hence $|A| \leq k - 1$. Let B be the set of colors that g_0 uses on intervals in $R_X(y)$. Because g' is canonical on $\overline{R}_{Y'}(v')$, $R_X(y) = \overline{R}_{Y'}(v') - \{v'\}$, and v' overlaps with every other interval in $\overline{R}_{Y'}(v')$, we have that $|B| \leq k - 1$. Let β be a color that g_0 uses but is not contained in $A \cup B$. Obtain g_1 from g_0 by applying Proposition 6.2.13 with colors $\{\alpha, \beta\}$ at point a . Note that because $\beta \notin B$, we have that g_1 does not use α on any interval in $R_X(y)$. Also, $g_1(z) = \alpha$ and an extension of g_1 to X is canonical on $\overline{R}_X(x)$ if and only if y is assigned color α . Therefore, we obtain a good coloring of X from g_1 by assigning y the color α .

Hence, we may assume that both x and y are inclusion-minimal. By Lemma 6.2.11, the canonical coloring on X is good. Because $X - x = \overline{R}_X(y)$ and Proposition 6.2.6 implies that the canonical coloring uses at most k colors on $\overline{R}_X(y)$, the canonical coloring on X uses at most $k + 1$ colors in total. \square

Theorem 6.2.15. *For each $k \geq 1$, there is a clean circle graph G with $\omega(G) = k$ such that every good coloring of G uses at least $2k - 1$ colors.*

Proof. We construct G in stages; see Figure 6.1 for the first stage. Our construction uses a set of $k - 1$ intervals V that induce a clique in the overlap graph and a set of $k - 1$ intervals V' that form a chain under inclusion. Let $V = \{v_1, \dots, v_{k-1}\}$ and let $V' = \{v'_1, \dots, v'_{k-1}\}$, indexed so that $v_1 < \dots < v_{k-1}$ and $v'_1 \supseteq \dots \supseteq v'_{k-1}$. The left endpoint of v'_j is placed slightly to the left of $l(v_j)$, and the right endpoints of intervals in V' satisfy $r(v'_1) \geq \dots \geq r(v'_{k-1})$. Next, add v_0 so that v_0 is a left-neighbor of all intervals in $V \cup V'$, and add v'_0 so that v'_0 is a right-neighbor of all intervals in V but contained in all intervals in V' .

Because a good coloring must be canonical on $\overline{R}(v_0)$, it follows that a good coloring assigns the same color to v_j and v'_j for $j \geq 1$, and hence $k - 1$ distinct colors are assigned to intervals in V' . Since v'_0 is a right-neighbor of each interval in V , it follows that k distinct colors are assigned to intervals in $V' \cup \{v'_0\}$. These intervals form an independent set in the overlap graph.

In the second stage, we add a set S of $k - 1$ pairwise overlapping intervals such that each interval in S overlaps with intervals in $V' \cup \{v'_0\}$ and no others. A good coloring must use $k - 1$ new colors on S , and hence at least $2k - 1$ colors in total. \square

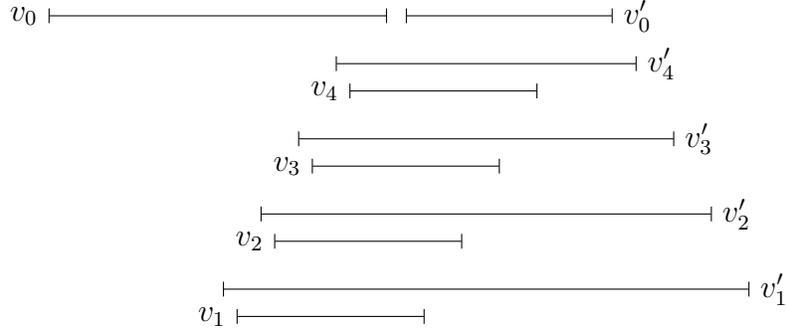


Figure 6.1: Construction in Theorem 6.2.15.

6.3 Chromatic number of K_4 -free circle graphs

In this section, we study the chromatic number of circle graphs with clique number at most 3. Our first task is to explore the structure of segments. A *segment* of a family X is an inclusion-maximal interval in the set of all centers of 2-cliques in X .

Lemma 6.3.1. *Let X be a family of intervals. If $[a, b]$ and $[c, d]$ are overlapping segments of X with $a < c < b < d$, then there exists $x \in X$ with $l(x) \in [a, c)$ and $r(x) \in (b, d]$.*

Proof. Let y_1 and y_2 be overlapping intervals in X with $y_1 < y_2$ and center $[a, b]$. Let z_1 and z_2 be overlapping intervals in X with $z_1 < z_2$ and center $[c, d]$. Note that $l(y_2) = a$ and $r(z_1) = d$. We claim that either $r(y_2) \in (b, d]$ or $l(z_1) \in [a, c)$. Because $r(y_2) > b$ and $l(z_1) < c$, failure requires $r(y_2) > d$ and $l(z_1) < a$. But then we have $l(z_1) < l(y_2) = a < d = r(z_1) < r(y_2)$ which implies that z_1 and y_2 are overlapping intervals in X with center $[a, d]$, contradicting that $[a, b]$ and $[c, d]$ are segments. Hence, either y_2 or z_1 is as required. \square

Lemma 6.3.2. *Let X be a family of intervals. If u_1, \dots, u_t are overlapping segments of X with $u_1 < u_2 < \dots < u_t$, then $[l(u_t), r(u_1)]$ is the center of a $(t + 1)$ -clique in X .*

Proof. For $1 \leq j < t$, apply Lemma 6.3.1 to the segments u_j and u_{j+1} to obtain $z_j \in X$ with $l(z_j) \in [l(u_j), l(u_{j+1}))$ and $r(z_j) \in (r(u_j), r(u_{j+1}))]$. Of the overlapping pair of intervals in X whose center is u_1 , let z_0 be the leftmost in the endpoint order. Similarly, of the overlapping pair of intervals in X whose center is u_t , let z_t be the rightmost in the endpoint order. It follows that $l(z_0) < l(z_1) < \dots < l(z_t) < r(z_0) < r(z_1) < \dots < r(z_t)$ and so $\{z_0, \dots, z_t\}$ is a $(t + 1)$ -clique in X with center $[l(u_t), r(u_1)]$. \square

As a consequence of Lemma 6.3.2, if X has clique number k and Y is the set of all segments of X , then Y has clique-number at most $k - 1$. Moreover, by definition, each interval in Y is inclusion-maximal. Hence the endpoint order on Y is a chain. Index Y as

$\{y_1, \dots, y_t\}$ so that $y_j < y_{j+1}$. If X has clique number at most 3, the structure of Y is highly constrained: each component in the overlap graph of Y is a path.

To find a proper coloring of a family of intervals X with clique number at most 3, we first partition X into distance levels. Let x be the interval in X that minimizes $l(x)$. For each $j \geq 0$, let $S_j = \{y \in X : \text{dist}(x, y) = j\}$, where $\text{dist}(x, y)$ is the distance between x and y in $G(X)$. When $G(X)$ is connected, we say that X is *connected*. Each S_j is called a *level* of X . When X is connected, the levels partition X . Note that if $u \in S_i$, $v \in S_j$, and $|i - j| \geq 2$, then u and v do not overlap. Thus, if $\chi(G(S_i)) \leq t$ for each level S_i , then $\chi(G(X)) \leq 2t$ by using one palette of t colors for $S_0 \cup S_2 \cup \dots$ and another palette of t colors for $S_1 \cup S_3 \cup \dots$. The following lemma is crucial.

Lemma 6.3.3 ([69, 55]). *Let X be a connected family of intervals with levels S_j defined as above. Let $i \geq 1$, let $[a, b]$ be an interval such that $[a, b] \subseteq \bigcup_{y \in S_i} y$. If $z \in S_{i-1}$ and one endpoint of z is in $[a, b]$, then the other endpoint of z is outside $[a, b]$.*

One common application of Lemma 6.3.3 is the following.

Lemma 6.3.4 ([69, 55]). *Let X be a connected family of intervals with levels S_i defined as above. Let $i \geq 1$, let $y \in S_i$, and let T be the set of all intervals in S_i that contain y . There is an interval $z \in S_{i-1}$ such that z overlaps every interval in T .*

Fix a level S_i and let S'_i be the vertices of a component in $G(S_i)$. We partition S'_i into sublevels. Let x be the interval in S'_i that minimizes $l(x)$ and, for each $j \geq 0$, define $S_{i,j} = \{y \in S'_i : \text{dist}(x, y) = j\}$. We say that $S_{i,j}$ is a *sublevel* of S'_i . In the following, we will often apply Lemma 6.3.4 twice; first, with respect to a component S'_i and its sublevels $S_{i,j}$, and next with respect to X and its levels S_i .

Let X be a family of intervals with clique number at most 3, and let S'_i be a component in level S_i . For each segment u of the sublevel $S_{i,j}$ in S'_i , we define the *left-pin* of u to be the point $\max\{r(w) : w \in S_{i,j-1}, w < u, \text{ and } w \text{ overlaps } u\}$ when this set is nonempty; otherwise, the left-pin of u is undefined. Similarly, we define the *right-pin* of u to be the point $\min\{l(w) : w \in S_{i,j-1}, w > u, \text{ and } w \text{ overlaps } u\}$ when this set is nonempty. We need the following facts about pins.

Proposition 6.3.5. *Let X be a family of intervals with clique number at most 3 with a level S_i , a component S'_i of S_i , and sublevels $S_{i,j}$ of S'_i . If u is a segment of $S_{i,j}$ with two well-defined pins, a is the left-pin of u , and b is the right-pin of u , then $a < b$.*

Proof. First, note that $a = b$ is impossible because all intervals in X have distinct endpoints. Suppose for a contradiction that $a > b$. Let x and y be overlapping intervals in $S_{i,j}$ with $x < y$ and center u . Obtain $w_1 \in S_{i,j-1}$ with $r(w_1) = a$ and $w_2 \in S_{i,j-1}$ with $l(w_2) = b$. Applying Lemma 6.3.3 to $[l(x), r(y)]$, we have that $l(w_1) < l(x) < l(y) < l(w_2) < r(w_1) < r(x) < r(y) < r(w_2)$, which implies that $\{w_1, x, y, w_2\}$ is a 4-clique in X , a contradiction. \square

Proposition 6.3.6. *Let X be a family of intervals with clique number at most 3 with a level S_i , a component S'_i of S_i , and sublevels $S_{i,j}$. If u is a segment of $S_{i,j}$ and $z \in S_{i,j}$ is contained in u , then z contains a (well-defined) pin of u .*

Proof. Let x and y be overlapping intervals in $S_{i,j}$ with $x < y$ and center u . Because $|S_{i,j}| > 1$, we have that $i \geq 1$ and $j \geq 1$. Applying Lemma 6.3.4 to S'_i and its sublevels $S_{i,j}$, obtain $w \in S_{i,j-1}$ such that w overlaps x , y , and z . By Lemma 6.3.3, it follows that w has one endpoint inside z and another outside $[l(x), r(y)]$.

Suppose that $l(w) \in z$. Then u has a right-pin b , and $b \leq l(w)$. We claim that $l(z) \leq b$. Indeed, if $b < l(z)$, then there is an interval $w' \in S_{i,j-1}$ with $l(w') = b$. By Lemma 6.3.3, we have that w' contains z and overlaps x and y . Thus, $\{x, y, w'\}$ is a 3-clique in S_i whose center contains z . A final application of Lemma 6.3.4 yields an interval $v \in S_{i-1}$ which overlaps x , y , and w' . Therefore $\{x, y, w', v\}$ is a 4-clique in X , a contradiction. Hence $l(z) \leq b$, and therefore $l(z) \leq b \leq l(w) \leq r(z)$, which implies that z contains the pin b .

The case that $r(w) \in z$ is similar and implies that u has a well defined right-pin a and that z contains a . \square

Proposition 6.3.7. *Let X be a family of intervals with clique number at most 3, a level S_i , a component S'_i of S_i , and sublevels $S_{i,j}$. If u and v are overlapping segments of $S_{i,j}$, then no pin is in the center of $\{u, v\}$.*

Proof. By Lemma 6.3.2, there is a 3-clique $\{x_1, x_2, x_3\}$ in $S_{i,j}$ whose center is the same as the center of $\{u, v\}$. If there is a pin in the center of $\{x_1, x_2, x_3\}$, then there exists $z \in S_{i,j-1}$ with one endpoint in the center of $\{x_1, x_2, x_3\}$. By Lemma 6.3.3, the other endpoint of z is outside $x_1 \cup x_2 \cup x_3$, and so z overlaps each of interval in $\{x_1, x_2, x_3\}$. Therefore $\{x_1, x_2, x_3, z\}$ is a 4-clique in X , a contradiction. \square

We now have the tools available to prove an upper bound on the chromatic number of a circle graph with clique number at most 3.

Theorem 6.3.8. *If X is a family of intervals and $\omega(G(X)) \leq 3$, then $\chi(G(X)) \leq 44$.*

Proof. We may assume that X is connected. We argue that $\chi(G(S_i)) \leq 22$ for each level S_i . Fix a level S_i , and let S'_i be the vertices of a component of $G(S_i)$. We argue that $\chi(G(S'_i)) \leq 22$. Let $S_{i,j}$ be the sublevels of S'_i . It suffices to show that $\chi(G(S_{i,j})) \leq 11$.

Let A be the set of all intervals in $S_{i,j}$ that are contained in the center of a 2-clique in $S_{i,j}$ and let $B = S_{i,j} - A$. Note that B is a clean family, and hence by Theorem 6.2.14, we have that $\chi(G(B)) \leq 5$. We use 5 colors for intervals in B .

It remains to show that $\chi(G(A)) \leq 6$. Let p_1, \dots, p_s be the pins of the segments of $S_{i,j}$, indexed so that $p_1 < \dots < p_s$. Each interval $x \in A$ is contained in a segment of $S_{i,j}$ and therefore contains at least one pin by Proposition 6.3.6. We claim that each interval $x \in A$

contains at most two pins. Proposition 6.3.7 implies that each pin is contained in exactly one segment; hence each segment contains at most two pins. The claim now follows.

For $1 \leq k \leq s$, let T_k be the set of intervals in A that contain p_k and no smaller pin. Since each interval contains a pin, $\{T_1, \dots, T_s\}$ is a partition of A . Note that if x and y are overlapping intervals in A with $x \in T_k$ and $y \in T_{k'}$, then $|k - k'| \leq 2$, since each interval contains at most two pins. Hence we may use the same set of colors on T_k and $T_{k'}$ whenever k and k' are congruent modulo 3. Therefore it suffices to show that $\chi(G(T_k)) \leq 2$ for each T_k .

Note that each T_k has clique number at most 2. Indeed, if T_k contained a 3-clique, then the interval in $S_{i,j-1}$ with endpoint p_k would complete a 4-clique in X . By Proposition 6.2.6, we have $\chi(G(T_k)) \leq 2$ via the canonical coloring. \square

Chapter 7

Acyclic Sets in k -majority Tournaments

A k -majority tournament is a model of the consensus preferences of a group of k individuals. When Π is a set of linear orders on a ground set X , the *majority digraph* of Π has vertex set X and has an edge from u to v if and only if a majority of the orders in Π rank u before v . When Π has size k and k is odd, the majority digraph is a *k -majority tournament*. In this chapter, we explore the maximum size of an acyclic set in k -majority tournaments.

This chapter is based on joint work with D. Schreiber and D. B. West that appears in [78].

7.1 Introduction

In studying generalized voting paradoxes, McGarvey [75] showed that every n -vertex tournament is realizable as a k -majority tournament with $k = 2\binom{n}{2}$. Erdős and Moser [38] improved this by showing that $k = O(n/\log n)$ always suffices, and Stearns [104] showed that $k = \Omega(n/\log n)$ is sometimes necessary.

A *dominating set* in a digraph is a set of vertices S such that for each vertex v not in S , there is a vertex $u \in S$ such that uv is an edge. The *domination number* of a directed graph D , denoted $\gamma(D)$, is the minimum size of a dominating set in D . In general, Erdős [36] showed that n -vertex tournaments can have domination number $\Omega(\log n)$. However, the domination number is bounded within the family of k -majority tournaments; Alon *et al.* [3] proved that every k -majority tournament has domination number at most $O(k \log k)$ and constructed k -majority tournaments with domination number at least $\Omega(k/\log k)$.

In this chapter, we study the extremal problem for acyclic sets in k -majority tournaments. Let $a(D)$ denote the maximum size of an acyclic set in D . Erdős and Moser [38] showed that if T is an n -vertex tournament, then $a(T) \geq \lfloor \lg n \rfloor + 1$, and they showed that almost every n -vertex tournament T satisfies $a(T) \leq 2\lfloor \lg n \rfloor + 1$, using a probabilistic argument. A similar argument shows that almost every n -vertex graph has clique number and independence number logarithmic in n .

Acyclic sets in tournaments are an analogue of cliques and independent sets in graphs. Given any linear ordering of the vertices of T , form a graph G on $V(T)$ by making vertices adjacent if the edge joining them in T is in increasing order. Cliques and independent

sets in G are acyclic in T . Hence, $a(T)$ is at least $\max\{\alpha(G), \omega(G)\}$, where $\alpha(G)$ is the independence number of G and $\omega(G)$ is the clique number of G . On the other hand, we will see that $a(T) \leq (\max\{\alpha(G), \omega(G)\})^2$.

Although tournaments may have only logarithmically large acyclic sets, k -majority tournaments always have polynomially sized acyclic sets. Let

$$f_k(n) = \min\{a(T) : T \text{ is an } n\text{-vertex } k\text{-majority tournament}\}.$$

We prove that $f_3(n) \geq \sqrt{n}$ and $f_3(n) \leq 2\sqrt{n} - 1$ when n is a perfect square. We also prove that $f_5(n) \geq n^{1/4}$. For general k , we prove that $n^{c_k} \leq f_k(n) \leq n^{d_k}$, where $c_k = 3^{(k-1)/2}$ and $d_k = \frac{(\ln \ln k) + 2}{(\ln k) - 1} (1 + \frac{\ln k}{n})$. We make heavy use of the Erdős–Szekeres Theorem.

Theorem (Erdős–Szekeres [40]). *Every list of more than $(r - 1)(s - 1)$ distinct integers contains a monotonically increasing sublist of length r or a monotonically decreasing sublist of length s .*

Let Π be a set of linear orderings of a ground set X . A set of elements of X is Π -consistent if it appears in the same order in each member of Π . When Π has even size, a set S of elements of X is Π -neutral if for all distinct $u, v \in S$, element u appears before element v in exactly half the members of Π . Note that if S is $\{\pi_1, \pi_2\}$ -neutral, then π_1 ranks the elements of S in reverse order from π_2 . In the following, we apply the Erdős–Szekeres Theorem in the following form.

Theorem (Erdős–Szekeres [40]). *If π_1 and π_2 are linear orderings of a ground set X and $|X| > (r - 1)(s - 1)$, then there is a $\{\pi_1, \pi_2\}$ -consistent set of size r or a $\{\pi_1, \pi_2\}$ -neutral set of size s .*

Proof. Rename the elements of X so that π_1 is the identity ordering $(1, \dots, n)$, and apply the Erdős–Szekeres Theorem to π_2 . □

Another elementary fact that we will use repeatedly is the following well-known characterization of acyclic tournaments. A *triangle* in a tournament is a (directed) 3-cycle.

Fact 7.1.1. Let T be a tournament. The following are equivalent:

1. T is acyclic.
2. The vertices of T can be indexed as v_1, \dots, v_n so that $v_i v_j$ is an edge in T if and only if $i < j$. The sequence v_1, \dots, v_n is called the *transitive order* of T .
3. T does not contain any triangles.

Of course, when S is an acyclic subset of T , Fact 7.1.1 applies to the subtournament of T induced by S , and S has a transitive order.

The Erdős–Szekeres Theorem allows us to justify our comment that $a(T) \leq \max\{\alpha(G), \omega(G)\}^2$ when G is the graph obtained from a linear order on $V(T)$ by making vertices adjacent when the edge joining them is increasing in the order on $V(T)$. For any acyclic set S in T , let π_1 be the restriction to S of the given linear order on $V(T)$, and let π_2 be the transitive order on S . Now $\alpha(G)$ is the maximum size of a $\{\pi_1, \pi_2\}$ -neutral set, and $\omega(G)$ is the maximum size of a $\{\pi_1, \pi_2\}$ -consistent set. Hence, the Erdős–Szekeres Theorem implies that $\max\{\alpha(G), \omega(G)\} \geq \sqrt{|S|}$.

7.2 Small odd k

In this section, we prove bounds on $f_k(n)$ when k is 3 or 5. When $k = 3$, our bounds differ only by a factor of 2.

The case $k = 3$

Beame and Huynh-Ngoc [8] gave a simple argument that if Π is a set of three orderings of a ground set of n elements, then there is a $\{\pi_i, \pi_j\}$ -consistent set of size $n^{1/3}$. When $|\Pi|$ is fixed, this is asymptotically best possible. Beame, Blais, and Huynh-Ngoc [7] proved that for integers n and k with $k \geq 3$ and $n \geq k^2$, there is a set Π of k orderings of an n -set in which no two orderings have a consistent set of size greater than $16(nk)^{1/3}$.

This implies that that $f_3(n) \geq n^{1/3}$ using only sets that are consistent in two of the orders. By considering also acyclic sets that are neutral in the first two orders, we obtain a better lower bound. The proof uses the Erdős–Szekeres Theorem.

Proposition 7.2.1. $f_3(n) \geq \sqrt{n}$.

Proof. Let T be an n -vertex 3-majority tournament realized by $\{\pi_1, \pi_2, \pi_3\}$. By the Erdős–Szekeres Theorem, there is a $\{\pi_1, \pi_2\}$ -consistent set of size at least \sqrt{n} or a $\{\pi_1, \pi_2\}$ -neutral set of size at least \sqrt{n} . In the first case, this set is an acyclic set.

Otherwise, let S be a $\{\pi_1, \pi_2\}$ -neutral set of size at least \sqrt{n} . Since S is $\{\pi_1, \pi_2\}$ -neutral, it follows that S induces a transitive subtournament of T with vertices in the same order as in π_3 . Hence S is an acyclic set. \square

While Proposition 7.2.1 is simple, it is within a constant factor of the true value of $f_3(n)$.

Theorem 7.2.2. *If n is a perfect square, then $f_3(n) \leq 2\sqrt{n} - 1$.*

Proof. Let $n = r^2$, and let $X = [r] \times [r]$. View X as integer points in the first quadrant in the plane. We describe three orderings of X and argue that they realize a tournament in

which every acyclic set has size at most $2r - 1$. Define linear orderings π_1 , π_2 , and π_3 as follows:

$$\begin{aligned} (u_1, u_2) <_1 (v_1, v_2) &\iff u_2 < v_2 \text{ or } (u_2 = v_2 \text{ and } u_1 < v_1) \\ (u_1, u_2) <_2 (v_1, v_2) &\iff u_2 > v_2 \text{ or } (u_2 = v_2 \text{ and } u_1 < v_1) \\ (u_1, u_2) <_3 (v_1, v_2) &\iff u_1 > v_1 \text{ or } (u_1 = v_1 \text{ and } u_2 < v_2). \end{aligned}$$

Note that these are all lexicographic orderings up to symmetries of the r -by- r grid, and hence they are linear orderings.

Let T be the 3-majority tournament realized by $\{\pi_1, \pi_2, \pi_3\}$. Consider distinct vertices u and v , where $u = (u_1, u_2)$ and $v = (v_1, v_2)$. If u and v differ in both coordinates, then $uv \in E(T)$ if and only if $u_1 > v_1$. Indeed, $\{u, v\}$ is $\{\pi_1, \pi_2\}$ -neutral, and π_3 brakes the tie by listing the vertex with larger first coordinate before the other. However, if $u_2 = v_2$, then $uv \in E(T)$ if and only if $u_1 < v_1$. Finally, if $u_1 = v_1$, then $uv \in E(T)$ if and only if $u_2 < v_2$.

For each $i \in [r]$, let R_i be the row $\{(u_1, u_2) \in X : u_2 = i\}$ and let C_j be the column $\{(u_1, u_2) \in X : u_1 = j\}$. Let S be an acyclic subset of T . We show that $|S| \leq 2r - 1$ by exhibiting an injection from S to the set $\{R_i : i \in [r] - \{1\}\} \cup \{C_j : j \in [r]\}$. For each column C_j , when S contains a vertex in C_j , we map the vertex in $S \cap C_j$ with smallest second coordinate to C_j . For every other vertex u in S , we map u to the row R_i that contains u . Note that no vertex is mapped to R_1 , because this vertex would have the smallest second coordinate in its column.

By construction, no two vertices are mapped to the same column. Hence, if the map fails to be an injection, there are two vertices u and v that are mapped to the same row R_i . It follows that $u = (u_1, i)$ and $v = (v_1, i)$; we may assume that $u_1 < v_1$. Because u is mapped to row R_i , it follows that there is a vertex w in S that is in the same column as u but has a smaller second coordinate. Hence, $w = (u_1, k)$ for some $k < i$. It follows that uvw is a directed triangle in S , contradicting that S is acyclic. \square

Proposition 7.2.1 and Theorem 7.2.2 combine to give the following general bounds on $f_3(n)$.

Theorem 7.2.3. $\sqrt{n} \leq f_3(n) < 2\sqrt{n} + 1$.

Proof. The lower bound is Proposition 7.2.1. For the upper bound, let n' be the smallest perfect square that is at least as large as n . Note that $\sqrt{n'} - \sqrt{n} < 1$. Using the monotonicity of f and Theorem 7.2.2, it follows that $f_3(n) \leq f_3(n') \leq 2\sqrt{n'} - 1 < 2\sqrt{n} + 1$. \square

The case $k = 5$

Because adding a linear ordering and its reverse to Π does not change the majority digraph, every k -majority tournament is a $(k + 2)$ -majority tournament. It follows that $f_{k+2}(n) \leq$

$f_k(n)$. Currently, our best upper bound on $f_5(n)$ is $f_5(n) \leq f_3(n) < 2\sqrt{n} + 1$.

Theorem 7.2.4. $f_5(n) \geq n^{1/4}$.

Proof. Let T be an n -vertex 5-majority tournament realized by $\{\pi_1, \dots, \pi_5\}$. Apply the Erdős–Szekeres Theorem to π_1 and π_2 to obtain a $\{\pi_1, \pi_2\}$ -consistent or a $\{\pi_1, \pi_2\}$ -neutral set S of size at least \sqrt{n} . Let $r = |S|$. If S is $\{\pi_1, \pi_2\}$ -neutral, then the subtournament on S is an r -vertex 3-majority tournament realized by $\{\pi_3, \pi_4, \pi_5\}$. By Proposition 7.2.1, S contains an acyclic set of size \sqrt{r} , and therefore T contains an acyclic set of size $n^{1/4}$.

Otherwise, S is $\{\pi_1, \pi_2\}$ -consistent. Let P be the poset that is the intersection of the orders in $\{\pi_3, \pi_4, \pi_5\}$, so $u <_P v$ if and only if all three orders list u before v . Let P' be the subposet of P induced by S . If P has no chain with more than t elements, then iteratively stripping off the antichain of maximal elements yields a partition of P into at most t chains. Hence P' contains a chain or an antichain of size at least \sqrt{r} . The elements of any chain of size at least \sqrt{r} in P' form a $\{\pi_3, \pi_4, \pi_5\}$ -consistent set, and this set is acyclic in T .

If there is no such chain, then P' has an antichain A of size at least \sqrt{r} . Any two elements in A appear ordered each way among the members of $\{\pi_3, \pi_4, \pi_5\}$. Therefore, A is acyclic in T , with transitive order given by the common order in which they appear in π_1 and π_2 . \square

7.3 General odd k

In this section, we present bounds on $f_k(n)$ for general k . Our bounds are far apart when k is large, but they do show that $f_k(n)$ has polynomial growth for all k , and the degree of the polynomial tends to zero as k grows.

For a family Π of linear orders, a set S is Π -homogeneous if there is a linear order on S and an integer α such that exactly α members of Π list u before v whenever $u <_L v$. Relative to this linear order, we say that α is the *signature* of S . Our lower bound finds a Π -homogeneous set inductively.

Theorem 7.3.1. *Let k be an odd integer. For any family Π of k linear orders of an n -set, there is a Π -homogeneous set of size at least n^{c_k} , where $c_k = 3^{-(k-1)/2}$.*

Proof. The proof is by induction on k . For $k = 1$, the claim is trivial. Suppose that $k \geq 3$, and let $\Pi = \{\pi_1, \dots, \pi_k\}$. By the Erdős–Szekeres Theorem, there is a $\{\pi_{k-1}, \pi_k\}$ -consistent set of size r , where $r \geq n^{2/3}$, or a $\{\pi_{k-1}, \pi_k\}$ -neutral set of size s , where $s \geq n^{1/3}$.

Suppose that S is a $\{\pi_{k-1}, \pi_k\}$ -neutral set of size s . For each $j \in [k]$, let π'_j be the restriction of π_j to S . By the induction hypothesis, there is a $\{\pi'_1, \dots, \pi'_{k-2}\}$ -homogeneous set S' of size at least $s^{c_{k-2}}$. Because S is $\{\pi_{k-1}, \pi_k\}$ -neutral, it follows that S' is Π -homogeneous. Note that S' has size at least $n^{c_{k-2}/3}$, and $c_{k-2}/3 = c_k$.

If there is no such set, then we obtain a $\{\pi_{k-1}, \pi_k\}$ -consistent set S of size r , where $r \geq n^{2/3}$. For each $j \in [k]$, let π'_j be the restriction of π_j to S . By the induction hypothesis,

there is a $\{\pi'_1, \dots, \pi'_{k-2}\}$ -homogeneous set S' of size q , where $q \geq r^{c_{k-2}}$. Let L_1 be the ordering of S' under which S' is $\{\pi'_1, \dots, \pi'_{k-2}\}$ -homogeneous, and let L_2 be the common ordering of S' in π_{k-1} and π_k . Apply the Erdős–Székere Theorem to L_1 and L_2 to obtain an $\{L_1, L_2\}$ -consistent or $\{L_1, L_2\}$ -neutral set S'' of size at least \sqrt{q} .

Let α be the signature of S' relative to L_1 . Whether S'' is $\{L_1, L_2\}$ -consistent or $\{L_1, L_2\}$ -neutral, S'' is Π -homogeneous with signature $\alpha + 2$ or $\alpha - 2$ relative to L_1 , respectively. Note that $|S''| \geq \sqrt{q} \geq r^{c_{k-2}/2} \geq n^{c_{k-2}/3} = n^{c_k}$. \square

Because a Π -homogeneous set is acyclic in the majority tournament, we immediately obtain the following.

Corollary 7.3.2. $f_k(n) \geq n^{c_k}$, where $c_k = 3^{-(k-1)/2}$.

Our upper bound on $f_k(n)$ for general odd k uses the following strategy. We begin with a k -vertex tournament having no large acyclic set; it is a k -majority tournament. We will compose this tournament with itself to obtain larger k -majority tournaments whose acyclic sets are small.

If T_1 and T_2 are tournaments, the *composition* of T_1 and T_2 , denoted $T_1 \circ T_2$, is the tournament T obtained by replacing each vertex u in T_1 with a copy $T_2(u)$ of T_2 and replacing each edge uv in T_1 with an orientation of a biclique so that all edges are directed from $T_2(u)$ to $T_2(v)$. Formally, if $V(T_1) = [r]$ and $V(T_2) = [s]$, then $V(T_1 \circ T_2) = [r] \times [s]$ and $(x_1, x_2)(y_1, y_2)$ is an edge in $T_1 \circ T_2$ if and only if $x_1 y_1 \in E(T_1)$ or $x_1 = y_1$ and $x_2 y_2 \in E(T_2)$.

Proposition 7.3.3. *If T_1 and T_2 are k -majority tournaments, then $T_1 \circ T_2$ is a k -majority tournament.*

Proof. Let T_1 and T_2 be k -majority tournaments with vertex sets $[r]$ and $[s]$ respectively. Suppose that T_1 is realized by $\{\pi_1, \dots, \pi_k\}$ and T_2 is realized by $\{\sigma_1, \dots, \sigma_k\}$. We construct a realizer $\{\tau_1, \dots, \tau_k\}$ for $T_1 \circ T_2$ in the natural way by letting τ_t be the linear ordering of $[r] \times [s]$ obtained by replacing the occurrence of $i \in [r]$ in the list π_t with the sequence $(i, \sigma_t(1)), (i, \sigma_t(2)), \dots, (i, \sigma_t(s))$, where $\sigma_t(j)$ is the j th element listed in σ_t .

Suppose that $(x_1, x_2)(y_1, y_2)$ is an edge in $T_1 \circ T_2$. If $x_1 \neq y_1$, then $x_1 y_1 \in E(T_1)$ and therefore more than half of the linear orders in $\{\pi_1, \dots, \pi_k\}$ list x_1 before y_1 . The corresponding orders in $\{\tau_1, \dots, \tau_k\}$ list all elements with first coordinate x_1 before elements with first coordinate y_1 . If $x_1 = y_1$, then $x_2 y_2 \in E(T_2)$ and therefore more than half of the linear orders in $\{\sigma_1, \dots, \sigma_k\}$ list x_2 before y_2 . The corresponding orders in $\{\tau_1, \dots, \tau_k\}$ list (x_1, x_2) before (y_1, y_2) . It follows that $\{\tau_1, \dots, \tau_k\}$ realizes $T_1 \circ T_2$. \square

Our next proposition is a simple modification of well-known product theorems.

Proposition 7.3.4. $a(T_1 \circ T_2) = a(T_1)a(T_2)$.

Proof. If S_j is an acyclic set in T_j , then $S_1 \times S_2$ is acyclic in $T_1 \circ T_2$. Therefore $a(T_1 \circ T_2) \geq a(T_1)a(T_2)$. Let S be an acyclic set in $T_1 \circ T_2$. The *projection* of S onto $V(T_1)$ is the set $\{u \in V(T_1) : (u, v) \in S \text{ for some } v \in V(T_2)\}$. Let S_1 be the projection of S onto $V(T_1)$. Note that S_1 is acyclic in T_1 , because a cycle induced by S_1 lifts to a cycle induced by S . Also, for each $u \in V(T_1)$, the set of all vertices in S with first coordinate u has size at most $a(T_2)$. It follows that $a(T_1 \circ T_2) \leq |S_1|a(T_2) \leq a(T_1)a(T_2)$. \square

Proposition 7.3.5. *Let T_1 be a tournament on n_1 vertices, and for each $j > 1$, let $T_j = T_{j-1} \circ T_1$. If $\alpha = a(T_1)$ and $n_j = |V(T_j)|$, then $a(T_j) = n_j^{\frac{\ln \alpha}{\ln n_1}}$.*

Proof. Note that $n_j = n_1^j$. Since $\alpha^{j \ln n_1} = n_1^{j \ln \alpha}$, Proposition 7.3.4 yields $a(T_j) = \alpha^j = n_j^{\frac{\ln \alpha}{\ln n_1}}$. \square

Proposition 7.3.5 provides a way of building larger k -majority tournaments from an initial k -majority tournament T_1 ; when $a(T_1)$ is small, the k -majority tournaments produced by Proposition 7.3.5 have no large acyclic set. A randomized construction produces a tournament with a given number of vertices that has no large acyclic set, but not every such tournament is a k -majority tournament. Nevertheless, when the given number of vertices is at most $(k-1)/2$, every tournament is a k -majority tournament. Stronger results are known, but our result only needs the following simple proposition.

Proposition 7.3.6. *Every n -vertex tournament is a $(2n-1)$ -majority tournament.*

Proof. Let T be an orientation of K_n . It is well known that K_n is n -edge-colorable. Let M_1, \dots, M_n be a decomposition of K_n into matchings. We first construct a realizer Π of T with $|\Pi| = 2n$. Each matching contributes two linear orders to Π . If uv is an edge in M_j , then u is listed immediately before v in both orders generated by M_j . The edges are listed arbitrarily in the first order contributed by M_j and in reverse in the second order. If $uv \in E(T)$, then u appears before v exactly $n+1$ times. Hence Π realizes T . Furthermore, deleting any one member of Π leaves u appearing before v at least n times out of the remaining $2n-1$ orders. \square

We now have the tools needed to prove our upper bound on $f_k(n)$ for general k .

Theorem 7.3.7. $f_k(n) \leq n^{d_k}$, where $d_k = \frac{\ln(3 \lg \frac{k+1}{2})}{\ln \frac{k+1}{2}} \left(1 + \frac{\ln \frac{k+1}{2}}{\ln n}\right)$.

Proof. Let $n_1 = (k+1)/2$. Erdős and Moser [38] proved that there is an n_1 -vertex tournament T_1 with $a(T_1) \leq 3 \lg n_1$. Let $\alpha = a(T_1)$. By Proposition 7.3.6, T_1 is a k -majority tournament. Let n be a positive integer, and let n' be the least power of n_1 that is at least as large as n . Note that $n' \leq nn_1$. By Proposition 7.3.5, there is an n' -vertex k -majority tournament T with $a(T) = n'^{\frac{\ln \alpha}{\ln n_1}}$. It follows that

$$f_k(n) \leq f_k(n') \leq (n')^{\frac{\ln \alpha}{\ln n_1}} = (nn_1)^{\frac{\ln \alpha}{\ln n_1}} = n^{\frac{\ln \alpha}{\ln n_1}} \left(1 + \frac{\ln n_1}{\ln n}\right) = n^{d_k},$$

as required. □

Theorem 7.3.7 implies the weaker but algebraically simpler bound $f_k(n) \leq n^{d'_k}$, where $d'_k = \frac{(\ln \ln k) + 2}{(\ln k) - 1} (1 + \frac{\ln k}{n})$. Erdős and Moser [38] also proved that every n -vertex tournament is a k -majority tournament for $k = O(n/\log n)$; equivalently, there is a constant c such that every tournament on $ck \log k$ vertices is a k -majority tournament. Thus we could let T_1 be a tournament with $ck \log k$ vertices such that $a(T_1) = 3 \lg(ck \log k)$. This would produce a very slight improvement in our bound.

Chapter 8

Cycle Spectra of Hamiltonian Graphs

In this chapter, we prove that every Hamiltonian graph with n vertices and m edges has cycles of at least $\sqrt{\frac{4}{7}(m-n)}$ different lengths. The coefficient $4/7$ cannot be increased above 1, since when $m = n^2/4$ there are $\sqrt{m-n+1}$ cycle lengths in $K_{n/2, n/2}$. For general m and n there are examples having at most $2 \left\lceil \sqrt{2(m-n+1)} \right\rceil$ different cycle lengths.

This chapter contains joint work with D. Rautenbach, F. Regen, and D. B. West [77].

8.1 Introduction

The *cycle spectrum* of a graph G is the set of lengths of cycles in G . A cycle containing all vertices of G is a *spanning* or *Hamiltonian cycle*, and a graph having such a cycle is a *Hamiltonian graph*. An n -vertex graph is *pancyclic* if its cycle spectrum is $\{3, \dots, n\}$. All our graphs have no loops or multiple edges. Let $d_G(x)$ denote the degree in G of a vertex x (its number of neighbors). A graph is k -regular if every vertex has degree k .

Interest in cycle spectra arose from Bondy's "Metaconjecture" (based on [17]) that sufficient conditions for existence of Hamiltonian cycles usually also imply pancyclicity, with possibly a small family of exceptions. For example, Bondy [17] showed that the sufficient condition on n -vertex graphs due to Ore [87] ($d_G(x) + d_G(y) \geq n$ whenever x and y are nonadjacent vertices) implies also that G is pancyclic or is the complete bipartite graph $K_{\frac{n}{2}, \frac{n}{2}}$. Schmeichel and Hakimi [98] showed that if a spanning cycle in an n -vertex graph G has consecutive vertices with degree-sum at least n , then G is pancyclic or bipartite or lacks only $n-1$ from the spectrum, with the latter cases occurring only when the degree-sum is exactly n . Bauer and Schmeichel [6] used this to give unified proofs that the conditions for Hamiltonian cycles due to Bondy [18], Chvátal [25], and Fan [41] also imply pancyclicity, with a small family of exceptions. Further results about the cycle spectrum under degree conditions on selected vertices in a spanning cycle appear in [43] and [99].

At the 1999 conference "Paul Erdős and His Mathematics", Jacobson and Lehel proposed the opposite question: *When sufficient conditions for spanning cycles are relaxed, how small can the cycle spectrum be if the graph is required to be Hamiltonian?* For example, consider regular graphs. Bondy's result [17] implies that $\lceil n/2 \rceil$ -regular graphs other than $K_{\frac{n}{2}, \frac{n}{2}}$ are

both Hamiltonian and pancyclic. On the other hand, 2-regular Hamiltonian graphs have only one cycle length. For $3 \leq k \leq \lceil n/2 \rceil - 1$, Jacobson and Lehel asked for the minimum size of the cycle spectrum of a k -regular n -vertex Hamiltonian graph, particularly when $k = 3$.

Let $s(G)$ be the size of the cycle spectrum of a graph G . At the SIAM Meeting on Discrete Mathematics in 2002, Jacobson announced that he, Gould, and Pfender had proved $s(G) \geq c_k n^{1/2}$ for k -regular graphs with n vertices. Others later independently obtained similar bounds, without seeking to optimize c_k . For an upper bound, Jacobson and Lehel constructed the 3-regular example below with only $n/6 + 3$ distinct cycle lengths (when $n \equiv 0 \pmod 6$ and $n > 6$), and they generalized it to the upper bound $\frac{n}{2} \frac{k-2}{k} + k$ for k -regular graphs.

Example 8.1.1. When $k = 3$ and 6 divides n , take $n/6$ disjoint copies of $K_{3,3}$ in a cyclic order, with vertex sets $V_1, \dots, V_{n/6}$. Remove one edge from each copy and replace it by an edge to the next copy to restore 3-regularity. A cycle of length different from 4 or 6 must visit each V_i , and in each V_i it uses 4 or 6 vertices. Hence the cycle lengths are 4, 6, and each even integer from $2n/3$ through n . For the generalization, use $K_{k,k}$ instead of $K_{3,3}$.

A related problem is the conjecture of Erdős [34] that $s(G) \geq \Omega(d^{\lfloor (g-1)/2 \rfloor})$ when G has girth g and average degree d . Erdős, Faudree, Rousseau, and Schelp [35] proved the conjecture for $g = 5$. Sudakov and Verstraëte [106] proved the full conjecture in a stronger form, obtaining $\frac{1}{8} (d^{\lfloor (g-1)/2 \rfloor})$ consecutive even integers in the cycle spectrum for graphs with fixed girth g and average degree $48(d+1)$. Gould, Haxell, and Scott [51] proved a similar result: for $c > 0$, there is a constant k_c such that for sufficiently large n , the cycle spectrum of every n -vertex graph G having minimum degree cn and longest even cycle length $2l$ contains all even integers from 4 up to $2l - k_c$ (see also [16]).

Prior arguments for lower bounds on $s(G)$ when G is regular and Hamiltonian used only the number of edges, m , not regularity. The complete bipartite graph $K_{n/2, n/2}$ shows that the coefficient c in a lower bound of the form $\sqrt{c(m-n)}$ cannot exceed 1. We give constructions for general m and n when m is above or below $n^2/4$; they are far from regular.

Example 8.1.2. For $m \leq n^2/4$, consider the graph G formed by replacing one edge of $K_{t,t}$ with a path having $n - 2t$ internal vertices; there are n vertices and $t^2 - 2t + n$ edges. The cycle spectrum of G consists of the $t - 1$ even numbers $\{4, \dots, 2t\}$ and the $t - 1$ numbers from $n - 2t + 4$ to n having the same parity as n . Letting m be the number of edges, this construction yields $s(G) \leq 2(t - 1) = 2\sqrt{m - n + 1}$.

Deleting edges cannot enlarge the cycle spectrum. Hence when we specify m as the number of edges, we can let m' be the next larger value such that $m' - n + 1$ is a square and apply the construction above for m' edges to obtain an upper bound. After discarding $m' - m$ edges, we obtain $s(G) \leq 2 \lceil \sqrt{m - n + 1} \rceil$.

The two parts of the spectrum remain separate when $n > 4(t - 1)$, which holds for $m < n^2/16 + n - 1$. When n is even and $m \geq n^2/16 + n$, the two parts overlap; in fact, $s(G) = \sqrt{m - n + 1}$ when $m = n^2/4$.

For $m > n^2/4$, consider the graph G formed by replacing one edge of K_t with a path having $n - t$ internal vertices; there are n vertices and $t(t - 3)/2 + n$ edges. The cycle spectrum of G consists of the $t - 2$ numbers $\{3, \dots, t\}$ and the $t - 2$ numbers $\{n - t + 3, \dots, n\}$. In terms of the number of edges, m , we have $t = \sqrt{2(m - n) + 9/4} + 3/2$. Thus $s(G) \leq 2(t - 2) < 2\sqrt{2}\sqrt{m - n + 1}$. The coefficient declines from $2\sqrt{2}$ toward $\sqrt{2}$ as m increases toward $\binom{n}{2}$. The worst case for the construction, in terms of $m - n$, occurs when $m = n^2/4 + 1$. Again we can interpolate for other values of m by inserting a ceiling function on the expression for t in terms of m and n , obtaining a general construction with $s(G) \leq 2 \left\lceil \sqrt{2(m - n + 1)} \right\rceil$.

Our main result is that $s(G) \geq \sqrt{\frac{4}{7}(m - n)}$ when G is an n -vertex Hamiltonian graph with m edges. A crucial tool is a lemma of Faudree, Flandrin, Jacobson, Lehel, and Schelp [42, Lemma 3]. We need a stronger version than their proof yields; we obtain this in Section 8.2. Section 8.3 applies it to $s(G)$.

8.2 Chords of a spanning path

A path with endpoints x and y is an x, y -*path*. A *chord* of a path (or cycle) P in a graph is an edge of the graph not in P whose endpoints are in P , and the *length* of the chord is the distance in P between its endpoints. Throughout this section, and in the hypotheses of its results, we let the graph G consist of the spanning x, y -path P plus q chords of the same length l , and we let r be the number of different lengths of x, y -paths in G . The vertices of P are v_1, \dots, v_n in order, with $v_1 = x$ and $v_n = y$. As defined above, the *length* of a chord $v_i v_j$ of P is $|j - i|$. Two chords $v_a v_c$ and $v_b v_d$ *overlap* if $a < b < c < d$. When v_a and v_b are vertices of P , we use $P[v_a, v_b]$ to denote the v_a, v_b -path contained in P .

Lemma 3 in [42] claims that in this setting, always $r \geq q/3 + 1$. However, the argument in [42] produces only $q/6 + 1$ path lengths in the following example.

Example 8.2.1. Let G be obtained from the 3-regular graph in Example 8.1.1 by deleting one edge that lies in no 6-cycle. Here $n = 6k = 2q$, and the endpoints x and y of the spanning path P are the endpoints of the deleted edge. The graph is bipartite, so all x, y -paths have odd length. All q chords have length 3, and they group into k sets of three overlapping chords. Each such set occupies six consecutive vertices along P , and an x, y -path visits exactly four or six vertices in each such group. Hence the lengths of x, y -paths are all odd numbers from $n - 1$ down to $2n/3 - 1$; there are $q/3 + 1$ of them. The argument in [42] discards half of these groups of three chords on six vertices and thus guarantees only $q/6 + 1$ path lengths.

Theorem 8.2.4 below will provide a lower bound on r that is always at least as large as $q/3 + 1$. The graph in Example 8.2.1 demonstrates sharpness.

Lemma 8.2.2. *If the chords other than the one nearest v_n pairwise overlap, then $r \geq q - 1$, with equality possible only when l is odd.*

Proof. Let the chords be e_1, \dots, e_q , indexed in order of the indices of their lower endpoints among v_1, \dots, v_n (e_q is nearest to v_n). Suppose first that e_q overlaps e_1 . For $2 \leq j \leq q$, let P_j be the v_1, v_n -path using the chords e_1 and e_j and no other chords. There is exactly one such path; if $e_1 = v_a v_c$ and $e_j = v_b v_d$, then P_j contains $P[v_1, v_a]$, $P[v_b, v_c]$, and $P[v_d, v_n]$. The length of P_j is $n + 1 - (b - a) - (d - c)$. As j increases, b and d increase, so P_2, \dots, P_q have distinct lengths. Furthermore, the v_1, v_n -path Q that contains e_1 and no other chord has length $n - l$. Since P_2, \dots, P_q have distinct lengths, $r \geq q - 1$. If l is even, then the length of Q has opposite parity from the lengths of P_2, \dots, P_q (since $(b - a) + (d - c)$ is even), and hence $r \geq q$.

Now suppose that e_q does not overlap e_1 . Let P' be the v_1, v_n -path using e_1 and e_q and no other chords; it has length $n + 1 - 2l$. Furthermore, this path is shorter than any of the paths P_2, \dots, P_{q-1} or Q constructed as in the previous case for $G - e_q$. \square

Proposition 8.2.3. *Let G be a graph with a distinguished spanning x, y -path P , let Q_1, \dots, Q_t be pairwise edge-disjoint subpaths of P , and let H_j be the subgraph of G induced by $V(Q_j)$. If H_j has r_j lengths of paths joining the endpoints of Q_j , then G has x, y -paths with at least $1 + \sum_{j=1}^t (r_j - 1)$ different lengths.*

Proof. Starting with a v_1, v_n -path that uses Q_1, \dots, Q_t (actually, this is P), one can iteratively shorten the path $\sum_{j=1}^t (r_j - 1)$ times. \square

Theorem 8.2.4. *If G is a graph consisting of a spanning path P with vertices v_1, \dots, v_n and q chords of length l , then the number r of v_1, v_n -paths in G is at least*

$$\max \left\{ \frac{q}{2} - \frac{n-1}{2l} + 1, \frac{q}{3} + 1 \right\}.$$

Moreover, if l is even, then there are at least $q/2 + 1$ such paths.

Proof. Choose chords e_1, \dots, e_k as follows. Let e_1 be the chord with lowest-indexed endpoint. Having chosen e_1, \dots, e_{j-1} , let e_j be the chord with lowest endpoint that overlaps none of e_1, \dots, e_{j-1} ; do this until no further chord can be added. Note that every chord of G coincides with or overlaps at least one of the chosen chords.

Let z_j and z'_j be the lower and upper endpoints of chord e_j , respectively. Let $z_0 = v_1$ and $z'_{k+1} = v_n$. For $0 \leq j \leq k$, let $Q_j = P[z_j, z'_{j+1}]$, let H_j be the subgraph of G induced by $V(Q_j)$, and let r_j be the number of z_j, z'_{j+1} -paths in H_j . Note that Q_j is a spanning path in

H_j . Since e_j is a chord in H_j and H_{j-1} for $1 \leq j \leq k$ and every other chord of G belongs to exactly one of these subgraphs, $\sum_{j=0}^k q_j = q + k$, where q_j is the number of chords in H_j . Each H_j has the form discussed in Lemma 8.2.2; hence $r_j \geq q_j - 1$ for $1 \leq j \leq k$. For H_0 we need the better value $r_0 = q_0 + 1 = 2$ (there is only one chord).

The odd-indexed graphs in H_0, \dots, H_k are pairwise disjoint, as are the even-indexed graphs. By applying Proposition 8.2.3 separately to the even and odd pieces and summing the resulting two inequalities, we obtain

$$2r \geq 2 + \sum_{j=0}^k (r_j - 1) \geq 2 + q_0 + \sum_{j=1}^k (q_j - 2) = 2 + q - k,$$

and thus $r \geq (q - k)/2 + 1$. Since e_1, \dots, e_k are pairwise non-overlapping, $n - 1 \geq kl$, so

$$r \geq \frac{q}{2} - \frac{n - 1}{2l} + 1. \quad (8.1)$$

Furthermore, the chords e_1, \dots, e_k by themselves yield $r \geq 1 + k$, and hence

$$r \geq \max \left\{ 1 + k, \frac{q - k}{2} + 1 \right\}.$$

Optimizing k yields $r \geq q/3 + 1$.

If l is even, then Lemma 8.2.2 yields $r_i \geq c_i$ for $1 \leq i \leq k$, and hence

$$2r \geq 2 + \sum_{i=0}^k (r_i - 1) \geq 2 + c_0 + \sum_{i=1}^k (c_i - 1) = 2 + q.$$

Thus $r \geq q/2 + 1$ in this case. □

Example 8.2.1 shows that the inequality $r \geq q/3 + 1$ is sharp; the example having $n - 3$ chords of length 3 also achieves equality. Similarly, a path with $q + 2$ vertices having q chords of length 2 shows that $r \geq q/2 + 1$ is best possible when l is even.

Corollary 8.2.5. *In the setting of Theorem 8.2.4, if $l \leq n/2$, then*

$$r \geq \frac{q}{3} \left(1 + \frac{l}{n} \right).$$

Proof. By using (8.1) to refine the second bound in Theorem 8.2.4, we have $r \geq \max\{f_1(q), f_2(q)\}$, where $f_1(x) = \frac{x}{2} - \frac{n}{2l} + 1$ and $f_2(x) = \frac{x}{3} + 1$.

Note that f_1 and f_2 are linear functions of x that intersect at a point $(x_0, y_0) = \left(\frac{3n}{l}, \frac{n}{l} + 1\right)$. Since $f_1(0) < 0 < f_2(0)$, the line $y = \frac{x}{3} \left(1 + \frac{l}{n}\right)$ that passes through $(0, 0)$ and (x_0, y_0) provides a uniform lower bound on $\max\{f_1(x), f_2(x)\}$. □

8.3 Cycle lengths in Hamiltonian graphs

In this section, G is a graph G with n vertices, m edges, and a distinguished Hamiltonian cycle C . A *chord of length l* is an edge uv in $E(G) - E(C)$ such that $d_C(u, v) = l$. The length of any chord is at least 2 and at most $\lfloor n/2 \rfloor$. Let the *normalized length* of a chord uv be $\frac{d_C(u, v)}{n/2}$. Two chords uv and xy *cross* if their endpoints are ordered u, x, v, y along C .

The next two lemmas give lower bounds on $s(G)$ (the size of the cycle spectrum of G). The first is strong when the average length of the chords is large, and the second is strong when the average length is small. The two bounds together imply our main result.

Let $\omega(G)$ and $\alpha(G)$ denote the *clique number* and *independence number* of a graph G ; these are the maximum sizes of sets of pairwise adjacent or pairwise nonadjacent vertices, respectively. A graph H is *perfect* if for every induced subgraph H' the vertices can be partitioned into $\omega(H')$ independent sets (and hence $\omega(H)\alpha(H) \geq |V(H)|$).

Lemma 8.3.1. *Let G be an n -vertex graph having m edges and a Hamiltonian cycle C . If the average normalized length of the chords of C is β , then $s(G) \geq \sqrt{\beta(m - n)}$.*

Proof. We seek a large set of chords in one of two special configurations. If C has q pairwise noncrossing chords, then $s(G) \geq q + 1$ (starting with C , we can iteratively obtain a shorter cycle q times). If C has q pairwise crossing chords, then $s(G) \geq q - 1$ (starting with a short cycle using two “closest” among the q crossing chords, we can iteratively obtain a longer cycle $q - 2$ times by replacing one of them).

To obtain a large set of pairwise noncrossing chords or a large set of pairwise crossing chords, we seek a large set of chords crossing a single diameter. With C drawn on a circle, let S be the vertex set of a path along C , with $\bar{S} = V(G) - S$. Let p be the number of chords of C having endpoints in S and \bar{S} . Let H be the graph whose vertex set is this set of p chords, with vertices being adjacent when the chords cross. It is well known ([90], for example) that a graph generated in this way is a perfect graph (specifically, a “permutation graph”). Since H is perfect, $\omega(H)\alpha(H) \geq p$. Hence $\omega(H) \geq \sqrt{p} + 1$ or $\alpha(H) \geq \sqrt{p} - 1$.

Choose a random set S of $\lfloor n/2 \rfloor$ consecutive vertices along C , with the n such sets being equally likely. The probability that a chord of length l has exactly one endpoint in S is exactly $\frac{2l}{n}$, which equals the normalized length of the chord. The expected number of chords with one endpoint in S is thus the sum of the normalized lengths, which equals $\beta(m - n)$. For some choice of S , there are at least $\beta(m - n)$ such chords. As argued above, there are thus at least $\sqrt{\beta(m - n)}$ distinct cycle lengths in G . \square

Lemma 8.3.2. *Let G be an n -vertex graph having m edges and a Hamiltonian cycle C . If the average normalized length of the chords of C is β , then*

$$s(G) \geq \sqrt{\frac{2}{3} \left(1 - \frac{\beta}{4}\right) (m - n)}.$$

Proof. Every chord of length l completes with C a cycle of length $l+1$ and a cycle of length $n-l+1$; these are both shorter than C , and they have the same length if and only if $l = n/2$. Letting t be the number of different lengths of chords of C , we find that C together with such cycles yields $s(G) \geq 2t$.

When t is small, many chords have equal length. Let $w(l) = (1 + l/n)/3$. To make use of the extra factor $w(l)$ in Corollary 8.2.5, assign each chord of C with length l the weight $w(l)$. For an edge uv of C chosen uniformly at random, let $P = C - uv$; we treat P as a distinguished Hamiltonian path of G . A chord of C is also a chord of P ; let the P -length of a chord xy of C be $d_P(x, y)$. For a chord of length l , the P -length is equal to l with probability $1 - l/n$. Let W be the expectation (over the choice of uv) of the total weight of all chords whose length and P -length coincide. Letting a_l be the number of chords with length l , the expected number of chords of length l that contribute to W is $a_l(1 - l/n)$. Thus

$$\begin{aligned} W &= \sum_{l \geq 2} \frac{1}{3} \left(1 + \frac{l}{n}\right) a_l \left(1 - \frac{l}{n}\right) = \frac{1}{3} \sum_{l \geq 2} a_l \left(1 - \frac{l^2}{n^2}\right) \\ &= \frac{1}{3} \left(\sum_{l \geq 2} a_l - \frac{1}{4} \sum_{l \geq 2} a_l \left(\frac{l}{n/2}\right)^2 \right) \geq \frac{1}{3} \left((m - n) - \frac{1}{4} \sum_{l \geq 2} a_l \frac{l}{n/2} \right) \\ &= \frac{1}{3} \left((m - n) - \frac{1}{4} \beta(m - n) \right) = \frac{1}{3} (m - n) \left(1 - \frac{\beta}{4}\right). \end{aligned}$$

For some choice of uv along C , the actual total weight of the chords whose length and P -length coincide is at least W . With t different chord lengths, some particular length contributes at least W/t to this total. Let l be this length. We have $a_l \geq W/(tw(l))$, and hence by Corollary 8.2.5 the chords of this length contribute at least W/t lengths of cycles.

We now have

$$s(G) \geq \max \left\{ 2t, \frac{(m - n)}{3t} \left(1 - \frac{\beta}{4}\right) \right\} \geq \sqrt{\frac{2}{3} \left(1 - \frac{\beta}{4}\right) (m - n)},$$

where the final inequality chooses t to minimize the maximum. □

Theorem 8.3.3. *If G is a n -vertex Hamiltonian graph with m edges, then $s(G) \geq \sqrt{\frac{4}{7}(m - n)}$. Furthermore, there is such a graph G with $s(G) \leq 2 \left\lceil \sqrt{2(m - n + 1)} \right\rceil$, and when $m = n^2/4$ there is such a graph with $s(G) = \sqrt{m - n + 1}$.*

Proof. By Lemmas 8.3.1 and 8.3.2, $s(G) \geq [(m - n) \max\{\beta, \frac{2}{3}(1 - \frac{\beta}{4})\}]^{1/2}$. Choosing $\beta = 4/7$ minimizes the larger lower bound. Example 8.1.2 provides the construction. □

References

- [1] A.A. Ageev. A triangle-free circle graph with chromatic number 5. *Discrete Math.*, 152(1-3):295–298, 1996.
- [2] N. Alon. Combinatorial Nullstellensatz. *Combin. Probab. Comput.*, 8(1-2):7–29, 1999. Recent trends in combinatorics (Mátraháza, 1995).
- [3] N. Alon, G. Brightwell, H.A. Kierstead, A.V. Kostochka, and P. Winkler. Dominating sets in k -majority tournaments. *J. Combin. Theory Ser. B*, 96(3):374–387, 2006.
- [4] N. Alon, G. Ding, B. Oporowski, and D. Vertigan. Partitioning into graphs with only small components. *Journal of Combinatorial Theory, Series B*, 87(2):231–243, 2003.
- [5] N. Alon, J. Grytczuk, M. Hałuszczak, and O. Riordan. Nonrepetitive colorings of graphs. *Random Structures Algorithms*, 21(3-4):336–346, 2002. Random structures and algorithms (Poznan, 2001).
- [6] D. Bauer and E.F. Schmeichel. Hamiltonian degree conditions which imply a graph is pancyclic. *Journal of Combinatorial Theory, Series B*, 48:111–116, 1990.
- [7] P. Beame, E. Blais, and D.T. Huynh-Ngoc. Longest common subsequences in sets of permutations. *Arxiv preprint arXiv:0904.1615*, 2009.
- [8] P. Beame and D.T. Huynh-Ngoc. On the value of multiple read/write streams for approximating frequency moments. In *IEEE 49th Annual IEEE Symposium on Foundations of Computer Science, 2008. FOCS'08*, pages 499–508, 2008.
- [9] J. Beck. On size Ramsey number of paths, trees, and circuits I. *J. Graph Theory*, 7(1):115–129, 1983.
- [10] J. Beck. On size Ramsey number of paths, trees and circuits. II. In *Mathematics of Ramsey Theory*, volume 5 of *Algorithms and Combin.*, pages 34–45. Springer, 1990.
- [11] J. Beck. Achievement games and the probabilistic method. In *Combinatorics, Paul Erdős is Eighty*, volume 1 of *Bolyai Soc. Math. Stud.*, pages 51–78. 1993.
- [12] L.W. Beineke and A.J. Schwenk. On a bipartite form of the Ramsey problem. In *Proceedings of the Fifth British Combinatorial Conference (Univ. Aberdeen, Aberdeen, 1975)*, pages 17–22. Congressus Numerantium, No. XV, Winnipeg, Man., 1976. Utilitas Math.

- [13] B. Bollobás. *The art of mathematics*, page 186. Cambridge University Press, New York, 2006. Coffee time in Memphis.
- [14] B. Bollobás and I. Leader. Sums in the grid. *Discrete Math.*, 162(1-3):31–48, 1996.
- [15] B. Bollobás, A. Saito, and N.C. Wormald. Regular factors of regular graphs. *Journal of Graph Theory*, 9(1):97–103, 2006.
- [16] B. Bollobás and A. Thomason. Weakly pancyclic graphs. *Journal of Combinatorial Theory, Series B*, 77:121–137, 1999.
- [17] J.A. Bondy. Pancyclic graphs. *Journal of Combinatorial Theory, Series B*, 11:80–84, 1971.
- [18] J.A. Bondy. Longest paths and cycles in graphs of high degree. Research Report CORR 80-16, Univ. Waterloo, 1980.
- [19] D.P. Bunde, K.G. Milans, D.B. West, and H. Wu. Parity and strong parity edge-coloring of graphs. In *Proceedings of the Thirty-Eighth Southeastern International Conference on Combinatorics, Graph Theory and Computing*, volume 187, pages 193–213, 2007.
- [20] D.P. Bunde, K.G. Milans, D.B. West, and H. Wu. Optimal strong parity edge-coloring of complete graphs. *Combinatorica*, 28(6):625–632, 2008.
- [21] S.A. Burr, P. Erdős, and L. Lovász. On graphs of Ramsey type. *Ars Combinatoria*, 1(1):167–190, 1976.
- [22] J. Butterfield, T. Grauman, B. Kinnersley, K.G. Milans, C. Stocker, and D.B. West. On-line Ramsey Theory for Bounded Degree Graphs. In preparation.
- [23] D. Cenzer and J.B. Remmel. Π_1^0 classes in mathematics. In Y.L. Ershov, S.S. Goncharov, A. Nerode, and J.B. Remmel, editors, *Handbook of Recursive Mathematics*, volume 139 of *Studies in Logic and the Foundations of Mathematics*, chapter 13, pages 623–681. Elsevier, 1998.
- [24] C. Chvátal, V. Rödl, E. Szemerédi, and W.T. Trotter, Jr. The Ramsey number of a graph with bounded maximum degree. *J. Combin. Theory Ser. B*, 34(3):239–243, 1983.
- [25] V. Chvátal. On Hamilton’s ideals. *Journal of Combinatorial Theory, Series B*, 12(2):163–168, 1972.
- [26] D. Conlon. A new upper bound for the bipartite Ramsey problem. *J. Graph Theory*, 58(4):351–356, 2008.
- [27] S.A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the third annual ACM symposium on Theory of computing*, pages 151–158. ACM, 1971.

- [28] G.A. Dirac. Some theorems on abstract graphs. *Proc. London Math. Soc. (3)*, 2:69–81, 1952.
- [29] J. Donadelli, P.E. Haxell, and Y. Kohayakawa. A note on the size-Ramsey number of long subdivisions of graphs. *RAIRO-Theoretical Informatics and Applications*, 39(1):191–206, 2005.
- [30] R.G. Downey, N. Greenberg, C.G. Jockusch Jr., and K.G. Milans. Binary subtrees with few labeled paths. Submitted, 2010+.
- [31] R.G. Downey and D. Hirschfeldt. *Algorithmic randomness and complexity*. Springer-Verlag, 2008. In preparation.
- [32] S. Eliahou and M. Kervaire. Sumsets in vector spaces over finite fields. *J. Number Theory*, 71(1):12–39, 1998.
- [33] S. Eliahou and M. Kervaire. Old and new formulas for the Hopf-Stiefel and related functions. *Expo. Math.*, 23(2):127–145, 2005.
- [34] P. Erdős. On some of my favourite problems in various branches of combinatorics. *Matematiche (Catania)*, 12:163–168, 1972.
- [35] P. Erdős, R. Faudree, C. Rousseau, and R. Schelp. The number of cycle lengths in graphs of given minimum degree and girth. *Discrete Mathematics*, 200:55–60, 1999.
- [36] P. Erdős. On a problem in graph theory. *Math. Gaz.*, 47:220–223, 1963.
- [37] P. Erdős, R.J. Faudree, C.C. Rousseau, and R.H. Schelp. The size Ramsey number. *Period. Math. Hungar.*, 9(1-2):145–161, 1978.
- [38] P. Erdős and L. Moser. On the representation of directed graphs as unions of orderings. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 9:125–132, 1964.
- [39] P. Erdős and H. Sachs. Reguläre Graphen gegebener Taillenweite mit minimaler Knotenzahl. *Wiss. Z. Martin-Luther-Univ. Halle-Wittenberg Math.-Natur. Reihe*, 12:251–257, 1963.
- [40] P. Erdős and G. Szekeres. A combinatorial problem in geometry. *Compositio Math.*, 2:463–470, 1935.
- [41] G. Fan. New sufficient conditions for cycles in graphs. *Journal of Combinatorial Theory, Series B*, 37(3):221–227, 1984.
- [42] R.J. Faudree, E. Flandrin, M.S. Jacobson, J. Lehel, and R.H. Schelp. Even cycles in graphs with many odd cycles.
- [43] M. Ferrara, A. Harris, and M. Jacobson. Cycle lengths in Hamiltonian graphs with a pair of vertices having large degree sum. Preprint, 2008.
- [44] M. Ferrara and C. Tennenhouse, 2009. Personal communication.

- [45] S.R. Finch. *Mathematical Constants*, chapter 2.30.1 Powers of $3/2$ modulo one, pages 194–199. Number 94 in *Encyclopedia of Mathematics and its Applications*. Cambridge, 2003.
- [46] J. Folkman. Graphs with monochromatic complete subgraphs in every edge coloring. *SIAM Journal on Applied Mathematics*, 18(1):19–24, 1970.
- [47] E. Friedgut, Y. Kohayakawa, V. Rödl, A. Ruciński, and P. Tetali. Ramsey games against a one-armed bandit. *Combinatorics, Probability and Computing*, 12:515–545, 2003.
- [48] J. Friedman and N. Pippenger. Expanding graphs contain all small trees. *Combinatorica*, 7(1):71–76, 1987.
- [49] Z. Füredi, C.G. Jockusch Jr., and L.A. Rubel. Difference sets and inverting the difference operator. *Combinatorica*, 16(1):87–106, 1996.
- [50] R. Goldblatt. The McKinsey axiom is not canonical. *The Journal of Symbolic Logic*, 56(2):554–562, 1991.
- [51] R.J. Gould, P.E. Haxell, and A.D. Scott. A note on cycle lengths in graphs. *Graphs and Combinatorics*, 18(3):491–498, 2002.
- [52] N. Greenberg and J.S. Miller. Diagonally non-recursive functions and effective hausdorff dimension. Submitted.
- [53] J.A. Grytczuk, M. Hałuszczak, and H.A. Kierstead. On-line Ramsey theory. *Electron. J. Combin.*, 11(1):Research Paper 57, 10 pp. (electronic), 2004. Paper number changed by the publisher from 60 to 57.
- [54] J.A. Grytczuk, H.A. Kierstead, and P. Prałat. On-line Ramsey numbers for paths and stars. *Discrete Math. Theor. Comp. Sci.*, 10:63–74, 2008.
- [55] A. Gyárfás. On the chromatic number of multiple interval graphs and overlap graphs. *Discrete Math.*, 55(2):161–166, 1985.
- [56] A. Gyárfás. Corrigendum: “On the chromatic number of multiple interval graphs and overlap graphs” [*Discrete Math.* **55** (1985), no. 2, 161–166; MR0798532 (86k:05052)]. *Discrete Math.*, 62(3):333, 1986.
- [57] I. Havel and J. Morávek. B -valuations of graphs. *Czechoslovak Math. J.*, 22(97):338–351, 1972.
- [58] P.E. Haxell, Y. Kohayakawa, and T. Łuczak. The induced size-Ramsey number of cycles. *Combin. Probab. Comput.*, 4(3):217–239, 1995.
- [59] H. Hopf. Ein topologischer Beitrag zur reellen Algebra. *Comment. Math. Helv.*, 13:219–239, 1941.

- [60] T. Jiang, K.G. Milans, and D.B. West. Degree Ramsey number of cycles is bounded. In preparation.
- [61] C.G. Jockusch Jr. Degrees of functions with no fixed points. In J.E. Fenstad, I. T. Frolov, and R. Hilpinen, editors, *Logic, Philosophy, and Methodology of Science VIII*, pages 191–201. North-Holland, 1989.
- [62] I.A. Karapetyan. Chordal graphs. *Mat. Voprosy Kibernet. Vychisl. Tekhn.*, (14):6–10, 1985.
- [63] G. Károlyi. A note on the Hopf-Stiefel function. *European J. Combin.*, 27(7):1135–1137, 2006.
- [64] P. Keevash and B. Sudakov. On a hypergraph Turán problem of Frankl. *Combinatorica*, 25(6):673–706, 2005.
- [65] H.A. Kierstead. An effective version of Dilworth’s theorem. *Transactions of the American Mathematical Society*, 268(1):63–77, 1981.
- [66] H.A. Kierstead and G. Konjevod. Coloring number and on-line Ramsey theory for graphs and hypergraphs. *Combinatorica*, 29:49–64, 2009.
- [67] B. Kinnersley, K.G. Milans, and D.B. West. Degree Ramsey number of double-stars and the 4-cycle. In preparation.
- [68] A. Kostochka. Coloring intersection graphs of geometric figures with a given clique number. In *Towards a theory of geometric graphs*, volume 342 of *Contemp. Math.*, pages 127–138. Amer. Math. Soc., Providence, RI, 2004.
- [69] A. Kostochka and J. Kratochvíl. Covering and coloring polygon-circle graphs. *Discrete Math.*, 163(1-3):299–305, 1997.
- [70] A.V. Kostochka. On upper bounds for the chromatic numbers of graphs. *Trudy Instituta Matematiki*, 10:204–226, 1988. (In Russian).
- [71] A.V. Kostochka and K.G. Milans. The chromatic number of K_4 -free circle graphs. In preparation.
- [72] D.W. Krumme, K.N. Venkataraman, and G. Cybenko. Hypercube embedding is NP-complete. In M.T. Heath, editor, *Proc. 1st Conf. Hypercube Multiprocessors*, pages 148–157. Soc. Ind. Appl. Math., 1986.
- [73] A. Kurek and A. Ruciński. Two variants of the size Ramsey number. *Discuss. Math. Graph Theory*, 25(1-2):141–149, 2005.
- [74] A. Kučera. Measure, Π_1^0 -classes and complete extensions of PA. In *Recursion Theory Week (Oberwolfach, 1984)*, volume 1141 of *Lecture Notes in Math.*, pages 245–259. Springer, 1985.

- [75] D.C. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21:608–610, 1953.
- [76] Y.T. Medvedev. Degrees of difficulty of mass problems. In *Doklady Akademii Nauk SSSR (N.S.)*, volume 104, pages 501–504, 1955.
- [77] K.G. Milans, D. Rautenbach, F. Regen, and D.B. West. Cycle spectra of Hamiltonian graphs. Submitted, 2010+.
- [78] K.G. Milans, D. Schreiber, and D.B. West. Acyclic sets in k -majority tournaments. In preparation.
- [79] J.R. Mileti. The canonical Ramsey theorem and computability theory. *Transactions of the American Mathematical Society*, 360(3):1309–1340, 2008.
- [80] J. Mitas and K. Reuter. Subgraphs of hypercubes and subdiagrams of Boolean lattices. *Order*, 13(1):41–64, 1996.
- [81] A. Montalbán. On the equimorphism types of linear orderings. *The Bulletin of Symbolic Logic*, 13(1):71–99, 2007.
- [82] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, Cambridge, New York, 1995.
- [83] A.A. Mučnik. On strong and weak reducibility of algorithmic problems. *Sibirsk. Mat. Ž.*, 4:1328–1341, 1963.
- [84] J. Mycielski. Sur le coloriage des graphes. *Colloq. Math.*, 3:161–162, 1955.
- [85] J. Nešetřil and V. Rödl. Type theory of partition properties of graphs. In *Recent advances in graph theory*, Proc. Second Czechoslovak Sympos (Prague 1974), pages 405–412, 1975.
- [86] A.M. Odlyzko and H.S. Wilf. Functional iteration and the Josephus problem. *Glasgow Math. J.*, 33(2):235–240, 1991.
- [87] O. Ore. Note on Hamilton circuits. *American Mathematical Monthly*, 67:55, 1960.
- [88] J. Petersen. Die Theorie der regulären graphs. *Acta Math.*, 15(1):193–220, 1891.
- [89] A. Plagne. Additive number theory sheds extra light on the Hopf-Stiefel \circ function. *Enseign. Math. (2)*, 49(1-2):109–116, 2003.
- [90] A. Pnueli, A. Lempel, and S. Even. Transitive orientation of graphs and identification of permutation graphs. *Canad. J. Math.*, 23:160–175, 1971.
- [91] G. Pólya and G. Szegő. *Problems and Theorems in Analysis I*. Springer-Verlag, Berlin, Heidelberg, New York, 1987.

- [92] P. Prałat. A note on small on-line Ramsey numbers for paths and their generalization. *Australasian Journal of Combinatorics*, 40:27–36, 2008.
- [93] P. Prałat. $\overline{\mathcal{R}}(3, 4) = 17$. *Electronic Journal of Combinatorics*, 15(R67):13pp, 2008.
- [94] P. Prałat. A note on off-diagonal small on-line Ramsey numbers for paths. *Ars Combinatoria*, 2010+.
- [95] V. Rödl and E. Szemerédi. On size Ramsey numbers of graphs with bounded degree. *Combinatorica*, 20(2):257–262, 2000.
- [96] V. Rödl and E. Szemerédi. On size Ramsey numbers of graphs with bounded degree. *Combinatorica*, 20(2):257–262, 2000.
- [97] D. Rolnick, 2010. Personal communication.
- [98] E.F. Schmeichel and S.L. Hakimi. A cycle structure theorem for Hamiltonian graphs. *Journal of Combinatorial Theory, Series B*, 45:99–107, 1988.
- [99] E.F. Schmeichel and J. Mitchem. Bipartite graphs with cycles of all even lengths. *J. Graph Theory*, 6:429–439, 1982.
- [100] H.S. Shapiro. The embedding of graphs in cubes and the design of sequential relay circuits. *Bell Telephone Laboratories Memorandum*, July 1953.
- [101] S.G. Simpson. Mass problems and randomness. *Bull. Symbolic Logic*, 11(1):1–27, 2005.
- [102] M. Sipser. *Introduction to the Theory of Computation*. Thomson Course Technology, Second edition, 2006.
- [103] R.I. Soare. *Recursively Enumerable Sets and Degrees*. Springer-Verlag, 1940.
- [104] R. Stearns. The voting problem. *Amer. Math. Monthly*, 66:761–763, 1959.
- [105] E. Stiefel. Über Richtungsfelder in den projektiven Räumen und einen Satz aus der reellen Algebra. *Comment. Math. Helv.*, 13:201–218, 1941. (In German).
- [106] B. Sudakov and J. Verstraëte. Cycle lengths in sparse graphs. *Combinatorica*, 28(3):357–372, 2008.
- [107] E. Szemerédi. Regular partitions of graphs. In *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*, volume 260 of *Colloq. Internat. CNRS*, pages 399–401. CNRS, Paris, 1978.
- [108] S.A. Terwijn. The Medvedev lattice of computably closed sets. *Archive for Mathematical Logic*, 45(2):179–190, 2006.
- [109] C. Thomassen. Two-coloring the edges of a cubic graph such that each monochromatic component is a path of length at most 5. *J. Combin. Theory Ser. B*, 75(1):100–109, 1999.

- [110] V.G. Vizing. On an estimate of the chromatic class of a p-graph. *Diskret. Analiz*, 3(7):25–30, 1964. In Russian.
- [111] A. Wagner and D.G. Corneil. Embedding trees in a hypercube is NP-complete. *SIAM J. Comput.*, 19(3):570–590, 1990.
- [112] Z. Xuding. Chromatic Ramsey numbers. *Discrete Mathematics*, 190:215–222, 1998.
- [113] S. Yuzvinsky. Orthogonal pairings of Euclidean spaces. *Michigan Math. J.*, 28(2):131–145, 1981.