On certain problems in extremal and additive combinatorics

Dissertation
zur Erlagung des Grades eines Doktors der Naturwissenschaften

am Fachbereich Mathematik und Informatik
der Freien Universität Berlin

Freie Universität Berlin

vorgelegt von
Codruţ-Miron Grosu

Berlin, 2015

Erstgutachter: Prof. Tibor Szabó, PhD (Betreuer)

Zweitgutachter: Prof. Oleg Pikhurko

Tag der Disputation: 09.05.2016

Contents

Preface

This thesis consists of five chapters.

The first chapter serves as an introduction, presenting the four problems studied in this thesis, and the results obtained. Each subsequent chapter then treats a separate problem.

The second chapter is about the existence of partial isomorphisms (i.e. bijective maps that preserve only a finite number of algebraic relations) between subsets of $\mathbb{F}_p$ and subsets of $\mathbb{C}$. The main result states that for any sufficiently small subset of $\mathbb{F}_p$ one can find a subset of $\mathbb{C}$ that algebraically behaves similarly. This has several applications, most importantly, it is shown that for small subsets of $\mathbb{F}_p$, the Szemerédi-Trotter theorem holds with optimal exponent $\frac{4}{3}$. Another application is towards an old question of A. Rényi concerning the number of terms of the square of a polynomial. The content of this chapter was published in [55].

The third chapter is about Turán densities of hypergraphs. The study of Turán densities was initiated by P. Turán in the 1940s, and has been an active area of research ever since. My main result proves that the set of Turán densities forms a non-trivial semigroup. Using this property several facts about Turán densities (which were previously proved by others) can be deduced in a streamlined fashion. The results presented in this chapter appear in [56].

The fourth chapter is about the Graceful Tree Conjecture, a generalization due to A. Rosa of the Ringel-Kotzig conjecture, with important ramifications in the field of graph decompositions. The main theorem is a proof of an approximative version of the conjecture for trees of bounded degree. The results in this chapter are joint work with Anna Adamaszek, Michał Adamaszek, Peter Allen and Jan Hladký, and appear in [1].

The fifth chapter is about the *Towers of Hanoi* problem with p pegs, a well-known variation on the classic puzzle of É. Lucas with 3 pegs. The question of determining the minimum number of moves needed to solve the puzzle has remained open for decades. Only recently a complete solution of the case of 4 pegs has been obtained by T. Bousch. The main result of the chapter, which appears in [57], is an asymptotic improvement on the best known lower bound for the minimum number of moves needed when $p \geq 5$.

Acknowledgements

First and foremost, I would like to express my thanks to Professor Tibor Szabó, my advisor, for supporting and encouraging me throughout my PhD. In addition, I would like to thank my coauthors Anna Adamaszek, Michał Adamaszek, Peter Allen and Jan Hladký for sharing their ideas with me and allowing the contents of our joint work to appear in this thesis.

I am grateful to Dennis, Roman, Anita, Lothar, Yury and Tuan for making a stimulating research group and for the many happy moments spent together. I fondly remember our yearly group workshops where we solved problems in a relaxed atmosphere.

This work would not have been possible without the support of the graduate school *Methods for Discrete Structures*, financially and otherwise. In particular I am thankful to Dorothea Kiefer for always answering my questions and helping me with various bureaucratic issues. I am also indebted to the *Berlin Mathematical School* for providing funding for my first two years in Berlin and for making me feel at home here.

I am thankful to all my friends for making my stay in Berlin more enjoyable. Albert, Gerrit, Claudius, Isabel, Han-Cheng, Arne, Alexandra, Julie, Agnes, Adrián, Atul, Ahmad, Giovanni, Emre: we had a great time together! Special thanks also go to Albert, for helping me with the German translation of my Preface.

Finally, I would like to thank my family, and especially my mom, for their everlasting support.

Notation

The following notation is used throughout the thesis.

$[n]$	the set $\{1, 2, \ldots, n\}$ in Chapters $2, 3, 4$; the set $\{0, 1, \ldots, n-1\}$ in Chapter 5
Δ_n	the $(n-1)$-dimensional simplex
$\mathbb{Z}_N$	the additive group of integers modulo N
$\mathbb{F}_p$	the finite field with p elements
$\mathbb{Z}_{(p)}$	the localization of $\mathbb{Z}$ at (p)
$A + B$	the sumset $\{a + b : a \in A, b \in B\}$
$v(G)$	the number of vertices of the (hyper)graph G
$e(G)$	the number of edges of the (hyper)graph G
$G[U]$	the induced subgraph of G on vertex set $U \subseteq V(G)$
K_n	the complete graph on n vertices
K_n^r	the complete r-uniform hypergraph on n vertices
K_n^{r-}	the complete r-uniform hypergraph on n vertices minus an edge

Chapter 1

Introduction

Combinatorics is a branch of mathematics devoted to the study of finite structures. Initially combinatorial problems were studied as a byproduct of number theory, algebra and topology. However, in the 20th century, partially due to the advent of computer science, these isolated problems were successfully unified in a general theory. Nowadays combinatorics has grown to encompass many smaller subfields, such as coding theory, algebraic and enumerative combinatorics, extremal graph theory, additive combinatorics, and many others.

Extremal graph theory is one of the oldest and more established subfields of combinatorics. The main topic of study is the maximum (or minimum) possible number of edges of a graph, subject to certain restrictions. One of the motivating questions in this area was posed by Turán in 1941. Suppose $k \geq 2$ is fixed and G is a graph with n vertices. What is the maximum possible number of edges of G, such that G does not contain a complete subgraph on k vertices? While this problem was solved by Turán in his original paper, the more general setting where we forbid an arbitrary subgraph H turned out to be more difficult to study and is not completely solved. Moreover, much less is known if one instead tries to maximize the number of edges of a uniform hypergraph with a forbidden subgraph. A less ambitious goal would be to find only the density of the extremal (hyper)graphs. Chapter 3 is devoted to the study of possible densities of extremal hypergraphs.

Graph decomposition problems can also be roughly classified as being part of extremal graph theory. In this case, given two graphs H and G, one tries to maximize the number of vertex (or edge) disjoint copies of H that can be embedded into G. In 1963, Ringel proposed the following conjecture: given a tree T with $n+1$ vertices, any complete graph on $2n+1$ vertices can be decomposed into $2n+1$ edge-disjoint copies of T. This problem was later strengthened by Rosa into a labelling problem for trees: given any n-vertex tree T, he conjectured that there exists a labelling of the vertices of T with distinct numbers from 1 to n such that the absolute value differences of the labels over the edges are pairwise distinct. Again, this can be thought of as an extremal problem, where we try to minimize the number of distinct vertex labels needed in a labelling of T such that the induced edge labels are all distinct. Chapter 4 shows that in this form it is possible to prove an approximative version of Rosa's conjecture: namely, given only εn additional labels, a labelling of T with the required properties exists.

Another more recent, but no less important, subfield is additive combinatorics. It has witnessed considerable growth since the proof of Green and Tao of the existence of arbitrarily long arithmetic progressions in the primes. The main theme of the field involves various estimates on the growth of finite subsets of numbers under elementary arithmetic operations.

Many times the choice of the underlying field in which addition and multiplication is carried over affects the statement of the theorem or the method of proof. For example, the Szemerédi-Trotter theorem in $\mathbb{C}$ provides stronger asymptotic bounds that the currently best known version of the theorem for $\mathbb{F}_p$. Hence it would be interesting to have transference theorems, that would give sufficient conditions allowing a change in the base field (from $\mathbb{C}$ to $\mathbb{F}_p$, or from $\mathbb{F}_p$ to $\mathbb{C}$, for example). Chapter 2 presents a result which allows the transfer of many theorems that hold for complex numbers to finite fields.

Finally, Chapter 5 presents progress on the Towers of Hanoi puzzle with more pegs. The usual setup of the Towers of Hanoi problem is 3 pegs and N disks arranged on the first peg in increasing order according to size. The goal is to move the N disks to another peg in as few moves as possible, such that larger disks are never placed on top of smaller ones. The problem can be generalized by allowing p pegs to be used instead of 3. Representing the possible states of the game as vertices of a graph, with edges between vertices corresponding to disk moves, it is easily seen that the Towers of Hanoi puzzle can be reformulated as a question about shortest distances in graphs. While the solution to the case of 3 pegs is folklore, the general case is still open. The purpose of Chapter 5 is to improve the best known lower bound for the minimum number of moves needed to solve the puzzle with p pegs.

We describe below in more detail the background, the problems studied and the results of this thesis.

1.1 Freiman ring isomorphisms

1.1.1 Definitions and previous results

Let $k \geq 1$ be an integer, Z and W two abelian groups and $A \subseteq Z, B \subseteq W$ finite non-empty subsets. We call Z *the ambient group* of A, and W *the ambient group* of B.

Definition 1. *A map $\phi : A \to B$ is a* Freiman isomorphism of order k*, or simply F_k-*isomorphism*, if for any $a_1, \ldots, a_{2k} \in A$ we have*

$$a_1 + \ldots + a_k = a_{k+1} + \ldots + a_{2k}$$

if and only if

$$\phi(a_1) + \ldots + \phi(a_k) = \phi(a_{k+1}) + \ldots + \phi(a_{2k}).$$

Example: Any (group) isomorphism $\phi : Z \to W$ is an F_k-isomorphism between A and $\phi(A)$, for all $k \geq 1$ and $A \subseteq Z$. In particular, if $(q, N) = 1$ then the map $\phi : \mathbb{Z}_N \to \mathbb{Z}_N$ given by $x \to qx$ is an F_k-isomorphism for all $k \geq 1$.

Note that a map $\phi : A \to B$ is an F_1-isomorphism if and only if it is a bijection. Furthermore, if a map is an F_{k+1}-isomorphism then it is also an F_k-isomorphism. To see this, choose $a_{k+1} = a_{2k+2}$ in Definition 1 to verify the condition for an F_k-isomorphism.

Translation does not affect Freiman isomorphisms. If $\phi : A \to B$ is an F_k-isomorphism and $u \in Z, v \in W$ arbitrary then the map $\psi : A + u \to B + v$ given by $a + u \to \phi(a) + v$ is also an F_k-isomorphism.

Freiman isomorphisms were first introduced by Freiman ([47]) in the proof of his eponymous theorem, and have since proved to be a useful tool in translating statements about the additive structure of sets of integers to $\mathbb{Z}_N$, and back. For example, it can be shown that any finite subset of a torsion-free group is F_k-isomorphic to a subset of $\mathbb{Z}_N$, for any large enough N. In the other direction it is well-known that small subsets of $\mathbb{Z}_p$, with p prime, are Freiman isomorphic to subsets of $\mathbb{Z}$.

Theorem 2 ([10]). *Let $A \subseteq \mathbb{Z}_p$, where p is a prime. If $|A| \leq \log_{2k} p$, then there exists a set of integers $A' \subset \mathbb{Z}$ such that the canonical homomorphism $\mathbb{Z} \to \mathbb{Z}_p$ induces an F_k-isomorphism of A' onto A.*

The theorem was further extended to subsets $A \subseteq \mathbb{Z}_p$ with $|A| \leq \log_{2k} p + \log_{2k} \log_{2k} p$ (see [10]). Nevertheless, the bound in Theorem 2 is almost tight, as in [10] it is shown the existence of a set $A \subset \mathbb{Z}_p$ of cardinality at most $2\log_k p + 1$ which is not F_k-isomorphic to any set of integers. Furthermore, we have the following.

Theorem 3 (Freiman rectification principle, [10], [54]). *For any $\sigma \in \mathbb{R}_{>0}$ and $k \geq 1$ there exists a constant $c > 0$ such that the following holds. If $A \subseteq \mathbb{Z}_p, |A| \leq cp$ and $|A+A| < \sigma|A|$, then there exists a set of integers $A' \subset \mathbb{Z}$ such that the canonical homomorphism $\mathbb{Z} \to \mathbb{Z}_p$ induces an F_k-isomorphism of A' onto A.*

It is now a natural question if it is possible to preserve both the additive and multiplicative structure. In this direction we have the following result of Vu, Wood and Wood.

Theorem 4 ([131]). *Let S be a finite subset of a characteristic zero integral domain D, and let L be a finite set of non-zero elements in the subring $\mathbb{Z}[S]$ of D. There exists an infinite sequence of primes with positive relative density such that for each prime p in the sequence, there is a ring homomorphism $\phi_p : \mathbb{Z}[S] \to \mathbb{F}_p$ satisfying $0 \notin \phi_p(L)$.*

Here $\mathbb{Z}[S]$ is the smallest subring of D containing S.

It was asked by Vu, Wood and Wood [131] whether given a small enough set $A \subseteq \mathbb{F}_p$, it is possible to map A to some characteristic zero integral domain, while preserving algebraic incidences. The purpose of Chapter 2 is to give an answer to this question.

1.1.2 The results

Let $f \in \mathbb{Z}[x_1, \ldots, x_n]$ be any polynomial with integer coefficients, and write $f = \sum_{\mathbf{x}} c_{\mathbf{x}}\mathbf{x}$, where the sum is taken over all monomials $\mathbf{x}$ in $x_1, \ldots, x_n$. We define

$$\|f\|_1 = \sum_{\mathbf{x}} |c_{\mathbf{x}}|,$$
$$\|f\|_\infty = \max_{\mathbf{x}} |c_{\mathbf{x}}|.$$

All rings considered are commutative and with 1. Consequently given $a_1, a_2, \ldots, a_n \in R$ in some ring R, it makes sense to evaluate f at $(a_1, \ldots, a_n)$, by carrying the operations in R, in the natural way (the integer k becomes $1 + 1 + \ldots + 1$, k times, where 1 is the unit in R).

Definition 5. *Let $k, t > 0$. A polynomial $f \in \mathbb{Z}[x_1, \dots, x_n]$ is called (k,t)-*bounded *if $\|f\|_1 \leq k$, and its degree is at most t. If $k = t$, we simply call f k-*bounded*. If $t = 1$, we say f is a k-*bounded linear polynomial*.*

Let R_1, R_2 be two rings and $A \subseteq R_1, B \subseteq R_2$ finite subsets.

Definition 6. *A bijection $\phi : A \to B$ is a* Freiman ring-isomorphism of order k*, or simply* F_k-ring-isomorphism*, if $A = \{a_1, \dots, a_n\}$ and for any k-bounded $f \in \mathbb{Z}[x_1, \dots, x_n]$ we have*

$$f(a_1, \dots, a_n) = 0$$

if and only if

$$f(\phi(a_1), \dots, \phi(a_n)) = 0.$$

The main result of Chapter 2 is the following.

Theorem 7. *Let $k \geq 2$ be an integer, p be a prime and $A \subseteq \mathbb{F}_p$. If*

$$|A| < \log_2 \log_{2k} \log_{2k^2} p \;-\; 1$$

then there exists a finite algebraic extension K of $\mathbb{Q}$ of degree at most $(2k)^{2^{|A|}}$, a subset $A' \subset K$ and a homomorphism $\phi_p : \mathbb{Z}[A'] \to \mathbb{F}_p$ such that ϕ_p is an F_k-ring-isomorphism between A' and A.

One can use the construction from [10] to see that for any $k \geq 2$ and any prime number p there exists a subset $A \subseteq \mathbb{F}_p$ of size $O(\log p)$, which is not F_k-ring-isomorphic to any subset of a characteristic zero integral domain. For $k \geq 3$ we can improve this bound to the following.

Theorem 8. *For any $k \geq 3$ and any prime number $p \geq 2^{32(k-1)^2 \log_2^2(16(k-1))}$ there exists a subset $A \subseteq \mathbb{F}_p$ of size $|A| \leq \frac{10}{k-1} \frac{\log_2 p}{\log_2 \log_2 p}$ which is not F_k-ring-isomorphic to any subset of a characteristic zero integral domain.*

It is an open problem if a better bound is possible. In this direction we make the following conjecture.

Conjecture 9. *For any $k \geq 3$ there is an infinite sequence of prime numbers, such that for each prime p in the sequence, there exists a subset $A \subseteq \mathbb{F}_p$ of size $O(\log \log p)$ which is not F_k-ring-isomorphic to any subset of a characteristic zero integral domain.*

As explained in Section 2.3, this conjecture would have a positive answer if, for example, there are infinitely many Mersenne primes (primes of the form $2^n - 1$; this would follow from the Lenstra–Pomerance–Wagstaff conjecture), or infinitely many Fermat primes (primes of the form $2^{2^n} + 1$; this is a question of Eisenstein).

1.1.3 Applications of the main theorem

Theorem 7 has several applications to subsets of $\mathbb{F}_p$ of size $O(\log \log \log p)$, which we now describe.

1.1.3.1 The Szemerédi-Trotter theorem

The well-known Szemerédi-Trotter theorem gives a tight upper bound on the number of incidences between a finite set of lines and a finite set of points in $\mathbb{R} \times \mathbb{R}$. This was extended to the complex plane $\mathbb{C}^2$ by Tóth.

Theorem 10 ([129])**.** *Let $\mathcal{P}$ and $\mathcal{L}$ be sets of points and lines in $\mathbb{C}^2$, with cardinalities $|\mathcal{P}|, |\mathcal{L}| \leq n$. Then there is a positive absolute constant c such that*

$$|\{(p,l) \in \mathcal{P} \times \mathcal{L} : p \in l\}| \leq cn^{4/3}.$$

Recently Zahl gave a different proof of Theorem 10 in [134]. Furthermore, if we allow an $\varepsilon > 0$ error in the exponent, and the constant c to depend on ε, then in this form Theorem 10 follows from a generalization of the Szemerédi-Trotter theorem to algebraic varieties due to Solymosi and Tao [118].

The problem of establishing a similar bound in $\mathbb{F}_p$ has been considered before ([15], [63]). We have the following result, due to Helfgott and Rudnev.

Theorem 11 ([63])**.** *Let p be a prime number, and $\mathcal{P}$ and $\mathcal{L}$ sets of points and lines in $\mathbb{F}_p^2$, with $|\mathcal{P}|, |\mathcal{L}| \leq n$ and $n < p$. Then there is a positive absolute constant c such that*

$$|\{(p,l) \in \mathcal{P} \times \mathcal{L} : p \in l\}| \leq cn^{\frac{3}{2}-\delta},$$

with $\delta = \frac{1}{10678}$.

The best (still unpublished) bound to date for $n < p$ is due to Jones [70], who proved that one can take $\delta = \frac{1}{662} - o(1)$ in the above.

We show that one can achieve optimal exponent $4/3$ in Theorem 11 provided n is sufficiently small compared to p.

Theorem 12. *Let p be a prime number, and $\mathcal{P}$ and $\mathcal{L}$ sets of points and lines in $\mathbb{F}_p^2$, with $|\mathcal{P}|, |\mathcal{L}| \leq n$ and $5n < \log_2 \log_6 \log_{18} p - 1$. Then there is a positive absolute constant c such that*

$$|\{(p,l) \in \mathcal{P} \times \mathcal{L} : p \in l\}| \leq cn^{4/3}.$$

Moreover, this inequality is sharp up to the constant c.

The proof of Theorem 12 is in Section 2.4.

One can now combine Theorem 12 with Theorem 4 to generalize Theorem 10 to any characteristic zero integral domain. As this statement can be proved directly with no recurse to Theorem 12, we do not discuss it here (see Theorem 2.3 and Lemma 7.1 from [131] for more details).

1.1.3.2 Sum-product estimates in $\mathbb{F}_p$

Suppose R is a commutative ring and $A \subset R$ a finite subset. We can define the sumset $A+A := \{a+b : a, b \in A\}$ and the product $A \cdot A := \{ab : a, b \in A\}$. Intuitively, the quantities $|A+A|$ and $|A \cdot A|$ can not both be small. The prototype theorem is a lower bound of the form $\max\{|A+A|, |A\cdot A|\} \geq c|A|^{1+\varepsilon_R}$, where $c>0$ is an absolute constant and ε_R depends on the ring R. The first sum-product estimate is due to Erdős and Szemerédi [40] for the case $R=\mathbb{Z}$ and it was followed by numerous improvements and generalizations ([32], [88], [43], [21], [117]). For $R=\mathbb{C}$, the best-known value $\varepsilon_{\mathbb{C}} = \frac{3}{11} - o(1)$ was for many years given by a result of Solymosi [116]. Using a beautiful geometric argument, Konyagin and Rudnev [76] have very recently improved this to $\varepsilon_{\mathbb{C}} = \frac{1}{3} - o(1)$, thus matching the lower bound for the reals.

Theorem 13 ([76])**.** *Suppose $A \subset \mathbb{C}$. Then there is a positive absolute constant c such that*

$$|A+A| + |A\cdot A| \geq c|A|^{1+\frac{1}{3}-o(1)}. \tag{1.1}$$

Bourgain, Katz and Tao [15] showed that a sum-product theorem holds in $\mathbb{F}_p$. Substantial work has gone into finding the best value for $\varepsilon_{\mathbb{F}_p}$. Garaev [49] showed that for $|A| < \sqrt{p}$ one can take $\varepsilon_{\mathbb{F}_p} = \frac{1}{14} - o(1)$. Katz and Shen [72] improved this to $\frac{1}{13} - o(1)$, and then Bourgain and Garaev [14] showed that $\frac{1}{12} - o(1)$ is in fact possible. Li [80] later removed the $o(1)$ term. The best result to date is due to Rudnev [107], who showed that

$$|A+A| + |A\cdot A| \geq c|A|^{1+\frac{1}{11}-o(1)}, \tag{1.2}$$

whenever $|A| < \sqrt{p}$.

We now improve (1.2) for small A.

Theorem 14. *Let p be a prime number and $A \subseteq \mathbb{F}_p$ with $|A| < \log_2 \log_8 \log_{32} p - 1$. Then*

$$|A+A| + |A\cdot A| \geq c|A|^{1+\frac{1}{3}-o(1)},$$

for some positive absolute constant c.

The proof of Theorem 14 is in Section 2.4.

1.1.3.3 Estimates for sets with small doubling constant

We gather in this section several miscellaneous results for the case when A has small doubling constant. We first have the following result, due to Solymosi.

Theorem 15 ([116])**.** *If $A \subset \mathbb{C}$ and $|A| = n$ with $|A+A| \leq Cn$, then $|A\cdot A| \geq cn^2/\log n$.*

This transfers immediately to $\mathbb{F}_p$ as follows.

Theorem 16. *If $A \subseteq \mathbb{F}_p$ and $|A| = n < \log_2 \log_8 \log_{32} p - 1$ with $|A+A| \leq Cn$, then $|A\cdot A| \geq cn^2/\log n$.*

The proof is similar to that of Theorem 14 and we omit it. We also have the following result due to Chang [20].

Theorem 17. *Let $A \subset \mathbb{C}$ with $|A| = n$ and $|A+A| \leq Cn$, for some $C > 0$. Then the following holds.*

(i) If $0 \notin A$ then $|A^{-1} + A^{-1}| > \exp^{-C' \frac{\log n}{\log \log n}} n^2$, for some C' depending only on C.

(ii) If $f(x) \in \mathbb{C}[x]$ is a polynomial of degree $t \geq 2$ then $|f(A) + f(A)| > \exp^{-C' \frac{\log n}{\log \log n}} n^2$, for some $C' := C'(C, t)$.

Here $A^{-1} = \{a^{-1} : a \in A\}$ and $f(A) = \{f(a) : a \in A\}$. The proof of Theorem 17 uses algebraic methods, in particular Lemma 56, but also relies crucially on facts specific to $\mathbb{C}$. We now transfer this theorem to small subsets of $\mathbb{F}_p$.

Theorem 18. *Let $A \subseteq \mathbb{F}_p$ with $|A| = n$ and $|A + A| \leq Cn$, for some $C > 0$. Then the following holds.*

(i) Suppose $2n < \log_2 \log_8 \log_{32} p - 1$ and $0 \notin A$. Then $|A^{-1} + A^{-1}| > \exp^{-C' \frac{\log n}{\log \log n}} n^2$, for some C' depending only on C.

(ii) Let $f(x) \in \mathbb{Z}[x]$ be a k-bounded polynomial of degree at least 2. If $n < \log_2 \log_{8k} \log_{32k^2} p - 1$ then $|f(A) + f(A)| > \exp^{-C' \frac{\log n}{\log \log n}} n^2$, for some $C' := C'(C, k)$.

The proof of Theorem 18 is in Section 2.4.

1.1.3.4 A question of Rényi

Let K be a field of characteristic zero. For a polynomial $f \in K[x]$ we define $N(f)$ to be the number of non-zero terms of f. For $k \geq 1$, let

$$Q_K(k) = \min_{f \in K[x] : N(f) = k} N(f^2). \tag{1.3}$$

As reported by Erdős [89], it was first asked by Rédei if $Q_{\mathbb{R}}(k) < k$ is possible, and Rényi [101] later constructed an example showing $Q_{\mathbb{Q}}(29) \leq 28$. Rényi made several conjectures about the behaviour of $Q_{\mathbb{R}}(k)$.

He conjectured that $\lim_{k \to \infty} \frac{Q_{\mathbb{R}}(k)}{k} = 0$, and this was proved by Erdős [89], who in fact showed that $Q_{\mathbb{Q}}(k) < ck^{1-\varepsilon}$, for some positive absolute constants c and ε.

Rényi further conjectured that $\lim_{k \to \infty} Q_{\mathbb{R}}(k) = \infty$, and this was proved many years later by Schinzel [109], using a very ingenious argument. Schinzel showed that $Q_K(k) \geq c \log \log k$, for some positive absolute constant c and any field K of characteristic zero. This lower bound was not improved for another 20 years, until recently Schinzel and Zannier [110], by an adaptation of the original method of Schinzel, proved that $Q_K(k) \geq c \log k$, for some positive absolute constant c.

Erdős [89] asked for the determination of the order of $Q_{\mathbb{R}}(k)$, and the general belief seems to be that $Q_{\mathbb{R}}(k)$ should be closer to the upper bound than the lower bound. Despite some work in this direction ([130], [28]), a solution to this problem seems at present out of reach.

From the definition we see that for any $k \geq 1$,

$$Q_{\mathbb{C}}(k) \leq Q_{\mathbb{R}}(k) \leq Q_{\mathbb{Q}}(k). \tag{1.4}$$

It is less known that Rényi [101] (see also [89]) asked whether equality holds in (1.4) everywhere for any k, and this problem seems to have received little attention.

For any $k \geq 1$ it also holds that

$$Q_{\mathbb{C}}(k) \leq Q_K(k) \leq Q_{\mathbb{Q}}(k), \tag{1.5}$$

for any finite algebraic extension K of $\mathbb{Q}$, and thus if we have equality in (1.4), then we also have equality in (1.5). In view of this we have the following result.

Theorem 19. *For any $k \geq 3$ there exists a finite algebraic extension K of $\mathbb{Q}$ such that $Q_{\mathbb{C}}(k) = Q_K(k)$, with degree at most k^{2^k}, if k is even, and at most $(k+1)^{2^k}$, if k is odd.*

The proof of Theorem 19 is in Section 2.4.

1.2 Turán densities of hypergraphs

1.2.1 The set of Turán densities

Let $r \geq 1$ and $\mathcal{F}$ be a (possibly infinite) family of r-graphs (i.e. hypergraphs where each edge has the same size r).

For any $n \geq 1$, we define the *Turán function* $\mathrm{ex}(n, \mathcal{F})$ as the maximum possible number of edges of an $\mathcal{F}$-free r-graph on n vertices. If no such r-graph exists, we set $\mathrm{ex}(n, \mathcal{F}) = 0$. If $\mathcal{F} = \{H\}$, we may simply write $\mathrm{ex}(n, H)$ instead of $\mathrm{ex}(n, \{H\})$.

The study of the Turán function was started by the following question: what is the maximum number of edges of a graph G not containing a k-clique? This question was posed, and answered, by Turán in 1941, essentially founding the field of extremal graph theory. His result was later generalized by Erdős and Stone, and Simonovits ([38], [37]).

Theorem 20 (Erdős-Stone-Simonovits theorem)**.** *For any graph H of chromatic number $k \geq 2$ and any $n \geq 1$ we have* $\mathrm{ex}(n, H) = (1 - \frac{1}{k-1})\frac{n^2}{2} + o(n^2)$.

As observed by Katona, Nemetz and Simonovits [71], for a family $\mathcal{F}$ of r-graphs we can define the Turán density of $\mathcal{F}$ as

$$\pi(\mathcal{F}) = \lim_{n \to \infty} \frac{\mathrm{ex}(n, \mathcal{F})}{\binom{n}{r}},$$

and this limit always exists. Thus $\pi(\mathcal{F})$ captures the asymptotic behaviour of $\mathrm{ex}(n, \mathcal{F})$ and so a first step towards understanding the behaviour of ex would be to determine $\pi(\mathcal{F})$ for all families $\mathcal{F}$.

It follows rather easily from Theorem 20 that for a family $\mathcal{F}$ of non-empty graphs, $\pi(\mathcal{F}) = 1 - \frac{1}{k-1}$, where $k := \min\{\chi(H) : H \in \mathcal{F}\}$. This determines π in the case of graphs ($r = 2$). In contrast, very little is known about Turán densities of hypergraphs ($r \geq 3$). In fact, to date no value $\pi(K_t^r)$ for $3 \leq r < t$ has been determined. It has been conjectured by Turán that $\pi(K_4^3) = 5/9$ and much effort was devoted to the resolution of this conjecture.

As a consequence it makes sense to study an easier problem. Let $\Pi_\infty^{(r)}$ consist of all possible Turán densities of r-graph families and

$$\Pi_{\mathrm{fin}}^{(r)} = \{\pi(\mathcal{F}) : \mathcal{F} \text{ is a finite family of } r\text{-graphs}\}.$$

Clearly $\Pi_{\mathrm{fin}}^{(r)} \subseteq \Pi_\infty^{(r)}$. Moreover, by Theorem 20,

$$\Pi_{\mathrm{fin}}^{(2)} = \Pi_\infty^{(2)} = \{1\} \cup \{1 - \frac{1}{k} : k \geq 1\}. \tag{1.6}$$

Not much is known about these sets if $r \geq 3$. Erdős [36] offered \$1000 for the complete determination of $\Pi_\infty^{(r)}$ for all r. We recount below the little information we have about $\Pi_{\mathrm{fin}}^{(r)}$ and $\Pi_\infty^{(r)}$.

One of the oldest results about $\Pi_\infty^{(r)}$ is due to Erdős [34], who proved that $\Pi_\infty^{(r)} \cap (0, r!/r^r) = \emptyset$. Erdős [36] then proposed the study of jumps of hypergraphs.

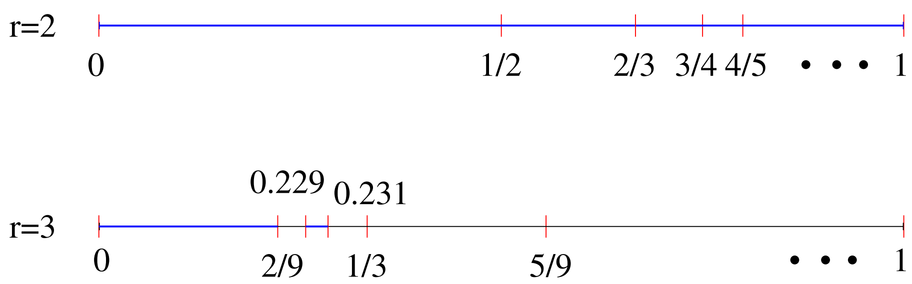

Figure 1.2.1: The sets $\Pi_\infty^{(2)}$ and $\Pi_\infty^{(3)}$

Definition 21. *A number $\alpha \in (0,1)$ is called a jump for r-graphs if there exists some $\varepsilon > 0$ with $\Pi_\infty^{(r)} \cap (\alpha, \alpha+\varepsilon) = \emptyset$.*

In particular all numbers in $[0, r!/r^r)$ are jumps for r-graphs. Erdős [36] conjectured that in fact all numbers in $(0,1)$ are jumps. Clearly this is the case for $r = 2$, but Frankl and Rödl [46] famously disproved the conjecture by showing that $1 - 1/\ell^{r-1}$ is a non-jump for r-graphs, for every $\ell > 2r, r \geq 3$. Erdős ([35], [36]) further conjectured that $r!/r^r$ is always a jump for r-graphs, and offered \$500 for a solution. This conjecture (called the *jumping constant conjecture*) is still open. Many examples of non-jumps were constructed using the method of Frankl and Rödl ([91], [92], [95], [93]), the smallest of which is $\frac{5}{2}\frac{r!}{r^r}$ [45].

Brown and Simonovits [17] proved that $\Pi_\infty^{(r)} \subseteq \overline{\Pi}_{\text{fin}}^{(r)}$. It is a non-trivial result that equality holds here.

Theorem 22 (Pikhurko, [96])**.** *The set $\Pi_\infty^{(r)}$ is closed in $[0,1]$.*

Furthermore, Pikhurko proved the following.

Theorem 23 (Pikhurko, [96])**.** *For every $r \geq 3$ the set $\Pi_\infty^{(r)}$ has cardinality of the continuum.*

In particular, as $\Pi_{\text{fin}}^{(r)}$ is countable, this means $\Pi_\infty^{(r)} \neq \Pi_{\text{fin}}^{(r)}$ for $r \geq 3$.

It is an open question if $\Pi_\infty^{(r)}$ contains an interval of positive length for $r \geq 3$. Proving that a certain number *does not belong* to $\Pi_\infty^{(r)}$ seems to be very hard. So far Baber and Talbot [5] proved that $[0.2299, 0.2316) \cap \Pi_\infty^{(3)} = \emptyset$, that $\pi(K_4^{3-})$ is a jump for 3-graphs, and by upper-bounding $\pi(K_4^{3-})$, they proved that $[0.2871, 8/27) \cap \Pi_\infty^{(3)} = \emptyset$. The proof uses flag algebras, introduced and developed by Razborov [99]. Flag algebras have been successfully used for computing Turán densities in certain special cases ([100], [6], [41]), and also for solving several open questions in graph theory ([61], [58], [62], [78], [7], [50]).

The (known) structure of $\Pi_\infty^{(2)}$ and $\Pi_\infty^{(3)}$ is displayed in Figure 1.2.1. The blue segments denote intervals of numbers that do not belong to $\Pi_\infty^{(r)}$, while the red marks denote Turán densities.

It was proved by Baber and Talbot [6] that $\Pi_{\text{fin}}^{(3)}$ contains irrational numbers, disproving a conjecture of Chung and Graham [26]. Pikhurko independently proved the following more general result.

Theorem 24 (Pikhurko, [96])**.** *For every $r \geq 3$ the set $\Pi_{\text{fin}}^{(r)}$ contains an irrational number.*

Finally, the following question is due to Jacob Fox.

Question 1 (Jacob Fox)**.** *Does $\Pi_{\text{fin}}^{(r)}$ contain a transcendental number?*

Note that $\Pi_{\infty}^{(r)}$ for $r \geq 3$, being uncountable, certainly contains a transcendental number.

1.2.2 The results

Clearly any sort of additional information about $\Pi_{\text{fin}}^{(r)}$ or $\Pi_{\infty}^{(r)}$ would be useful in understanding the structure of these sets. In light of this, the results described in Chapter 3 concern the algebraic structure of $\Pi_{\text{fin}}^{(r)}$ and $\Pi_{\infty}^{(r)}$. An additional motivation for this study is Question 1: a sufficiently strong algebraic structure might be enough to settle the existence of transcendental Turán densities.

We will also briefly have something to say about the topological properties of $\Pi_{\infty}^{(r)}$.

Define $\mathfrak{h}_2 : [0,1) \to \mathbb{R}$ by $\mathfrak{h}_2(x) = \frac{1}{1-x}$. Then $\mathfrak{h}_2(\Pi_{\text{fin}}^{(2)} \setminus \{1\}) = \mathbb{N}$. The starting observation is that $\mathbb{N}$ is a semiring with two operations $+$ and $\cdot$. These two operations correspond to the following combinatorial constructions: given an n-clique and an m-clique, one can form an $(n+m-1)$-clique (by identifying a vertex and adding $(n-1)(m-1)$ additional edges), or an nm-clique (by blowing up every vertex of the n-clique with a copy of the m-clique). Thus it would be interesting to study whether these two operations (and their corresponding constructions) generalize to arbitrary r.

For any $r \geq 2$, define $\mathfrak{h}_r : [0,1) \to \mathbb{R}$ by $\mathfrak{h}_r(x) = \sqrt[r-1]{\frac{1}{1-x}}$. Note that in the case $r = 2$ this is the same as the previous definition.

Theorem 25. *The sets $\mathfrak{h}_r(\Pi_{\text{fin}}^{(r)} \setminus \{1\})$ and $\mathfrak{h}_r(\Pi_{\infty}^{(r)} \setminus \{1\})$ are closed under addition.*

This shows that the addition operation generalizes in a natural way.

As in the case $r = 2$, Theorem 25 has a combinatorial interpretation. If $\mathcal{F}_\alpha$ and $\mathcal{F}_\beta$ are two families of r-graphs with $\pi(\mathcal{F}_\alpha) = \alpha$ and $\pi(\mathcal{F}_\beta) = \beta$, we wish to show that there exists a family of r-graphs $\mathcal{F}$ with $\mathfrak{h}_r(\pi(\mathcal{F})) = \mathfrak{h}_r(\alpha) + \mathfrak{h}_r(\beta)$. In the case when the families are allowed to be infinite, a rather simple argument gives the existence of $\mathcal{F}$. For finite families, however, $\mathcal{F}$ is roughly constructed in the following way. For any $F_\alpha \in \mathcal{F}_\alpha, v \in V(F_\alpha)$ and any $F_\beta \in \mathcal{F}_\beta, w \in V(F_\beta)$, a hypergraph $H(F_\alpha, F_\beta, v, w)$ is formed by identifying v with w and then *appending* a special structure. The family $\mathcal{F}$ contains all such hypergraphs $H(F_\alpha, F_\beta, v, w)$, and some more, due to the use of the Strong Removal Lemma. Although the special structure appended is rather complicated to describe, in the case of $F_\alpha = K_n^r, F_\beta = K_m^r$ one can imagine $H(F_\alpha, F_\beta)$ as K_{n+m-1}^r, thus giving a correspondence with the case $r = 2$.

Theorem 25 has several consequences.

Theorem 26. *For every $r \geq 3$ the set $\Pi_{\text{fin}}^{(r)}$ contains the irrational number*

$$1 - \frac{r^{r-1} - (r-1)!}{\left(r + \sqrt[r-1]{r^{r-1} - (r-1)!}\right)^{r-1}}.$$

This in particular provides a new proof of Theorem 24, as well as explicit examples of irrational Turán densities. For even values of r simpler examples can be given.

Theorem 27. *For every even* $r \geq 4$ *the set* $\Pi_{\mathrm{fin}}^{(r)}$ *contains the irrational number* $1 - \frac{1}{(1+\sqrt[r-1]{2})^{r-1}}$.

From the proof of Theorem 25 one can also extract the following.

Theorem 28. *For all* $r \geq 2$ *and* $t \geq r$ *there exists a finite family of* r*-uniform hypergraphs* $\mathcal{F}_{r,t}$ *such that*

$$\mathfrak{h}_r(\pi(\{K_{t+r-1}^r\} \cup \mathcal{F}_{r,t})) = 1 + \mathfrak{h}_r(\pi(K_t^r)).$$

Thus, for example, $\pi(\{K_5^3\} \cup \mathcal{F}) = \frac{3}{4}$ for some finite $\mathcal{F}$. A family $\mathcal{F}$ with this property for K_5^3 was found by Zhou [135]. Furthermore, from Theorem 28, for $r-1|t-1$ one recovers the classical lower bound for $\pi(K_t^r)$,

$$\pi(K_t^r) \geq 1 - \left(\frac{r-1}{t-1}\right)^{r-1}. \tag{1.7}$$

It is known that this bound is not always sharp (see [113] for a construction), so one needs the family $\mathcal{F}_{r,t}$ in Theorem 28.

It seems very hard to find a proper generalization of the multiplication operation. We will discuss this problem at length in Section 3.5. Nevertheless, it is possible to define a multiplication on *the set of all Turán densities*, defined as

$$\Pi_\infty := \{(\alpha, r) : \alpha \in \Pi_\infty^{(r)}, r \geq 0\},$$

We furthermore define *the set of finite Turán densities* as

$$\Pi_{\mathrm{fin}} := \{(\alpha, r) : \alpha \in \Pi_{\mathrm{fin}}^{(r)}, r \geq 0\}.$$

For technical reasons we set here $\Pi_\infty^{(0)} = \Pi_{\mathrm{fin}}^{(0)} = \{1\}$.

We now define a binary operation $*$ on the set $\mathbb{R} \times \mathbb{N}$, which obviously contains Π_∞:

$$\begin{aligned} *\colon (\mathbb{R} \times \mathbb{N}) \times (\mathbb{R} \times \mathbb{N}) &\to \mathbb{R} \times \mathbb{N} \\ (\alpha, r) \times (\beta, s) \quad &\mapsto (\alpha\beta \binom{r+s}{r} \frac{r^r s^s}{(r+s)^{r+s}}, r+s). \end{aligned}$$

This corresponds to the following combinatorial construction: given an r-graph G and an s-graph H, one can form an $(r+s)$-graph on vertex set $V(G)\dot\cup V(H)$ and edge set $\{e \cup f : e \in E(G), f \in E(H)\}$.

We have the following result.

Theorem 29. $(\Pi_\infty, *)$ *is a commutative cancellative monoid.*

In particular, $\mathbb{Z}\Pi_\infty \simeq \bigoplus_{r\geq 0} \mathbb{Z}\Pi_\infty^{(r)}$ is a graded ring under $*$ (here $\mathbb{Z}\Pi_\infty^{(r)}$ is the free abelian group generated by $\Pi_\infty^{(r)}$). Theorems 25 and 29 have the following consequence.

Theorem 30. *For any $r \geq 2$ the following statements are equivalent:*

(i) $\Pi_\infty^{(r)}$ *has positive Lebesgue measure.*

(ii) $\Pi_\infty^{(r)}$ *contains an open interval.*

(iii) For any $r' \geq r$, $\Pi_\infty^{(r')}$ *has positive Lebesgue measure.*

(iv) For any $r' \geq r$, $\Pi_\infty^{(r')}$ *contains an open interval.*

Finally, it is possible to deduce the following from Theorem 29.

Theorem 31. *Let $r \geq 3$ and $c > 0$. Suppose $c\frac{r!}{r^r}$ is a non-jump for r-graphs. Then $c\frac{q!}{q^q}$ is a non-jump for q-graphs, for any $q \geq r$.*

Theorem 31 was originally proved by Peng in [94]. Theorems 25 and 29 can be used to give many new examples of non-jumps. In fact, starting from any non-jump and applying the addition operation (see the definition of $\oplus_r$ in Section 1.2.3) or the operation $*$, gives a new non-jump. These are perhaps the first examples of non-jumps which are not constructed by the method of Frankl and Rödl.

Our study of Turán densities of hypergraphs led us to the following conjecture, which was later proved by Pikhurko [97].

Conjecture 32 (Theorem 1, [97])**.** $\overline{\cup_{r\geq 2}\Pi_{\text{fin}}^{(r)}} = \overline{\cup_{r\geq 2}\Pi_\infty^{(r)}} = [0,1]$.

1.2.3 A reformulation of the main theorem

Theorem 25 says that $\mathfrak{h}_r(\Pi_\infty^{(r)} \setminus \{1\})$ is closed under addition. It is possible to pull back the addition operation via $\mathfrak{h}_r$ in order to obtain a semigroup operation for $\Pi_\infty^{(r)}$.

For any $r \geq 2$, let us define $\oplus_r : [0,1] \times [0,1] \to [0,1]$ by

$$\alpha \oplus_r \beta = 1 - \frac{1 - \alpha - \beta + \alpha\beta}{(\sqrt[r-1]{1-\alpha} + \sqrt[r-1]{1-\beta})^{r-1}}, \tag{1.8}$$

and $1 \oplus_r 1 = 1$.

It is easy to see that $\mathfrak{h}_r(\alpha \oplus_r \beta) = \mathfrak{h}_r(\alpha) + \mathfrak{h}_r(\beta)$, for any $\alpha, \beta \in [0,1)$. Hence Theorem 25 can be rewritten as follows.

Theorem 33. $(\Pi_\infty^{(r)}, \oplus_r)$ *is a commutative topological semigroup, and* $\Pi_{\text{fin}}^{(r)}$ *is closed under* $\oplus_r$.

We shall find this statement more suitable for a proof.

1.3 The Graceful Tree Conjecture

1.3.1 The conjectures of Ringel, Kotzig and Rosa

In 1963, at the Czechoslovak Symposium on Graph Theory in Smolenice, Ringel [102] proposed the following conjecture:

Conjecture 34 (Ringel's conjecture, [102]). *For any $(n+1)$-vertex tree T the complete graph K_{2n+1} can be decomposed into $2n+1$ edge-disjoint subgraphs isomorphic to T.*

To this day the conjecture is still unsolved. What makes this problem so hard is the small amount of "free space" available when embedding the copies of T. Indeed, if instead we would require a decomposition of K_m into edge-disjoint copies of an $(n+1)$-vertex tree, where $n|\binom{m}{2}$, then such a decomposition always exists provided m is large enough. This follows from a general result of Wilson [133].

It is reported by Rosa [106] that Kotzig later conjectured a stronger statement, which we now describe.

Identify the vertices of K_{2n+1} with the integers $0, 1, \ldots, 2n$. Then for any subgraph G of K_{2n+1} we may define the *cyclic shift* of G as the subgraph $S(G)$ given by

$$S(G) = (\{x+1 : x \in V(G)\}, \{(x+1, y+1) : (x, y) \in E(G)\})$$

where all addition is performed modulo $2n+1$.

If G is any graph with n edges, we say that K_{2n+1} can be *cyclically decomposed* into copies of G if there is a subgraph $G' \simeq G$ of K_{2n+1} such that the cyclic shifts $G', S(G'), \ldots, S^{2n}(G')$ are edge-disjoint (and thus form a decomposition of K_{2n+1}).

Kotzig conjectured the following.

Conjecture 35 (Ringel–Kotzig conjecture). *For any $(n+1)$-vertex tree T the complete graph K_{2n+1} can be cyclically decomposed into copies of T.*

In an effort to tackle Conjecture 35, Rosa introduced β-valuations. The name was later changed to *graceful labellings* by Golomb [51].

Let G be a graph with q edges and $\psi : V(G) \to [q+1]$ a labelling of the vertices of G with numbers from 1 to $q+1$. Then any edge $e = (x, y) \in E(G)$ has an *induced labelling* given by $|\psi(x) - \psi(y)|$. The map ψ is called *graceful* if it is injective and all the edge labels are pairwise distinct. It then follows that the set of edge labels is precisely the set $[q]$.

Rosa [106] showed that a graph with every vertex of even degree and number of edges congruent to 1 or 2 (mod 4) is not graceful. It was also shown by Graham and Sloane [53] that almost all graphs are not graceful.

Nevertheless, the case of trees is still wide open. In this direction we have the following conjecture of Rosa.

Conjecture 36 (Rosa, [106]). *Every tree has a graceful labelling.*

This conjecture became known as *Rosa's conjecture* or the *Graceful Tree Conjecture.*

Rosa further showed the following.

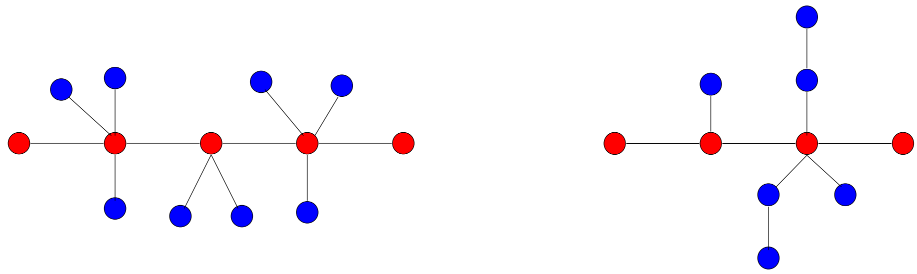

Figure 1.3.1: A caterpillar and a lobster tree

Lemma 37. *Conjecture 36 implies Conjecture 35.*

Proof. Let T be any $(n+1)$-vertex tree and $\psi : V(T) \to [n+1]$ the graceful labelling of T guaranteed by Conjecture 36. We can modify ψ so that the image is the set $\{0, 1, \ldots, n\}$.

Let T' be a subgraph of K_{2n+1} with vertices in the set $\{0, 1, \ldots, n\}$ and isomorphic to T under ψ^{-1}. We claim that the cyclic shifts of T' decompose K_{2n+1}: otherwise there exist integers $a < b$ and distinct edges $(u_1, v_1), (u_2, v_2) \in E(T')$ such that $u_1 + a \equiv u_2 + b \pmod{2n+1}$ and $v_1 + a \equiv v_2 + b \pmod{2n+1}$. But then $|v_1 - u_1| = |v_2 - u_2|$ so by the definition of ψ the two edges are the same, a contradiction. $\square$

Due to Lemma 37, the interest in Ringel's conjecture shifted towards searching for a proof of Conjecture 36.

1.3.2 Classes of graceful trees

Despite attracting a large amount of attention, the Graceful Tree Conjecture is only known in several special cases. A detailed survey of the current status of the conjecture can be found in [48]. We will try to summarize below some of the most important results.

A *caterpillar* is a tree T with the following property: there exists a path $P \subseteq T$ such that every vertex of T is eiter on P or is adjacent to a vertex of P. Paths and stars are important special cases of this construction. Rosa [106] showed that all caterpillars are graceful.

A *lobster* is a tree T having a path P such that every vertex of T is eiter on P or is at distance at most 2 from P. Lobsters can be seen as a natural generalization of caterpillars. Bermond [9] conjectured that all lobsters are graceful and despite a substantial effort and the relatively simple structure of these trees, the conjecture is still open.

A *firecracker* is a tree formed by joining the centers of a collection of stars to vertices of a path, in such a way that every vertex of the path is joined to the center of at most one star. Firecrackers are known to be graceful [22].

An *olive tree* is constructed by taking k paths, of increasing lengths from 1 up to k, and identifying one endpoint from each to form the root of the tree. Olive trees were shown to be graceful by Pastel and Raynaud [90].

All trees of diameter up to five are graceful [65]. Very recently all trees of diameter at most 7 were shown to be graceful [132].

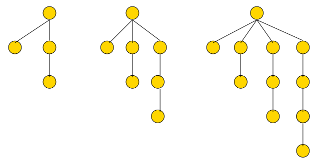

Figure 1.3.2: Several olive trees

Aldred and McKay [2] used a computer to show that all trees with at most 27 vertices are graceful. Recently Fang [42] claimed that all trees up to 35 vertices are graceful, but the result is still unpublished.

A major obstacle towards a proof by induction is that there is no known way of decomposing the problem of finding a graceful labelling into smaller subproblems. On the other hand, it is sometimes possible to combine two (or more) gracefully labelled trees into a larger graceful tree. To this end, several operations that produce larger graceful trees have been proposed. Nevertheless, all these operations are either specialised on some classes of trees or have rather heavy additional constraints, leaving no hope that all trees can be generated in this manner. More details on this approach can be found in [31].

1.3.3 The results

Given the difficulty of Conjecture 36, it is natural to try to study a relaxed version where the set of vertex labels is slightly larger than $[n]$.

We call a map $\psi : V(G) \to [m]$ from the vertex set of a graph *m-graceful* if ψ is injective, and the map ψ_* induced on the edges, $\psi_* : E(G) \to [m-1]$, $\psi_*(xy) := |\psi(x) - \psi(y)|$, is injective as well. If m is clear from the context, we simply call ψ *graceful.*

In joint work with Anna Adamaszek, Michał Adamaszek, Peter Allen and Jan Hladký we have proved the following.

Theorem 38. *For every $\Delta \in \mathbb{N}$ and every $\varepsilon > 0$ there exists a number $n_0 \geq 1$ such that the following holds for every $n > n_0$. Suppose that T is an n-vertex tree with $\Delta(T) \leq \Delta$. Then there exists a graceful labelling $\psi : V(T) \to [(1+\varepsilon)n]$.*

Theorem 38 shows that the Graceful Tree Conjecture almost holds for trees of bounded maximum degree.

We further extended this result to random trees. For our purposes a *random tree of order n* is a uniformly random element from the set of all n^{n-2} labelled trees on vertex set $[n]$. We say that a property P holds *asymptotically almost surely (a.a.s.) for a random tree* if

$$\lim_{n\to\infty} \mathbf{P}[\text{a random tree of order } n \text{ has property P}] = 1.$$

Theorem 39. *For any $\varepsilon > 0$, a tree chosen uniformly at random from the set of labelled n-vertex trees has a graceful labelling $\psi : V(T) \to [(1+\varepsilon)n]$ asymptotically almost surely.*

1.4 The Towers of Hanoi

1.4.1 The puzzle

The Towers of Hanoi is a puzzle invented by the French mathematician Édouard Lucas in 1883 ([81]). The setup consists of 3 pegs and N disks of different sizes, arranged on the first peg in increasing order according to size. The goal is to move the disks from the first peg to another in as few moves as possible, such that the following three rules are always obeyed:

(R1) only one disk can be moved at a time;

(R2) each move consists of taking the topmost disk on a peg and placing it on another peg;

(R3) a smaller disk is always moved on top of a larger one, or on an empty peg.

One can prove by induction that the solution requires $2^N - 1$ moves. The puzzle is very popular, and is frequently used to teach recursive algorithms to first-year computer science students.

Several variations of the original problem have been proposed ([123]), with one possibility being to increase the number of pegs available in the game. The puzzle with 4 pegs was first introduced by Dudeney in 1908 in his book *The Canterbury Puzzles*, under the name "Reve's Puzzle". In 1939, the general problem with p pegs and N disks was proposed in the *American Mathematical Monthly* in the *Advanced Problems* section, as Problem 3918 ([121]). Two years later, the journal published the proposer's (B.M. Stewart) claimed solution [122], as well as one solution submitted by a reader (J. S. Frame) [44]. The two solutions presented essentially the same algorithm, which we will call the Frame-Stewart algorithm, and which we know describe.

Given N disks and p pegs, the algorithm chooses an integer $1 \leq \ell < N$ that minimizes the number of steps in the following formula:

- Move the top ℓ disks from the start peg to an intermediate peg, using p pegs.
- Move the bottom $N - \ell$ disks from the start peg to the goal peg, using $p - 1$ pegs (one peg is blocked by the ℓ smaller disks sitting on it).
- Move the initial ℓ disks from the intermediate peg to the goal peg, using p pegs.

Let $\Phi(p, N)$ denote the number of steps taken by the Frame-Stewart algorithm for N disks and p pegs. Then we have the recursive formula

$$\Phi(p, N) = \min_{1 \leq \ell < N} \{2\Phi(p, \ell) + \Phi(p - 1, N - \ell)\}, \tag{1.9}$$

with initial data $\Phi(3, N) = 2^N - 1$.

Frame [44] and Stewart [122] both derived closed-form formulas for $\Phi(p, N)$. However, as already noted by the Editors of the *Monthly*, the two proofs tacitly assumed the optimality of the algorithm proposed.

In fact, proving that the Frame-Stewart algorithm is best possible has since become a notorious open problem ([82]). However, in 2014, more than a century after Dudeney's book appeared, the case $p = 4$ was finally solved by Bousch ([16]) in a very elegant way.

Let $H(p, N)$ denote the minimum number of steps needed to move N disks frome one peg to another, using p pegs, according to the rules (R1)-(R3). We already know that $H(3, N) = \Phi(3, N)$ (this is the classic puzzle of Lucas). Bousch proved the following.

Theorem 40 (Bousch, [16])**.** *For all $N \geq 1$ we have $H(4, N) = \Phi(4, N)$.*

The general case $p \geq 5$ is still open.

By definition $H(p, N) \leq \Phi(p, N)$. Rather than proving that equality holds, one can instead try to give a lower bound for $H(p, N)$ which hopefully closely matches the upper bound $\Phi(p, N)$. Building upon a result of Szegedy ([125]), Chen and Shen showed the following.

Theorem 41 (Chen-Shen, [23])**.** *For all $p \geq 3$ and $N \geq 1$ we have $H(p, N) \geq 2^{m-1}$, where $m \geq 0$ is the largest integer such that $\binom{m+p-3}{p-2} < N$.*

Note that

$$\sqrt[p-2]{(p-2)!N} - (p-2) \leq m < \sqrt[p-2]{(p-2)!N}$$

and hence $\log_2 H(p, N) \geq c_p \sqrt[p-2]{N}$ for some $c_p > 0$ depending only on p.

On the other hand, it is well-known that $\Phi(p, N) = \Theta(\frac{1}{(p-3)!} m^{p-3} 2^m)$, and so by the above theorem, $\log_2 H(p, N)$ is asymptotically the same as $\log_2 \Phi(p, N)$ for p fixed and N tending to infinity. Theorem 41 gives the best known lower bound on $H(p, N)$.

1.4.2 The results

The main result of Chapter 5 is the following asymptotic improvement of Theorem 41.

Theorem 42. *Let $p \geq 4$ and $N \geq 1$. Write $N - 1 = \binom{m+p-3}{p-2} + \binom{t+p-4}{p-3} + r$, with $m \geq t \geq 0$ and $0 \leq r < \binom{t+p-4}{p-4}$ (this decomposition exists and is unique). Then we have $H(p, N) \geq (m + t)2^{m-2(p-2)}$.*

Note that Theorem 42 essentially improves the lower bound of Theorem 41 by a factor of m, where $m = \Theta_p(\sqrt[p-2]{N})$. Consequently this gives $H(p, N) = \Omega_p(\sqrt[p-2]{N} 2^{\sqrt[p-2]{N}})$. This should be compared with the upper bound $\Phi(p, N) = O_p(N^{\frac{p-3}{p-2}} 2^{\sqrt[p-2]{N}})$.

In particular, for $p = 5$, the previous lower bound of $\Omega(2^{\sqrt[3]{N}})$ is improved by Theorem 42 to $\Omega(\sqrt[3]{N} 2^{\sqrt[3]{N}})$, while the upper bound $\Phi(5, N)$ is of the order $O(N^{\frac{2}{3}} 2^{\sqrt[3]{N}})$.

The proof of Theorem 42 relies on the following idea, introduced by Szegedy. Rather than finding a lower bound for $H(p, N)$, one can try to bound the length $\Gamma(p, N)$ of the shortest sequence of steps that moves every disk at least once (here we also minimize over all possible starting configurations). Clearly $\Gamma(p, N)$ is then a lower bound for $H(p, N)$, as every disk must move at least once from the initial peg to the destination peg in the Hanoi problem. Szegedy has shown the following.

Theorem 43 (Szegedy, [125])**.** *If $N \leq 1$ then $\Gamma(3, N) = N$. Otherwise $\Gamma(3, N) = 1 + 2^{N-2}$.*

The main step in the proof of Theorem 42 is the following result, which may be of independent interest.

Theorem 44. *For all $N \geq 0$ we have*

$$\Gamma(4, N) = \begin{cases} N, & \text{if } N \leq 2, \\ 3 + \frac{\Phi(4,N)-5}{4}, & \text{otherwise.} \end{cases}$$

In fact we believe that the following holds.

Conjecture 45. *For all $p \geq 3$ and $N \geq 0$ we have*

$$\Gamma(p, N) = \begin{cases} N, & \text{if } N \leq p-2, \\ p - 1 + \frac{\Phi(p,N)-(2(p-2)+1)}{4}, & \text{otherwise.} \end{cases}$$

Theorems 43 and 44 show that Conjecture 45 holds for $p \in \{3, 4\}$.

Chapter 2

Freiman ring isomorphisms

2.1 Preparations for the proof

2.1.1 Preserving the additive structure

For comparison reasons we start by sketching a proof of Theorem 2, following [10].

Proof of Theorem 2. We first choose $0 < t < p$ such that multiplying every element of A by t (modulo p) results in a set $A^* \subseteq \{-\lfloor \frac{p}{2k} \rfloor, \dots, \lfloor \frac{p}{2k} \rfloor\}$. The existence of t follows from the Kronecker approximation theorem (Corollary 3.2.5, [128]). Let $m \in \mathbb{Z}$ be such that $mt \equiv 1 \pmod{p}$. We multiply every element of A^* by m to obtain A'. Then the canonical homomorphism maps A' onto A, and one easily sees that this is also an F_k-isomorphism. □

We will now consider the problem of preserving bounded linear polynomials. As we allow non-zero constant terms, we will have to find a proof different from that of Theorem 2.

We first prove an inequality.

Lemma 46. *Suppose $M = (m_{ij})$ is an $n \times n$ matrix with entries $m_{ij} \in \mathbb{Z}[x_1, \dots, x_r]$. If for any i, $\sum_j \|m_{ij}\|_1 \leq k$, then $\|\det(M)\|_1 \leq k^n$. Furthermore, for any matrix M with integer entries, $|\det(M)|$ is at most the product of the $\|\cdot\|_1$-norms of the rows.*

Proof. We use the easily verified inequality $\|fg\|_1 \leq \|f\|_1\|g\|_1$, which holds for any $f, g \in \mathbb{Z}[x_1, \dots, x_r]$, to see that

$$\begin{aligned}\|\det(M)\|_1 &\leq \sum_{\pi \in S_n} \|m_{1\pi(1)}\|_1 \dots \|m_{n\pi(n)}\|_1 \leq \sum_{1 \leq i_1, \dots, i_n \leq n} \|m_{1i_1}\|_1 \dots \|m_{ni_n}\|_1 \\ &= (\sum_j \|m_{1j}\|_1) \dots (\sum_j \|m_{nj}\|_1) \leq k^n.\end{aligned}$$

□

The last statement of Lemma 46 is also a consequence of Hadamard's inequality.

We now have the following technical result.

Lemma 47. *Let $k > 1$ be an integer and p be a prime. Suppose $A = \{a_1, \dots, a_n\} \subseteq \mathbb{Z}_p$ and let $\mathcal{L}_1, \mathcal{L}_2 \subset \mathbb{Z}[x_1, \dots, x_n]$ be collections of k-bounded linear polynomials, such that any $f \in \mathcal{L}_1$ is zero when evaluated at $(a_1, \dots, a_n)$, and any $f \in \mathcal{L}_2$ is non-zero when evaluated at $(a_1, \dots, a_n)$. If $|A| < \log_k p - 1$, then there exists $A' = \{b_1, \dots, b_n\} \subset \mathbb{Z}_{(p)}$ such that* *the canonical homomorphism $\mathbb{Z}_{(p)} \to \mathbb{Z}_p$ maps b_i to a_i, and $f(b_1, \ldots, b_n) = 0$ for $f \in \mathcal{L}_1$, $f(b_1, \ldots, b_n) \neq 0$ for $f \in \mathcal{L}_2$.*

This directly implies Theorem 2, with almost the same bound.

Corollary 48. *Let $k \geq 1$ be an integer and p be a prime. Then for any $A \subseteq \mathbb{Z}_p$ with $|A| < \log_{2k} p - 1$ there exists $A' \subset \mathbb{Z}$ F_k-isomorphic with A via the canonical homomorphism.*

Proof. We consider all linear polynomials in $n := |A|$ variables having $\|\cdot\|_1$-norm at most $2k$, and split them into $\mathcal{L}_1$ and $\mathcal{L}_2$ according to the result of evaluation with elements from A. This includes all polynomials used in the definition of the usual Freiman isomorphism. Applying Lemma 47, we get a subset $A' \subset \mathbb{Z}_{(p)}$, which by definition must be F_k-isomorphic with A via the canonical homomorphism. Multiplying all values of A' by a large enough integer, which is 1 modulo p and cleares all denominators, will ensure that A' lies in $\mathbb{Z}$, while still being F_k-isomorphic with A via the canonical homomorphism. □

Proof of Lemma 47. We can express $\mathcal{L}_1$ as the system $M\mathbf{x} = \mathbf{b}$, for some $m \times n$ matrix M and vector $\mathbf{b}$. We then form the augmented matrix $M' = (M|\mathbf{b})$. By assumption, the $\|\cdot\|_1$-norm of any row of M' is at most k.

The system $\mathcal{L}_1$ is solvable in a field $\mathbb{F}$ if and only if $\mathrm{rk}_{\mathbb{F}} M = \mathrm{rk}_{\mathbb{F}} M'$. We will show that this is the case in $\mathbb{Q}$.

As the rank of M' is the maximum size of one of its square submatrices with non-zero determinant, we see that $\mathrm{rk}_{\mathbb{Q}} M' \geq \mathrm{rk}_{\mathbb{F}_p} M'$. On the other hand, let M_1' be any square submatrix of M' of full rank in $\mathbb{Q}$. By Lemma 46, $|\det(M_1')| \leq k^{n+1} < p$. Hence $\det(M_1')$ is also non-zero in $\mathbb{F}_p$, and consequently $\mathrm{rk}_{\mathbb{Q}} M' \leq \mathrm{rk}_{\mathbb{F}_p} M'$. But then M' has the same rank t in $\mathbb{Q}$ and in $\mathbb{F}_p$. Similarly we obtain that M has the same rank in both $\mathbb{Q}$ and $\mathbb{F}_p$. However, the system $\mathcal{L}_1$ is solvable in $\mathbb{F}_p$, and so we must have $t = \mathrm{rk}\, M \leq n$. Consequently $\mathcal{L}_1$ is solvable in $\mathbb{Q}$. This is nevertheless not enough for our purposes; we must further show that a solution A' with the desired properties exists.

We may assume w.l.o.g. that

$$M = \begin{pmatrix} M_1 & M_2 \\ M_3 & M_4 \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{pmatrix}$$

where M_1 is a square matrix of full rank $t = \mathrm{rk}\, M$ in both $\mathbb{Q}$ and $\mathbb{F}_p$, and $\mathbf{b}$ is partitioned accordingly. Let M_1^* be the adjoint of M_1.

We get

$$\begin{pmatrix} M_1^* & 0 \\ 0 & I \end{pmatrix} \begin{pmatrix} M_1 & M_2 \\ M_3 & M_4 \end{pmatrix} x = \begin{pmatrix} \det(M_1) I & M_1^* M_2 \\ M_3 & M_4 \end{pmatrix} x = \begin{pmatrix} M_1^* b_1 \\ b_2 \end{pmatrix}.$$

By Lemma 46, $|\det(M_1)| \leq k^n$.

Consequently we can express the first t variables in terms of the last $n - t$ variables, involving fractions with denominator bounded by $k^n < p$. By letting $b_i := a_i$ and replacing

x_i with b_i in these equations for $t+1 \leq i \leq n$, we obtain values $b_1, \ldots, b_t$ in $\mathbb{Z}_{(p)}$ for $x_1, \ldots, x_t$ such that b_i is mapped to a_i by the canonical homomorphism, for any $1 \leq i \leq n$. Furthermore, as $\operatorname{rk}_{\mathbb{Q}} M' = t$, by replacing x_i with b_i in the last $m-t$ equations we obtain the identity $0 = 0$ in $\mathbb{Q}$ everywhere.

We conclude that $A' := \{b_1, \ldots, b_n\}$ is a solution for $\mathcal{L}_1$ in $\mathbb{Z}_{(p)}$. Furthermore, no polynomial $f \in \mathcal{L}_2$ can be zero when evaluated at A', for otherwise it would also be zero modulo p, hence 0 when evaluated at A, a contradiction. Then we are done. $\square$

2.1.2 Resultants, subresultants and the gcd

As in the case of linear polynomials, we must bound the complexity of solving a system of multivariate polynomials. We gather in this section all the tools required for the proof.

In what follows we shall introduce and make substantial use of subresultants, an alternative to Euclid's algorithm for computing the greatest common divisor of two polynomials. This approach will be essential in obtaining any reasonable quantitative bound in Theorem 7, as Euclid's algorithm leads to an explosive growth of the coefficients involved in the polynomial division.

A more comprehensive treatment of subresultants is given in Appendix A. We will only restate the necessary definitions.

Suppose R is an integral domain. If $R \subseteq D$, D is a commutative ring and $d \in D$, we shall denote by ev_d the evaluation homomorphism $\mathrm{ev}_d : R[x] \to D$ mapping $f(x)$ to $f(d)$. If $0 \neq r \in R$, we shall denote by $R[\frac{1}{r}]$ the ring of polynomials $R[x]$ evaluated at $\frac{1}{r}$. This is the same as the ring of fractions of R with respect to $\{r^n : n \geq 0\}$, and is sometimes denoted by R_r. If D is another integral domain and $\phi : R \to D$ is a homomorphism, ϕ extends to a homomorphism from $R[x]$ to $D[x]$, which we shall also denote by ϕ.

Let $f, g \in R[x]$. We say $g|f$ if there exists $h \in R[x]$ such that $f = hg$. Hence $h|0$ for any $h \in R[x]$, but 0 divides only 0. Moreover if R is a unique factorization domain (UFD), then $\gcd_R(f,g)$ is well-defined. Here we use the conventions $\gcd_R(h,0) = \gcd_R(0,h) = h$, for any polynomial h. Note that $\gcd_R(f,g)$ is unique only up to a unit of R. If no confusion may occur, we shall drop the subscript R. Furthermore if $f_1, \ldots, f_m \in R[x]$ we let $\gcd(f_1, \ldots, f_m)$ denote their greatest common divisor, where for $m = 1$ this is by convention f_1.

We also make the convention $\deg(0) = -\infty$.

We shall need the following easy fact.

Lemma 49. *Suppose $R \subseteq D$ are integral domains, $f, g \in R[x]$ non-zero and $g|f$ in $D[x]$. Then $g|f$ in $R[\frac{1}{b}]$, where b is the leading coefficient of g.*

Proof. By replacing R with $R[\frac{1}{b}]$ and D with $D[\frac{1}{b}]$, we may suppose $\frac{1}{b} \in R$.

Assume $p := \deg(f), q := \deg(g)$ and $a \neq 0$ is the leading coefficient of f. By assumption, $f = hg$, for some $h \in D[x]$.

We prove by induction on $\deg(h) \geq 0$ that $h \in R[x]$.

Let $c \neq 0$ be the leading coefficient of h. Note that $\deg(h) = p - q$. Then $cb = a$, and so $c = \frac{a}{b} \in R$. If $\deg(h) = 0$, we are done, otherwise $f - cx^{p-q}g = (h - cx^{p-q})g$, and so by induction $h - cx^{p-q} \in R[x]$. Thus the claim is proved. $\square$

Now let R be an integral domain, $f, g \in R[x]$ be non-zero polynomials and suppose $f = a_p x^p + \ldots + a_0, g = b_q x^q + \ldots + b_0$ with $a_p, b_q \neq 0$. The *Sylvester matrix* of f and g is the $(p+q) \times (p+q)$ matrix

$$S_{f,g} := \begin{pmatrix} a_p & \ldots & a_0 & & \\ & \ddots & & \ddots & \\ & & a_p & \ldots & a_0 \\ b_q & \ldots & b_0 & & \\ & \ddots & & \ddots & \\ & & b_q & \ldots & b_0 \end{pmatrix},$$

where the first q lines are formed by shifting the first row to the right, and the last p lines are formed by shifting the $(q+1)$th row to the right. If $p = q = 0$, we define $S_{f,g} = (1)$. The *resultant* of f and g, denoted by $\operatorname{res}(f, g)$, is the determinant of $S_{f,g}$. We also define $\operatorname{res}(0, h) = \operatorname{res}(h, 0) = 0$, for any polynomial h, so that the resultant is now properly defined for any two polynomials in $R[x]$.

The main application of resultants is to determine when two polynomials have a common root (see Theorem A.1). However, we will have to deal with more than two polynomials and more than one variable. We therefore make the following definition, following [64].

Let $f_1, \ldots, f_m \in R[x_1, \ldots, x_n], m \geq 1$. Let $y_3, \ldots, y_m$ be new indeterminates and define $R' := R[x_2, \ldots, x_n], R'' := R'[y_3, \ldots, y_m]$. Let F_1, F_2 be polynomials in $R''[x_1]$ defined as follows:

$$\begin{aligned} F_1 &:= f_1 \\ F_2 &:= f_2 + y_3 f_3 + \ldots + y_m f_m. \end{aligned} \tag{2.1}$$

If $m = 1$, we take $F_2 := 0$. We define the resultant of the polynomials $f_1, \ldots, f_m$ in terms of x_1, denoted by $\operatorname{res}_{x_1}(f_1, \ldots, f_m)$, as the resultant of F_1 and F_2. Note that this is a polynomial in $x_2, \ldots, x_n$ and $y_3, \ldots, y_m$.

Theorem 50. *Assume R is a field and let K be its algebraic closure. Let $(a_2, \ldots, a_n) \in K^{n-1}$ and suppose that the leading coefficient of x_1 in $f_1 \in R[x_1, \ldots, x_n]$, a polynomial in $x_2, x_3, \ldots, x_n$, does not vanish when replacing x_2 with a_2, x_3 with $a_3, \ldots, x_n$ with a_n. Then there exists an $a_1 \in K$ such that $(a_1, \ldots, a_n)$ is a common zero for $f_1, \ldots, f_m$ if and only if $\operatorname{res}_{x_1}(f_1, \ldots, f_m)(a_2, \ldots, a_n) = 0$.*

A proof of Theorem 50 can be found in the Appendix (see Theorem A.6).

We now turn to subresultants.

Let R be an integral domain, $f, g \in R[x]$ non-zero as before and again suppose $f = a_p x^p + \ldots + a_0, g = b_q x^q + \ldots + b_0$ with $a_p, b_q \neq 0$. The *subresultant sequence* for f and g is a list of polynomials $S_i(f, g) := \sum_{j=0}^{i} s_{ij}(f, g) x^j, 0 \leq i \leq \min\{p, q\}$, where $s_{ij}(f, g)$ is the determinant of the matrix built with rows $1, \ldots, q-i$ and $q+1, \ldots, q+p-i$ of $S_{f,g}$, and columns $1, 2, \ldots, p+q-2i-1, p+q-i-j$ of $S_{f,g}$. This is well-defined except when $i = p = q$. Thus when $p = q \neq 0$ we set $S_q(f, g) = g$ and define s_{qj} in the obvious way. For $p = q = 0$ we set $S_0(f, g) = 1$.

Due to technical reasons we define the subresultant sequence also for the case when one of f or g (but not both) is 0. If $g = 0$, we let $S_i(f, g) := S_i(f, f), 0 \leq i \leq \deg(f)$, and we proceed similarly if $f = 0$.

We now have the following result.

Theorem 51 (Theorem A.10). *Suppose R is a UFD and $f, g \in R[x]$ are not both zero. If $k \geq 0$ is minimal such that $s_{kk}(f, g) \neq 0$ then there exists non-zero $u, v \in R$ such that $u \gcd(f, g) = vS_k(f, g)$.*

In a similar form, Theorem 51 was already known in the 19th Century. Collins [27] introduced the terminology of subresultants, leading to the modern formulation of Theorem 51, in conjuction with the problem of efficiently computing the gcd of two polynomials. The theory was subsequently refined and simplified by Brown and Traub [18]. A proof of Theorem 51 can be found in the Appendix.

In the proof of the main result we will encounter rings which are not UFD, and so we will not be able to apply Theorem 51 directly. We deal with this situation below.

Let R be an integral domain and $f_1, \ldots, f_m \in R[x], m \geq 1$. We define F_1 and F_2 as in (2.1). We first make a simple observation.

Lemma 52. *Assume $R \subseteq K \subseteq \overline{K}$, where $K, \overline{K}$ are fields, and $\overline{K}$ is algebraically closed. Suppose $G := \gcd_K(f_1, \ldots, f_m)$ has degree $\delta \geq 1$, and let $b_1, \ldots, b_d$ be the distinct roots of G in $\overline{K}$, each appearing with multiplicity $\mu_i, 1 \leq i \leq d$. Then*

$$S_\delta(F_1, F_2) = \ell \prod_{i=1}^{d} (x - b_i)^{\mu_i}, \tag{2.2}$$

where ℓ is the leading coefficient of $S_\delta(F_1, F_2)$ as a polynomial in x.

As $\delta \geq 1$ we have $\deg(F_1) = \deg(f_1) \geq 1$ and so $S_i(F_1, F_2)$ is well-defined (nevertheless it may happen that F_2 is 0 if $m = 1$). Further recall that $S_i(F_1, F_2)$ is a polynomial in $y_3, \ldots, y_m$ and x.

Proof of Lemma 52. By Lemma A.5, $G = \gcd_{K[y_3,\ldots,y_m]}(F_1, F_2)$. Hence by Theorem 51, there are non-zero $u, v \in K[y_3, \ldots, y_m]$ such that $uG = vS_\delta(F_1, F_2)$. But for any $1 \leq i \leq d$, $(x - b_i)^{\mu_i} | uG$ in $\overline{K}[y_3, \ldots, y_m, x]$. Hence $(x - b_i)^{\mu_i} | S_\delta(F_1, F_2), 1 \leq i \leq d$. As $S_\delta(F_1, F_2)$ has degree exactly δ as a polynomial in x, (2.2) must hold, thus proving the lemma. □

The main consequence of Theorem 51 is the following.

Lemma 53. *Suppose $R \subseteq \mathbb{C}, G := \gcd_{\mathbb{C}}(f_1, \ldots, f_m)$ has degree $\delta \geq 1$, $\ell := s_{\delta\delta}(F_1, F_2)$ and $\phi : R \to \mathbb{F}_p$ is a homomorphism such that*

$$\deg_x(\phi(F_1)) = \deg_x(F_1), \quad \deg_x(\phi(F_2)) = \deg_x(F_2) \quad \text{and} \quad \phi(\ell) \neq 0. \tag{2.3}$$

Then for any root $b' \in \mathbb{F}_p$ of $\gcd_{\mathbb{F}_p}(\phi(f_1), \ldots, \phi(f_m))$ *there exists a root b of G and a homomorphism $\Phi : R[b] \to \mathbb{F}_p$ such that the following diagram commutes*

$$\begin{array}{ccc} R[x] & \xrightarrow{\mathrm{ev}_b} & R[b] \\ \phi\Big\downarrow & & \Big\downarrow\Phi \\ \mathbb{F}_p[x] & \xrightarrow{\mathrm{ev}_{b'}} & \mathbb{F}_p \end{array} \tag{2.4}$$

Proof. By definition the case $m = 1$ is equivalent to the case $m = 2$ where $f_2 = f_1$, and so we will assume w.l.o.g. that $m \geq 2$ and $F_2 \neq 0$. Let $G' := \gcd_{\mathbb{F}_p}(\phi(f_1), \ldots, \phi(f_m))$.

As $\deg_x(\phi(F_1)) = \deg_x(F_1)$ and $\deg_x(\phi(F_2)) = \deg_x(F_2)$, we have $\phi(S_i(F_1, F_2)) = S_i(\phi(F_1), \phi(F_2))$. Hence by Theorem 51 and the fact that $\phi(\ell) \neq 0$, we have $\deg(G) = \deg(G') = \delta \geq 1$.

Let b' be any root of G' in $\mathbb{F}_p$. By Lemma 52 we have $\phi(S_\delta(F_1, F_2))(b') = 0$. Define $\psi := ev_{b'} \circ \phi : R[x] \to \mathbb{F}_p$.

Let $b_1, \ldots, b_d$ be the distinct roots of G in $\mathbb{C}$, each appearing with multiplicity $\mu_i, 1 \leq i \leq d$. Assume for a contradiction that for any root b_i of G there is no homomorphism Φ making the diagram (2.4) commutative. This means $\ker \mathrm{ev}_{b_i} \not\subseteq \ker \psi$, so there exists a polynomial $g_i \in R[x]$ such that $g_i(b_i) = 0$, but $(\phi \circ g_i)(b') \neq 0$.

Define

$$H := \ell \prod_{i=1}^{d} g_i^{\mu_i}.$$

Then $H \in R[x, y_3, \ldots, y_m]$. As $\phi(\ell) \neq 0$, we have $\phi(H)(b') \neq 0$ in $\mathbb{F}_p[y_3, \ldots, y_m]$. But by Lemma 52,

$$S_\delta(F_1, F_2) = \ell \prod_{i=1}^{d} (x - b_i)^{\mu_i}$$

in $\mathbb{C}[x, y_3, \ldots, y_m]$. Then $S_\delta(F_1, F_2) | H$ in $\mathbb{C}[x, y_3, \ldots, y_m]$. Hence by Lemma 49, $S_\delta(F_1, F_2)|H$ in $R[x, y_3, \ldots, y_m, \frac{1}{\ell}]$. But $\phi(\ell) \neq 0$, so ϕ extends to a homomorphism

$$\phi : R[x, y_3, \ldots, y_m, \frac{1}{\ell}] \to \mathbb{F}_p[x, y_3, \ldots, y_m].$$

This implies $\phi(S_\delta(F_1, F_2))|\phi(H)$. As $\phi(S_\delta(F_1, F_2))(b') = 0$, we obtain $\phi(H)(b') = 0$, a contradiction. This finishes the proof of the lemma. □

2.2 Proof of Theorem 7

We have the following technical result.

Lemma 54. *Let $k, t \geq 2$ be integers and p be a prime. Suppose $A = \{a_1, \ldots, a_n\} \subseteq \mathbb{F}_p$ and let $\mathcal{L}_1, \mathcal{L}_2 \subset \mathbb{Z}[x_1, \ldots, x_n]$ be collections of (k,t)-bounded polynomials, such that any $f \in \mathcal{L}_1$ is zero when evaluated at $(a_1, \ldots, a_n)$, and any $f \in \mathcal{L}_2$ is non-zero when evaluated at $(a_1, \ldots, a_n)$. If*

$$|A| < \log_2 \log_{2t} \log_{2kt} p - 1 \tag{2.5}$$

then there exists a finite algebraic extension K of $\mathbb{Q}$ of degree at most $(2t)^{2^n}$ and a subset $A' = \{b_1, \ldots, b_n\} \subset K$ such that $f(b_1, \ldots, b_n) = 0$ for $f \in \mathcal{L}_1$, and $f(b_1, \ldots, b_n) \neq 0$ for $f \in \mathcal{L}_2$. Furthermore, the map $\phi_p : \mathbb{Z}[A'] \to \mathbb{F}_p$ sending b_i to a_i is a ring homomorphism.

Proof. We first give a rough overview of the proof.

The proof has three steps.

In the first step we eliminate the variables one by one. We start with the collection of polynomials $\mathcal{L}^0 := \mathcal{L}_1$ and we compute the resultant R_1 in terms of x_1. We then form a new collection of polynomials $\mathcal{L}^1$ in $x_2, \ldots, x_n$ by taking the coefficients of the y-monomials in R_1. By Theorem 50, there is at least one choice for x_1 iff there exists a common solution to the polynomials in $\mathcal{L}^1$. We then eliminate x_2 and proceed further in the same manner to construct collections $\mathcal{L}^i$. After at most n steps we have eliminated all variables, and only constant polynomials remain. However, the same procedure could have been carried over in $\mathbb{F}_p$, with the same starting collection of polynomials, and there it is guaranteed that a solution exists. Hence if the final constants are less than p, they must in fact be 0, and so a solution exists in $\mathbb{C}$ as well.

In the second step we go back, trying to determine the b_i's. Suppose for example that we have only polynomials in one variable, say x_n, and we know that a common root exists. Then their gcd is non-constant, and we can use Lemma 53 to pick one of the roots of the gcd as b_n. The hypothesis of Lemma 53 will be satisfied by adding some more polynomials to $\mathcal{L}^i$ in the first step. We then adjoin b_n to $\mathbb{Q}$, replace x_n by b_n, and proceed similarly to determine b_{n-1}. Theorem 50 will ensure that once $b_{i+1}, \ldots, b_n$ are picked, there is still a choice for b_i.

Note that once the homomorphism ϕ_p is constructed, the conditions imposed by $\mathcal{L}_2$ are automatically satisfied. For if $f \in \mathcal{L}_2$ then $\phi_p(f(b_1, \ldots, b_n)) = f(a_1, \ldots, a_n) \not\equiv 0 \,(\text{mod } p)$, hence $f(b_1, \ldots, b_n) \neq 0$ as well.

In the last step we will estimate the degree of the extension.

We now present the proof in detail.

Step 1. We let $u_0 := k, v_0 := t$ and for any $1 \leq i \leq n$ we define u_i and v_i inductively by

$$\begin{aligned} u_i &:= u_{i-1}^{2v_{i-1}} v_{i-1}^{v_{i-1}}, \\ v_i &:= 2v_{i-1}^2. \end{aligned}$$

We shall prove in Step 3 that for $0 \leq i \leq n$ we have

$$u_i < p. \tag{2.6}$$

Assume for the moment that this is indeed the case. For $0 \leq i \leq n$ let $\sigma_i : \mathbb{Z}[x_{i+1}, \ldots, x_n] \to \mathbb{F}_p[x_{i+1}]$ be the homomorphism mapping x_j to $a_j, i+1 < j \leq n$. We similarly define $\sigma : \mathbb{Z}[x_1, \ldots, x_n] \to \mathbb{F}_p$ as the homomorphism mapping x_j to a_j for all $1 \leq j \leq n$.

We will construct by induction on $i \geq 0$ sets $\mathcal{L}_1 = \mathcal{L}^0, \mathcal{L}^1, \ldots, \mathcal{L}^r, r \leq n$, such that $\mathcal{L}^i \subset \mathbb{Z}[x_{i+1}, \ldots, x_n]$ is a collection of (u_i, v_i)-bounded polynomials satisfying $\sigma(f) = 0$ for any $f \in \mathcal{L}^i, 0 \leq i \leq r$. Furthermore, it will be necessary at every step $i < r$ to slightly modify the set $\mathcal{L}^i$ into another one $\mathcal{A}_i$ by altering some of the polynomials. $\mathcal{A}_i$ will still contain only (u_i, v_i)-bounded polynomials f verifying $\sigma(f) = 0$.

The construction of the sets $\mathcal{L}^i$ will be done in three stages, indicated by the bold letters **(A), (B)** and **(C)**.

For $i = 0$, by assumption $\mathcal{L}^0$ is a collection of (u_0, v_0)-bounded polynomials mapped to 0 by σ.

Now suppose $n \geq i \geq 0$ and we have constructed $\mathcal{L}^i$. If $i = n$ or $\mathcal{L}^i$ is empty or $\{0\}$, we set $r = i$ and stop. Otherwise, let $\mathcal{L}^i = \{f_1, \ldots, f_m\}$ and $f_j = \sum_{\ell=0}^{d_j} c_{j\ell} x_{i+1}^{\ell}$. By assumption we have $i \leq n-1$.

(A) For any $1 \leq j \leq m$ and $\deg_{x_{i+1}}(\sigma_i(f_j)) < \ell \leq \deg_{x_{i+1}}(f_j)$ we put $c_{j\ell}$ into $\mathcal{L}^{i+1}$.

We then set $d_j' = \deg_{x_{i+1}}(\sigma_i(f_j))$ and define

$$f_j' := \sum_{\ell=0}^{d_j'} c_{j\ell} x_{i+1}^{\ell}.$$

Note that $d_j' \neq 0$, otherwise $\sigma(f_j) = \sigma_i(f_j) \neq 0$, a contradiction.

Let $\mathcal{A}_i := \{f_1', \ldots, f_m'\}$. Clearly every polynomial in $\mathcal{A}_i$ is still (u_i, v_i)-bounded. Furthermore if x_{i+1} does not appear in any polynomial in $\mathcal{A}_i$, then $\mathcal{A}_i$ contains only 0 by the above. In this case there is nothing else to be done.

So assume w.l.o.g. that x_{i+1} appears in f_1'. Let

$$\begin{aligned} F_1 &:= f_1' \\ F_2 &:= f_2' + y_3 f_3' + \ldots + y_m f_m' \end{aligned}$$

for unknowns $y_3, \ldots, y_m$, where $F_2 := 0$ if $m = 1$.

(B) We take $\mathrm{res}_{x_{i+1}}(f_1', \ldots, f_m')$ and put into $\mathcal{L}^{i+1}$ the coefficient of every monomial in y_j, which must be a polynomial in $x_{i+2}, \ldots, x_n$.

Set $R_1 := \mathbb{Z}[x_{i+2}, \ldots, x_n, y_3, \ldots, y_m]$ and $R_2 := \mathbb{F}_p[y_3, \ldots, y_m]$. Note that σ_i induces a homomorphism between $R_1[x_{i+1}]$ and $R_2[x_{i+1}]$. We have $F_1, F_2 \in R_1[x_{i+1}]$ and by (A), $\deg_{x_{i+1}}(\sigma_i(F_1)) = \deg_{x_{i+1}}(F_1)$ and $\deg_{x_{i+1}}(\sigma_i(F_2)) = \deg_{x_{i+1}}(F_2)$. So let $q_1 := \deg_{x_{i+1}}(F_1)$ and $q_2 := \deg_{x_{i+1}}(F_2)$. By assumption, $q_1 \geq 1$.

Let $\delta \geq 0$ be minimal such that $\sigma_i(s_{\delta\delta}(F_1, F_2)) \neq 0$, where $s_{k\ell}$ are the coefficients of the subresultant sequence.

(C) We put into $\mathcal{L}^{i+1}$ the coefficients of $s_{jj}(F_1, F_2)$ (polynomials in $x_{i+2}, \ldots, x_n$), for $1 \leq j < \delta$. For $j = 0$ this has already been done, as $s_{00}(F_1, F_2) = \mathrm{res}_{x_{i+1}}(f_1', \ldots, f_m')$ by definition.

The construction of $\mathcal{L}^{i+1}$ is now over. We must show that any polynomial in $\mathcal{L}^{i+1}$ is indeed (u_{i+1}, v_{i+1})-bounded.

This is certainly the case for the polynomials added in stage (A). So consider the stage (B) of the construction.

Fix an arbitrary monomial M in $y_3, \ldots, y_m$ of degree at most q_1. This has a coefficient g in $\mathrm{res}_{x_{i+1}}(f'_1, \ldots, f'_m)$ and we must estimate $\|g\|_1$ and $\deg(g)$. Since $q_1, q_2 \leq v_i$, the degree of g is at most $(q_1 + q_2)v_i \leq 2v_i^2 = v_{i+1}$, as desired.

Now let $2 \leq j_1, j_2, \ldots, j_{q_1} \leq m$ and define $S_{F_1,F_2}(j_1, \ldots, j_{q_1})$ by writing on line $q_2 + k'$ of S_{F_1,F_2}, instead of the coeficients of F_2, the corresponding coefficients of $f'_{j_{k'}}, 1 \leq k' \leq q_1$. Then g is a sum of $\det(S_{F_1,F_2}(j_1, \ldots, j_{q_1}))$, for certain q_1-tuples $j_1, j_2, \ldots, j_{q_1}$ depending on M. The number of such q_1-tuples is

$$\binom{q_1}{\deg(M)} \frac{\deg(M)!}{\deg_{y_3}(M)! \ldots \deg_{y_m}(M)!} \leq q_1^{\deg(M)} \leq v_i^{v_i}.$$

Recall that $\|f'_j\|_1 \leq u_i, 1 \leq j \leq m$. So by Lemma 46 applied to $S_{F_1,F_2}(j_1, \ldots, j_{q_1})$ (a square matrix of size $q_1 + q_2 \leq 2v_i$), we obtain $\|\det(S_{F_1,F_2}(j_1, \ldots, j_{q_1}))\|_1 \leq u_i^{2v_i}$. Hence $\|g\|_1 \leq u_i^{2v_i} v_i^{v_i} = u_{i+1}$, as desired.

Finally, as subresultants are defined using submatrices of S_{F_1,F_2}, all the above estimates apply to subresultants as well. Hence any polynomial added to $\mathcal{L}^{i+1}$ in stage (C) is also (u_{i+1}, v_{i+1})-bounded. Consequently any polynomial in $\mathcal{L}^{i+1}$ is (u_{i+1}, v_{i+1})-bounded, as claimed.

We must further check that σ maps all the polynomials in $\mathcal{L}^{i+1}$ to 0. This is certainly the case with the polynomials added in stages (A) and (C) of the construction. As $\deg_{x_{i+1}}(\sigma_i(f'_j)) = \deg_{x_{i+1}}(f'_j), 1 \leq j \leq m$, we have that $\mathrm{res}_{x_{i+1}}(\sigma_i(f'_1), \ldots, \sigma_i(f'_m)) = \sigma_i(\mathrm{res}_{x_{i+1}}(f'_1, \ldots, f'_m))$. By Theorem 50 and the fact that the polynomials $\sigma_i(f'_j)$ have the common root a_{i+1}, we obtain $\mathrm{res}_{x_{i+1}}(\sigma_i(f'_1), \ldots, \sigma_i(f'_m)) = 0$. This shows that all the polynomials added in stage (B) of the construction are indeed mapped to 0 by σ. Thus the induction step is verified.

Step 2. If $\mathcal{L}^r$ is empty, all the sets $\mathcal{L}^i$ were empty, in particular $\mathcal{L}^0 = \mathcal{L}_1 = \emptyset$. Then we set $b_i = a_i, 1 \leq i \leq n$, take ϕ_p to be the canonical homomorphism, and we are done.

So we may assume that $\mathcal{L}^r$ is non-empty. Let $f \in \mathcal{L}^r$. By construction f is an integer constant at most u_r in absolute value, and $u_r < p$ by (2.6). However, $\sigma(f) = 0$, and as σ is a homomorphism, we must have $f = 0$. Hence $\mathcal{L}^r = \{0\}$.

By decreasing induction on $r \geq i \geq 0$ we shall find algebraic numbers $b_{i+1}, \ldots, b_n$ such that for any $f \in \mathcal{L}^i$, $f(b_{i+1}, \ldots, b_n) = 0$, and furthermore the map $\phi_p^i : \mathbb{Z}[b_{i+1}, \ldots, b_n] \to \mathbb{F}_p$, sending b_j to $a_j, i < j \leq n$, is a well-defined homomorphism.

For any $j > r$, we let b_j be the integer in $\{0, 1, \ldots, p-1\}$ satisfying $b_j \equiv a_j \pmod{p}$. Then $\phi_p^r = \sigma\big|_{\mathbb{Z}}$ is a homomorphism. As $\mathcal{L}^r = \{0\}$, the base case $i = r$ is verified.

Now assume $0 \leq i < r$ and we have found $b_{i+2}, \ldots, b_n$ satisfying the induction hypothesis.

Suppose $\mathcal{L}^i = \{f_1, \ldots, f_m\}$ and $\mathcal{A}_i = \{f'_1, \ldots, f'_m\}$. We replace $x_{i+2}, \ldots, x_n$ with their values b_j in the polynomials $f_1, \ldots, f_m$ and $f'_1, \ldots, f'_m$. By (A), $f_j = f'_j$ and furthermore $\deg_{x_{i+1}}(\phi_p^{i+1}(f_j)) = \deg_{x_{i+1}}(f_j), 1 \leq j \leq m$. If x_{i+1} does not appear in any of these polynomials, then all of them are in fact 0. In this case we let b_{i+1} be the integer in $\{0, 1, \ldots, p-1\}$ satisfying $b_{i+1} \equiv a_{i+1} \pmod{p}$. We have $\phi_p^{i+1}(b_{i+1}) = a_{i+1}$. Thus $\phi_p^i = \phi_p^{i+1}$ is a well-defined homomorphism, and the claim holds.

So assume x_{i+1} appears in f_1. Here we use the same indexing scheme as in Step 1; in particular, f_1 corresponds to the polynomial f_1' selected in Step 1.

By (B) and Theorem 50, at least one choice b_{i+1} for x_{i+1} exists, such that replacing x_{i+1} with this value vanishes all polynomials in $\mathcal{L}^i$. In other words, $G := \gcd_{\mathbb{C}}(f_1, \dots, f_m)$ has degree $\delta \geq 1$.

Now recall our construction of F_1 and F_2. By (A), $\deg_{x_{i+1}}(\phi_p^{i+1}(F_2)) = \deg_{x_{i+1}}(F_2)$. Let $\ell := s_{\delta\delta}(F_1, F_2)$. By (C), Lemma A.5 and Theorem 51 applied to F_1 and F_2 in $\mathbb{C}[x_{i+1}, y_3, \dots, y_m]$, we see that $\phi_p^{i+1}(\ell) \neq 0$.

Hence the hypothesis of Lemma 53 is satisfied for the ring $A := \mathbb{Z}[b_{i+2}, \dots, b_n]$, the polynomials $f_1, \dots, f_m$ and the homomorphism $\phi := \phi_p^{i+1}$. This implies that for the root a_{i+1} of $\gcd_{\mathbb{F}_p}(\phi_p^{i+1}(f_1), \dots, \phi_p^{i+1}(f_m))$ there exists a root b_{i+1} of G and a homomorphism $\phi_p^i : \mathbb{Z}[b_{i+1}, \dots, b_n] \to \mathbb{F}_p$ making the diagram (2.4) commutative. Then ϕ_p^i still maps b_j to a_j for $i+1 < j \leq n$. Furthermore by construction, replacing x_{i+1} with b_{i+1} in the polynomials in $\mathcal{L}^i$ vanishes all of them. This proves the induction step.

Continuing in this way we obtain all algebraic numbers $b_1, \dots, b_n$ and in the last step $\phi_p := \phi_p^0$ maps b_j to a_j as desired.

Step 3. We now compute the degree of the extension and verify (2.6).

First note that $r \leq n$ and $v_i = 2^{2^i-1}t^{2^i}, 0 \leq i \leq n$. Then the degree of the extension is at most

$$\prod_{i=0}^{r-1} v_i \leq \prod_{i=0}^{n-1} 2^{2^i-1}t^{2^i} \leq 2^{2^n-(n+1)}t^{2^n} \leq (2t)^{2^n}.$$

Further note that

$$\prod_{i=0}^{n-1} 2v_i \leq 2^{2^n-1}t^{2^n}.$$

We also have $u_0 = k$ and

$$u_{i+1} = u_i^{2v_i} v_i^{v_i}, \tag{2.7}$$

and so by iterating (2.7), and using the above estimates, we obtain

$$
\begin{aligned}
u_n &= u_{n-1}^{2v_{n-1}} v_{n-1}^{v_{n-1}} \\
&= u_{n-2}^{(2v_{n-2})(2v_{n-1})} v_{n-2}^{v_{n-2}(2v_{n-1})} v_{n-1}^{v_{n-1}} \\
&= \dots \\
&= \exp\left\{ \left(\prod_{i=0}^{n-1} 2v_i \right) \log k + \sum_{i=0}^{n-1} v_i (2v_{i+1}) \dots (2v_{n-1}) \log v_i \right\} \\
&\leq \exp\left\{ \left(\prod_{i=0}^{n-1} 2v_i \right) (\log k + n \log v_{n-1}) \right\} \\
&\leq \exp\left\{ 2^{2^n-1} t^{2^n} (\log k + n \log (2t)^{2^{n-1}}) \right\} \\
&\leq k^{2^{2^n} t^{2^n}} (2t)^{n 2^{n-1} 2^{2^n-1} t^{2^n}} \\
&\leq k^{(2t)^{2^n}} (2t)^{2^{2^n+2n-2} t^{2^n}} \\
&\leq k^{(2t)^{2^n}} (2t)^{(2t)^{2^{n+1}}} \\
&\leq (2kt)^{(2t)^{2^{n+1}}}.
\end{aligned}
$$

Thus the condition $u_n < p$ is satisfied if $n < \log_2 \log_{2t} \log_{2kt} p - 1$. This shows that (2.6) holds, and hence the proof is finished. $\square$

Proof of Theorem 7. We consider all k-bounded polynomials in $n := |A|$ variables, and we split them into $\mathcal{L}_1$ and $\mathcal{L}_2$ according to the result of evaluation with elements from A. Applying Lemma 54, we get a finite algebraic extension K of $\mathbb{Q}$ of degree at most $(2k)^{2^n}$, a subset $A' \subset K$ and a homomorphism $\phi_p : \mathbb{Z}[A'] \to \mathbb{F}_p$ which by definition is an F_k-ring-isomorphism between A' and A. This proves the theorem. $\square$

2.3 Sharpness of the main result

In this section we prove Theorem 8. For $k \geq 2, t \geq 1$ we say that a positive integer r is (k,t)-*constructible in at most n steps* if there exists a sequence of non-negative integers $0 = a_0, a_1, \ldots, a_m = r, m \leq n$, such that for any $i \geq 1, a_i = f_i(a_0, \ldots, a_{i-1})$, with $f_i \in \mathbb{Z}[x_0, \ldots, x_{i-1}]$ a (k,t)-bounded polynomial.

The main step is to prove the following lemma.

Lemma 55. *Let $k \geq 2$. Any $p \geq 2^{32(k\log_2(16k))^2}$ is (k,k)-constructible in at most $\frac{10}{k}\frac{\log_2 p}{\log_2\log_2 p}$ steps, and moreover this is sharp up to a constant not depending on k.*

Proof. Let $p \geq 2^{32(k\log_2(16k))^2}$ arbitrary. We first note the following inequality:

$$\log_2\log_2 p \geq 2\log_2(k\log_2\log_2 p). \tag{2.8}$$

Indeed, this is true if $\log_2 p \geq k^2(\log_2\log_2 p)^2$, which in turn is true if $\log p \geq \frac{2k^2}{\log 2}(\log\log p)^2$. By derivation this holds whenever $p \geq 2^{32(k\log_2(16k))^2} \geq e^{8(k\log_2(16k))^2}$.

Now set

$$s := \left\lceil \log_2\left(\frac{\log_2 p}{k\log_2\log_2 p}\right)\right\rceil \quad \text{and} \quad N := \lfloor \log_2 p \rfloor .$$

Note that $s \geq 1$, as $\log_2 p > k\log_2\log_2 p$ by (2.8).

Consider the base-2 representation $(b_0 b_1 \ldots b_N)$ of p, with b_0 being the least significant bit. We break it into $\ell := \lceil \frac{N+1}{sk} \rceil \geq 1$ contiguous subsequences

$$(b_0 b_1 \ldots b_{sk-1}), \ldots, (b_{(\ell-1)sk} b_{(\ell-1)sk+1} \ldots b_N),$$

all of them except possibly the last one of length sk, defining in base-2 numbers $p_0, p_1, \ldots, p_{\ell-1}$. Note that

$$p = \sum_{i=0}^{\ell-1} 2^{ski} p_i$$

and $p_i < 2^{sk}, 0 \leq i < \ell$. We further write

$$p_i = \sum_{j=0}^{k-1} 2^{sj} p_{ij},$$

with $0 \leq p_{ij} < 2^s$.

We now define the sequence $a_0, \ldots, a_{2^s+\ell+2(\ell-1)}$ as follows.

We start by setting $a_0 := 0$ and $a_i := a_{i-1} + 1, 1 \leq i \leq 2^s$. Note that $a_i = i, 1 \leq i \leq 2^s$. For any $0 \leq i \leq \ell - 1$ we define

$$a_{2^s+1+i} := \sum_{j=0}^{k-1} a_{2^s}^j a_{p_{ij}}.$$

Hence $a_{2^s+1+i} = p_i$. For any $1 \le i \le \ell - 1$ we further define $a_{2^s+\ell+2(i-1)+1}$ and $a_{2^s+\ell+2(i-1)+2}$ as follows:

$$a_{2^s+\ell+2(i-1)+1} := \begin{cases} a_{2^s}^k, & \text{if } i = 1, \\ a_{2^s+\ell+2(i-2)+1}a_{2^s+\ell+1}, & \text{otherwise.} \end{cases}$$

$$a_{2^s+\ell+2(i-1)+2} := \begin{cases} a_{2^s+\ell+1}a_{2^s+2} + a_{2^s+1}, & \text{if } i = 1, \\ a_{2^s+\ell+2(i-1)+1}a_{2^s+i+1} + a_{2^s+\ell+2(i-2)+2}, & \text{otherwise.} \end{cases}$$

Hence

$$\begin{aligned} a_{2^s+\ell+2(i-1)+1} &= 2^{ski}, \\ a_{2^s+\ell+2(i-1)+2} &= \sum_{j=0}^{i} 2^{skj} p_j. \end{aligned}$$

In particular, $a_{2^s+\ell+2(\ell-1)} = p$. Hence p is (k,k)-constructible in at most $2^s + 3\ell - 2$ steps. But

$$\begin{aligned} 2^s + 3\ell - 2 &\le 2^s + 1 + 3\frac{N+1}{sk} \\ &\le 2^s + 4\frac{N}{sk}, \quad \text{as } sk + 3N + 3 \le 4N, \\ &\le \frac{2}{k}\frac{\log_2 p}{\log_2\log_2 p} + \frac{4}{k}\frac{\log_2 p}{\log_2\log_2 p - \log_2(k\log_2\log_2 p)} \\ &\le \frac{10}{k}\frac{\log_2 p}{\log_2\log_2 p}, \quad \text{by (2.8).} \end{aligned}$$

This proves the first part of the lemma. To show that this bound is essentially best possible, we fix n and count the number of positive integers (k,k)-constructible in at most n steps.

First note that for given $\ell \ge 1$, the number of monomials in ℓ variables $x_1, \ldots, x_l$ of degree at most k is $\binom{\ell+k}{k} \le (kl)^k$. Hence the number of (k,k)-bounded polynomials in ℓ variables is at most $3^k\binom{\ell+k}{k} \le (3k\ell)^k$, as any such polynomial is a sum of k monomials in ℓ variables of degree at most k, with coefficients $1, -1$ or 0.

Now to any number which is (k,k)-constructible in at most n steps corresponds a sequence of (k,k)-bounded polynomials $f_1, \ldots, f_m, m \le n$, such that f_i is a polynomial in i variables. Thus the number of integers (k,k)-constructible in at most n steps is upper bounded by the number of such sequences, which for $n \ge 3k$ is at most

$$\prod_{i=1}^{n}(3ki)^k \le (3k)^{kn}n^{kn} \le n^{2kn}.$$

However if p is given, then for $n \le \frac{\log p}{2k\log\log p}$ we have

$$n^{2kn} \le \left(\frac{\log p}{2k\log\log p}\right)^{\frac{\log p}{\log\log p}} < p.$$

Hence not all numbers between 1 and p are (k,k)-constructible in at most $\frac{\log p}{2k \log\log p}$ steps. This finishes the proof of the lemma. $\square$

Proof of Theorem 8. Given $p \geq 2^{32(k-1)^2 \log_2^2(16(k-1))}$ a prime number, we apply Lemma 55 to find a sequence of non-negative integers $0 = a_0, \ldots, a_n = p, n \leq \frac{10}{k-1} \frac{\log_2 p}{\log_2 \log_2 p}$, which shows that it is $(k-1, k-1)$-constructible. Let $A' := \{a_0, a_1, a_2, \ldots, a_n\}$. Taking the residues modulo p of the numbers in A' we obtain a set $A \subseteq \mathbb{F}_p$ of size at most n.

Now suppose for a contradiction that there exists an F_k-ring-isomorphism ϕ of A into an integral domain R of characteristic 0.

There is a natural embedding of $\mathbb{Z}$ into R, and we can identify $\mathbb{Z}$ with the image of this embedding. Let $x_i \in A$ be the image of a_i in $\mathbb{F}_p, 0 \leq i \leq n$. By induction on $i \geq 0$ we see that $\phi(x_i)$ must equal a_i.

This is certainly the case for $x_0 = 0$. For $i \geq 1$ there exists a $(k-1)$-bounded polynomial f_i such that $a_i = f_i(a_0, \ldots, a_{i-1})$. Hence $f_i(x_0, \ldots, x_{i-1}) - x_i = 0$ in $\mathbb{F}_p$. As this is a k-bounded polynomial, it must be preserved by ϕ. Therefore the induction hypothesis implies $\phi(x_i) = a_i$, as claimed.

However, A has size at most n, while A' has size $n+1$. Therefore the image of ϕ can not contain the whole of A', a contradiction. This proves the theorem. $\square$

The proof of Lemma 55 tells us that for given $M \geq 1$ there are only $(\log M)^{O(\log\log\log M)}$ positive integers less than M which are $(2,2)$-constructible in $O(\log\log M)$ steps. Nevertheless any Mersenne prime (a prime p of the form $2^n - 1$) is $(2,2)$-constructible in $O(\log n) = O(\log\log p)$ steps, by using the base-2 representation of n and an approach similar to that of Lemma 55. Furthermore any Fermat prime (a prime p of the form $2^{2^n} + 1$) is $(2,2)$-constructible in $O(n) = O(\log\log p)$ steps. Thus the existence of infinitely many such primes would imply Conjecture 9. Unfortunately proving or disproving such a statement seems at present to be an unreachable goal.

2.4 Proof of Theorems 12, 14, 18 and 19

Proof of Theorem 12. We may assume w.l.o.g. that $|\mathcal{P}| = |\mathcal{L}| = n$, by adding some points and lines if necessary. Let $\mathcal{P} = \{(x_i, y_i) : 1 \leq i \leq n\}$. By uniquely parametrizing each line $l \in \mathcal{L}$ defined by $a_i y + b_i x + c_i = 0$, by the ordered triple (a_i, b_i, c_i), let $\mathcal{L} = \{(a_i, b_i, c_i) : 1 \leq i \leq n\}$. Now form the set $A := \cup_{i=1}^{n} \{x_i, y_i, a_i, b_i, c_i\}$. As $|A| \leq 5n$, we may apply Theorem 7 to find a subset $A' \subset \mathbb{C}$ and an F_3-ring-isomorphism ϕ between A and A'. By definition we have

$$a_j y_i + b_j x_i + c_j = 0 \Leftrightarrow \phi(a_j)\phi(y_i) + \phi(b_j)\phi(x_i) + \phi(c_j) = 0, \forall 1 \leq i, j \leq n,$$

hence the number of incidences between $\mathcal{P}$ and $\mathcal{L}$ in $\mathbb{F}_p^2$ is the same as the number of incidences between $\phi(\mathcal{P})$ and $\phi(\mathcal{L})$ in $\mathbb{C}$. Note that $\phi(\mathcal{P})$ and $\phi(\mathcal{L})$ have cardinality exactly n as ϕ is bijective. Hence by Theorem 10, the number of incidences is $O(n^{4/3})$, as desired.

To show that the bound is sharp, we use a standard construction that proves sharpness of the Szemerédi-Trotter theorem in $\mathbb{R}^2$. Let $r := \lfloor \frac{1}{2} n^{1/3} \rfloor$. We set $\mathcal{P}$ to be the points of the lattice $[r] \times [2r^2]$ in $\mathbb{F}_p^2$, and $\mathcal{L}$ to be all lines $y = mx + b$, with $1 \leq m \leq r, 1 \leq b \leq r^2$. Then every line from $\mathcal{L}$ is incident with exactly r points from $\mathcal{P}$, for a total of $r^4 = \Theta(n^{4/3})$ incidences. □

Proof of Theorem 14. We apply Theorem 7 to find a subset $A' \subset \mathbb{C}$ and an F_4-ring-isomorphism ϕ between A and A'. Then $|\phi(A) + \phi(A)| = |A + A|$ and $|\phi(A) \cdot \phi(A)| = |A \cdot A|$. By (1.1) applied to $A' = \phi(A)$, the theorem follows. □

Proof of Theorem 18. We first prove (i).

We apply Theorem 7 to find a subset $A' \subset \mathbb{C}$ and an F_4-ring-isomorphism ϕ between $A \cup A^{-1}$ and A'. Then $|\phi(A)| = n$, $|\phi(A) + \phi(A)| = |A + A|$ and $|\phi(A^{-1}) + \phi(A^{-1})| = |A^{-1} + A^{-1}|$. Moreover, all identities of the form $a^{-1}a = 1, a \in A$, must be preserved by the ring-isomorphism, and hence $\phi(a^{-1}) = \phi(a)^{-1}, \forall a \in A$. Then by applying Theorem 17, (i), the result follows.

We now prove (ii).

We apply Theorem 7 to find a subset $A' \subset \mathbb{C}$ and an F_{4k}-ring-isomorphism ϕ between A and A'. Then $|\phi(A)| = n$ and $|\phi(A) + \phi(A)| = |A + A|$. We further have

$$f(\phi(a)) + f(\phi(b)) - f(\phi(c)) - f(\phi(d)) = 0 \Leftrightarrow f(a) + f(b) - f(c) - f(d) = 0,$$

for any $a, b, c, d \in A$, as ϕ is an F_{4k}-ring-isomorphism. Hence $|f(\phi(A)) + f(\phi(A))| = |f(A) + f(A)|$. Then by applying Theorem 17, (ii), the result follows. □

Proof of Theorem 19. Set $s := \lfloor \frac{k+1}{2} \rfloor$. Note that $s \geq 2$.

Let $f \in \mathbb{C}[x]$ be a polynomial with k non-zero terms minimizing $N(f^2)$. Suppose $f = a_0 + a_1 x^{n_1} + \ldots + a_{k-1} x^{n_{k-1}}$ and set $A := \{a_0, \ldots, a_{k-1}\} \subset \mathbb{C}$.

We now apply Theorem 4 in order to find a sufficiently large prime p (compared to k) and a homomorphism $\phi : \mathbb{Z}[A] \to \mathbb{F}_p$ which is an F_s-ring-isomorphism between A and $\phi(A)$. We then apply Theorem 7 to the set $\phi(A)$ in order to find a finite algebraic extension K of $\mathbb{Q}$ of degree at most $(2s)^{2^k}$, a subset $B \subset K$ and a map ψ between $\phi(A)$ and B, which is an F_s-ring-isomorphism. Then $\psi \circ \phi$ is an F_s-ring-isomorphism between A and B by construction.

Let $g = (\psi \circ \phi)(a_0) + (\psi \circ \phi)(a_1)x^{n_1} + \ldots + (\psi \circ \phi)(a_{k-1})x^{n_{k-1}}$. Then $g \in K[x]$ and $N(g) = k$. As any coefficient of g^2 is given by a polynomial with integer coefficients of degree at most 2 and $\|\cdot\|_1$-norm at most s, evaluated at $((\psi \circ \phi)(a_0), (\psi \circ \phi)(a_1), \ldots, (\psi \circ \phi)(a_{k-1}))$, we see that $N(g^2) = N(f^2)$. Consequently $Q_K(k) \leq N(f^2) = Q_{\mathbb{C}}(k)$, thus proving the theorem. $\square$

2.5 Remarks on the main result

The proof of Theorem 7 uses elimination theory. This is not the first time when elimination theory is applied to additive combinatorics: similar techniques were used by Chang in the proof of Lemma 2.14 from [20]. We state this lemma below in an equivalent form.

Lemma 56 (Lemma 2.14, [20]). *Let $f_1, \ldots, f_s \in \mathbb{Z}[x_1, \ldots, x_n]$ be polynomials of degree at most t and $\|\cdot\|_\infty$-norm at most k. If the system*

$$f_1(x) = \ldots = f_s(x) = 0$$

has a solution $(a_1, \ldots, a_n) \in \mathbb{C}^n$, then it also has a solution $(b_1, \ldots, b_n)$, where each b_i is the root of an integer polynomial of degree at most C and $\|\cdot\|_\infty$-norm at most Ck^C, with $C := C(t, n, s)$ depending only on t, n and s.

This lemma is discussed by Tao on his blog [126], in particular he gives a proof of it using nonstandard analysis. Neither this proof nor the proof in [20] provides a bound on the constant C.

The proof of Lemma 56 from [20] shows in fact a bit more; namely that if we are further given a polynomial $g \in \mathbb{Z}[x_1, \ldots, x_n]$ which does not vanish at $(a_1, \ldots, a_n)$, and has degree at most t and $\|\cdot\|_\infty$-norm at most k, then it is possible to choose $(b_1, \ldots, b_n)$ such that $g(b_1, \ldots, b_n) \neq 0$. On close examination of the proof it turns out that translated into the correct setting it implies the following weak version of Theorem 7.

Theorem 57. *For any $k \geq 2$ there exists a function $\nu_k : \mathbb{N} \to \mathbb{N}$ with $\lim_{n\to\infty} \nu_k(n) = \infty$, such that the following holds. If p is a prime and $A \subseteq \mathbb{F}_p$ with $|A| \leq \nu_k(p)$ then there exists a finite algebraic extension K of $\mathbb{Q}$ and a subset $A' \subset K$ such that A' is F_k-ring-isomorphic with A.*

An upper bound for the constant C implies a lower bound for $\nu_k(n)$; however from the proof of Lemma 56 one can only extract a rather poor bound for C.

It is also important to note that Theorem 57 does not provide any bound on the degree of the field extension K, nor does it guarantee that the F_k-ring-isomorphism is the restriction of a genuine ring homomorphism, as in Theorem 7. In fact it is easy to construct an example of a Freiman ring-isomorphism ϕ between a subset $A' \subset \mathbb{C}$ and a subset $A \subset \mathbb{F}_p$ such that ϕ is not the restriction of any ring homomorphism between $\mathbb{Z}[A']$ and $\mathbb{F}_p$. Indeed, consider $A' := \{-\frac{1}{2}, 2\} \subset \mathbb{C}$ and $A := \{3, 7\} \subset \mathbb{F}_{11}$. The map ϕ sending $-\frac{1}{2}$ to 3 and 2 to 7 is an F_2-ring-isomorphism, but it is obviously not the restriction of a ring homomorphism between $\mathbb{Z}[A']$ and $\mathbb{F}_{11}$ (as any such homomorphism would send 2 to 2). Examples for arbitrarily large k and p can be constructed as well.

Remark. After submitting the results presented in this chapter, I was informed by Pierre Simon that one can use the arithmetic Nullstellensatz stated in [79] to prove a good lower bound for the function ν_k in Theorem 57. With his idea, my own computations show that one can take $\nu_k(p) = \Omega(\frac{\log\log p}{\log\log\log p})$. This would improve the upper bound for n in Theorems 12, 14, 16 and 18 to $O(\frac{\log\log p}{\log\log\log p})$.

Remark. In his blog post *Rectification and the Lefschetz principle* [127], Tao presented a short proof of the following version of Theorem 7.

Theorem 58. *Let $k, n \geq 1$. If $\mathbb{F}$ is a field of characteristic at least $C_{k,n}$ for some $C_{k,n}$ depending only on k and n, and A is a subset of $\mathbb{F}$ of cardinality n, then there exists a map $\phi : A \to A'$ into a subset A' of the complex numbers which is a Freiman ring-isomorphism of order k.*

The proof uses non-standard analysis, and hence does not offer any bound on $C_{k,n}$. However, unlike Theorem 7, it also applies to fields of prime power order.

Remark. Theorem 8 does not cover the case $k = 2$, and in fact here I believe, but can not prove, that the correct bound is $\Theta(\log p)$; that is, any subset $A \subseteq \mathbb{F}_p$ of size $O(\log p)$ is F_2-ring-isomorphic to a subset of $\mathbb{C}$. Neither the proof of Theorem 2 nor that of Lemma 47 properly adapt to this situation, as one would have to work over the multiplicative group $\mathbb{F}_p^*$ of order $p - 1$.

Remark. Lemma 54 implies the following weaker version of Lemma 56: under the hypothesis of Lemma 56, there exists a solution $(b_1, \ldots, b_n) \in K^n$ to the polynomials $f_1, \ldots, f_s$, where K is a finite algebraic extension of $\mathbb{Q}$ of degree at most $(2t)^{2^n}$. Indeed, suppose each f_i has degree at most t and $\| \cdot \|_\infty$-norm at most k. Then each f_i is $(k(nt)^t, t)$-bounded. Fix $A := \{a_1, \ldots, a_n\}$, the coordinates of a complex solution of the system of polynomials $\{f_i : 1 \leq i \leq s\}$. We first apply Theorem 4 in order to find a sufficiently large prime p (compared to n, k and t) and a homomorphism $\phi : \mathbb{Z}[A] \to \mathbb{F}_p$. We then apply Lemma 54 to the collections $\mathcal{L}_1 := \{f_1, \ldots, f_s\}$ and $\mathcal{L}_2 := \emptyset$, in order to find a finite algebraic extension K of degree at most $(2t)^{2^n}$, a subset $A' \subset K$ and a map ψ between $\phi(A)$ and A'. Then $((\psi \circ \phi)(a_i))_{i=1}^n$ are the coordinates of a solution $(b_1, \ldots, b_n) \in K^n$ of the system of polynomials $\{f_i : 1 \leq i \leq s\}$.

Remark. In view of Theorem 12 one may ask what is the largest number $n(p)$ of points and lines in $\mathbb{F}_p^2$ for which the upper bound $cn(p)^{4/3}$ on the number of incidences holds. I have only proved $n(p) = \Omega(\log \log \log p)$, and I am not aware of any non-trivial upper bound for this function.

Chapter 3

Turán densities of hypergraphs

3.1 Useful notation

We introduce some notation needed in the sequel.

An *r-multiset* D is an unordered collection of r elements $x_1, \ldots, x_r$ with repetitions allowed. The multiplicity $D(x)$ of x in D is the number of times that x appears in D.

A pair $G = (V, E)$ with $E \subseteq V^{(r)}$ is called an *r-multigraph*. V is the set of *vertices* and E the set of *edges*. Note that every edge is an r-multiset. Furthermore, note that our definition is different from the usual definition of a multigraph where the same edge may appear multiple times in E. If all edges in E are proper sets, then G is called a (simple) r-graph. We let $v(G) := |V(G)|$ be the number of vertices and $e(G) := |E(G)|$ be the number of edges. The *density* of an r-graph G is

$$d(G) = \frac{e(G)}{\binom{n}{r}}.$$

We do not define the density of an r-multigraph.

We allow graphs without edges, and we also consider $\emptyset$ to be an r-graph without vertices. We call $\emptyset$ *the empty graph*.

If G and H are r-graphs, we say H is a *subgraph* of G if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. H is *induced* if it has the edge set $\{X \in E(G) : X \subseteq V(H)\}$.

If G is an r-graph and $U \subseteq V(G)$ we let $G[U]$ be the induced subgraph on vertex set U. If no confusion can arise, we may identify $G[U]$ with U. We let $G \setminus U$ denote the induced subgraph on $V(G) \setminus U$. Furthermore if $x \in V(G)$ we let $d_U(x) = d_{G[U]}(x)$ denote the degree of x with respect to U, i.e. the number of edges of G containing x and intersecting $U \setminus \{x\}$ in $r-1$ vertices.

If G and H are r-graphs on disjoint vertex sets we let $G \dot\cup H$ denote the r-graph on vertex set $V(G) \dot\cup V(H)$ and edge set $E(G) \dot\cup E(H)$. We call $G \dot\cup H$ the *disjoint union* of G and H. As G can be replaced by an identical r-graph on vertex set $V(G) \times \{1\}$, and H by an identical r-graph on vertex set $V(H) \times \{2\}$, the definition of $G \dot\cup H$ extends naturally to pairs of r-graphs which are not necessarily disjoint.

If F and G are r-graphs, a map $f : V(F) \to V(G)$ is a *homomorphism* if it maps edges to edges. An *embedding* is an injective homomorphism. We shall frequently abuse the notion of subgraph and say F is a subgraph of G if there exists an embedding of F into G. We will denote this by $F \subseteq G$, and if no confusion can arise we may identify F with the image of its embedding in G.

If $\mathcal{F}$ is a family of r-graphs, the *closure of* $\mathcal{F}$ *under homomorphisms* is the family $\overline{\mathcal{F}}$ containing all r-graphs G for which there exists $F \in \mathcal{F}$ and a surjective homomorphism $f : V(F) \to V(G)$ (here f is surjective on $V(G)$, but G may contain edges not in the image of f). If $\mathcal{F} = \overline{\mathcal{F}}$ then $\mathcal{F}$ is *closed under homomorphisms.* If for any $G \in \overline{\mathcal{F}}$ there exists $F \subseteq G$ with $F \in \mathcal{F}$ we say $\mathcal{F}$ is *weakly closed under homomorphisms.*

If $\mathcal{F}$ is a family of r-graphs, an r-graph G is $\mathcal{F}$*-free* if no subgraph of G belongs to $\mathcal{F}$.

3.2 The Infinity Principle

If G is an r-multigraph on $[n]$, we define a polynomial $p_G(x)$ as follows:

$$p_G(x_1, \dots, x_n) := r! \sum_{D \in E(G)} \prod_{i=1}^{n} \frac{x_i^{D(i)}}{D(i)!}.$$

The *Lagrangian* of G is defined to be

$$\lambda(G) := \max \{p_G(\mathbf{x}) : \mathbf{x} \in \Delta_n\}. \tag{3.1}$$

The maximum is attained as it is taken over a compact set and p_G is continuous. An element $\mathbf{x} \in \Delta_n$ such that $p_G(\mathbf{x}) = \lambda(G)$ is an *optimal vector for* G. Note that $\lambda(G) = 0$ implies that G has no edges. For technical reasons we also define $\lambda(\emptyset) = 0$.

For $r \geq 1$, let $\Lambda^{(r)}$ be the set of values $\lambda(G)$, with G an r-graph. Note that we do not take into account non-simple r-multigraphs. Pikhurko proved the following.

Theorem 59 (Pikhurko, [96])**.** $\Lambda^{(r)} \subseteq \Pi^{(r)}_{\text{fin}}$.

The weaker statement $\Lambda^{(r)} \subseteq \Pi^{(r)}_{\infty}$ is much simpler to prove. In particular if e is an r-edge then $\lambda(e) = \frac{r!}{r^r} \in \Pi^{(r)}_{\infty}$. It was shown by Brown and Simonovits [17] that $\Lambda^{(r)}$ is dense in $\Pi^{(r)}_{\infty}$. As the latter is a closed set, this in fact proves the following.

Lemma 60. $\overline{\Lambda}^{(r)} = \Pi^{(r)}_{\infty}$.

We shall frequently rely on Lemma 60 to transfer statements about $\Lambda^{(r)}$ to the whole of $\Pi^{(r)}_{\infty}$ via continuity.

Pikhurko further proved that $\lambda(G) \in \Pi^{(r)}_{\text{fin}}$ for any r-multigraph G. We shall only need the following weaker statement.

Lemma 61. *For any r-multigraph G we have $\lambda(G) \in \Pi^{(r)}_{\infty}$.*

As Pikhurko's proof is long and difficult, we include here a short proof of Lemma 61.

First we need a definition introduced by Pikhurko in [96]. We reproduce it here in a simplified variant that better suits our needs.

Let $G = (S, E)$ be an r-multigraph. Identify S with $[m]$ and let $V_1, \dots, V_m$ be disjoint sets with $V := V_1 \cup \ldots \cup V_m$. The *profile* of an r-set $X \subseteq V$ (with respect to $V_1, \dots, V_m$) is the r-multiset on $[m]$ that contains $i \in [m]$ with multiplicity $|X \cap V_i|$. For an r-multiset $Y \subseteq [m]$ let $Y((V_1, \dots, V_m))$ consist of all r-subsets of V whose profile is Y. We call this r-graph the *blow-up* of Y and the r-graph

$$E((V_1, \dots, V_m)) := \bigcup_{Y \in E} Y((V_1, \dots, V_m))$$

is called the *blow-up of* G (with respect to $V_1, \dots, V_m$). If all sets V_i have the same size t, we denote $E((V_1, \dots, V_m))$ by $G(t)$.

A G-construction on a set V is any r-graph $E((V_1, \ldots, V_m))$ obtained by taking a partition $V = V_1 \cup \ldots \cup V_m$. Let p_n be the maximum number of edges of a G-construction on n vertices. Then Pikhurko defined

$$\Lambda_G := \lim_{n \to \infty} \frac{p_n}{\binom{n}{r}},$$

and proved that this limit always exists. It is easy to see that $\Lambda_G = \lambda(G)$. In fact Pikhurko defined a much larger class of G-constructions, where one is allowed to recursively apply the construction into some of the parts; this was a key step in his proof of Theorem 24.

The main observation is now the following, which is implicit in [46] and [17].

Lemma 62 (The Infinity Principle)**.** *Let $\{G_n\}_{n \geq 1}$ be a sequence of r-graphs with $v(G_n) = n$ and $d(G_n) \to \alpha$. Suppose that for any sequence of r-graphs H_n with $H_n \subseteq G_n$ and $v(H_n)$ tending to infinity, we have $\limsup_{n \to \infty} d(H_n) \leq \alpha$. Then $\alpha \in \Pi_\infty^{(r)}$.*

Proof. Define $\mathcal{F}_\infty := \{H : H \not\subseteq G_n \text{ for any } n \geq 1\}$. We claim $\pi(\mathcal{F}_\infty) = \alpha$.

Indeed, for any $m \geq 1$, let T_m be a maximum $\mathcal{F}_\infty$-free r-graph on m vertices. Then for each $m \geq 1$, $T_m \notin \mathcal{F}_\infty$, and hence $T_m \subseteq G_{n(m)}$, for some $n(m)$ depending on m. Define $\{H_n\}_{n \geq 1}$ in the following way. If there exists m with $n(m) = n$, let $H_n := T_m$ (if several choices exists, choose one with maximum density). Otherwise let $H_n := G_n$. Then $\lim_{m \to \infty} d(T_m) \leq \limsup_{n \to \infty} d(H_n) \leq \alpha$, by assumption. Hence $\pi(\mathcal{F}_\infty) \leq \alpha$. However, by construction $\pi(\mathcal{F}_\infty) \geq \alpha$, and so equality holds. $\square$

Proof of Lemma 61. Let G_n be a maximum[1] G-construction on n vertices. Let $H_n \subseteq G_n$ be any sequence of subgraphs with number of vertices tending to infinity. W.l.o.g. we may assume that H_n is an induced subgraph on $m(n)$ vertices. Then H_n is by definition also a G-construction, and hence has no more than $p_{m(n)}$ edges. Thus by definition of Λ_G, we must have $\limsup d(H_n) \leq \Lambda_G$. Consequently by the Infinity Principle, $\lambda(G) = \Lambda_G \in \Pi_\infty^{(r)}$. $\square$

If $G = (V, E)$ is an r-multigraph, we define $\overline{G} = (V, V^{(r)} \setminus E)$. One of the advantages of working with multigraphs is the following.

Lemma 63. *For any r-multigraph G on $[n]$ and any $\mathbf{x} \in \Delta_n$ we have*

$$p_G(\mathbf{x}) + p_{\overline{G}}(\mathbf{x}) = 1. \tag{3.2}$$

Proof. Note that $p_G(\mathbf{x}) + p_{\overline{G}}(\mathbf{x}) = (\sum_{i=1}^n x_i)^r = 1$. $\square$

3.3 The multiplicative structure

In this section we prove Theorem 29.

It is an easy exercise to check that $*$ is commutative, associative and cancellative. Furthermore the unit is the element $(1, 0)$, under the convention $0^0 = 1$. Thus we only need to show that Π_∞ is closed under $*$. To this end we make the following definition.

[1] i.e. with a maximum number of edges.

Definition 64. *Let $r, s \geq 0$, G be an r-graph and H an s-graph on disjoint vertex sets. We define $G * H$ as the $(r+s)$-graph on vertex set $V(G) \cup V(H)$ and edge set $\{e \cup f : e \in E(G), f \in E(H)\}$.*

This definition was introduced by Emtander [33] in connection with Betti numbers of hypergraphs. It was also considered by Bollobás, Leader and Malvenuto in the context of Turán densities [11]. The definition of $G * H$ extends naturally to any two (not necessarily disjoint) uniform hypergraphs G and H.

Proposition 65. *Let $r, s \geq 1$ and $\mathfrak{f} : [0,1] \to \mathbb{R}$ be given by $\mathfrak{f}(x) = x^r(1-x)^s$. Then $\mathfrak{f}$ has a unique maximum $x_0 := \frac{r}{r+s}$ and furthermore $\mathfrak{f}(x_0) = \frac{r^r s^s}{(r+s)^{r+s}}$.*

Proof. We see that $\mathfrak{f}'(x) = rx^{r-1}(1-x)^s - sx^r(1-x)^{s-1}$, and so $\mathfrak{f}'(x) = 0$ only happens for $x_0 := \frac{r}{r+s}$. As $\mathfrak{f}(0) = 0$, $\mathfrak{f}(x_0)$ must be a maximum point, and the claim follows. □

Lemma 66. *Let G be an r-graph and H be an s-graph. Then $\lambda(G*H) = \lambda(G)\lambda(H)\binom{r+s}{r}\frac{r^r s^s}{(r+s)^{r+s}}$.*

Proof. We assume w.l.o.g. that G has vertex set $\{1, \dots, n\}$ and H has vertex set $\{n+1, \dots, n+m\}$. Then $G * H$ has vertex set $[n+m]$.

By definition

$$p_{G*H}(x_1, \dots, x_n, y_1, \dots, y_m) = \binom{r+s}{r} p_G(x_1, \dots, x_n) p_H(y_1, \dots, y_m).$$

Let $\mathbf{a} \in \Delta_n$ be an optimal vector for G and $\mathbf{b} \in \Delta_m$ be an optimal vector for H. Let $\theta := \frac{r}{r+s}$. Then

$$\begin{aligned} \lambda(G * H) &\geq p_{G*H}(\theta a_1, \dots, \theta a_n, (1-\theta)b_1, \dots, (1-\theta)b_m) \\ &= \binom{r+s}{r} \theta^r \lambda(G)(1-\theta)^s \lambda(H) \\ &= \lambda(G)\lambda(H)\binom{r+s}{r}\frac{r^r s^s}{(r+s)^{r+s}}. \end{aligned} \tag{3.3}$$

On the other hand, let $\mathbf{z} \in \Delta_{n+m}$ be an optimal vector for $G * H$. Set $M := \sum_{i=1}^n z_i$. Then

$$\begin{aligned} \lambda(G * H) &= p_{G*H}(z_1, \dots, z_{n+m}) \\ &= \binom{r+s}{r} M^r p_G\left(\frac{z_1}{M}, \dots, \frac{z_n}{M}\right)(1-M)^s p_H\left(\frac{z_{n+1}}{1-M}, \dots, \frac{z_{n+m}}{1-M}\right) \\ &\leq \binom{r+s}{r} \lambda(G)\lambda(H)\mathfrak{f}(M) \\ &\leq \lambda(G)\lambda(H)\binom{r+s}{r}\frac{r^r s^s}{(r+s)^{r+s}}, \text{by Proposition 65.} \end{aligned}$$

Together with (3.3) this proves the claim. □

Proof of Theorem 29. Let $(\alpha, r), (\beta, s) \in \Pi_\infty$. We want to show that $(\alpha, r) * (\beta, s) \in \Pi_\infty$.

We may assume that $r, s \geq 1$.

By Lemma 60, there exists a sequence of r-graphs G_n with $\lambda(G_n) \to \alpha$. Similarly there exists a sequence of s-graphs H_n with $\lambda(H_n) \to \beta$. Then

$$\begin{aligned} \lim_{n\to\infty} \lambda(G_n * H_n) &= \lim_{n\to\infty} \left(\lambda(G_n)\lambda(H_n) \binom{r+s}{r} \frac{r^r s^s}{(r+s)^{r+s}} \right) \\ &= \alpha\beta \binom{r+s}{r} \frac{r^r s^s}{(r+s)^{r+s}}. \end{aligned}$$

Thus $\alpha\beta\binom{r+s}{r}\frac{r^r s^s}{(r+s)^{r+s}} \in \overline{\Lambda}^{(r+s)}$. Then Lemma 60 completes the proof. $\square$

3.4 Proof of the main result

In this section we prove Theorem 33. The proof is naturally divided into two parts. We first prove the semigroup structure of $\Pi_\infty^{(r)}$ and then the closure of $\Pi_{\text{fin}}^{(r)}$ under $\oplus_r$.

3.4.1 The semigroup structure

Let G and H be two r-graphs on disjoint vertex sets. We define $G \oplus_r H$ as the r-multigraph with vertex set $V(G)\dot\cup V(H)$ and edge set

$$E(G \oplus_r H) = E(G)\dot\cup E(H)\dot\cup\{e \in (V(G) \cup V(H))^{(r)} : e \text{ intersects both } V(G) \text{ and } V(H)\}.$$

We then extend this definition to pairs of r-graphs with intersecting vertex sets in the same manner as before.

Lemma 67. *Let $r \geq 2$ and $\alpha, \beta \in [0,1]$. Define $\mathfrak{g}_{\alpha,\beta} : [0,1] \to \mathbb{R}$ by*

$$\mathfrak{g}_{\alpha,\beta}(x) = \alpha x^r + \beta(1-x)^r + r!\sum_{i=1}^{r-1} \frac{x^i(1-x)^{r-i}}{i!(r-i)!}.$$

Then $\alpha \oplus_r \beta = \sup_{x\in[0,1]} \mathfrak{g}_{\alpha,\beta}(x)$, and moreover for $(\alpha, \beta) \neq (1,1)$, $\mathfrak{g}_{\alpha,\beta}$ is strictly concave and has a unique maximum at $x_{\alpha,\beta} := \frac{\sqrt[r-1]{1-\beta}}{\sqrt[r-1]{1-\alpha}+\sqrt[r-1]{1-\beta}}$.

Proof. Note that for $x \in [0,1]$ we have that

$$\begin{aligned}\alpha x^r + \beta(1-x)^r + r!\sum_{i=1}^{r-1} \frac{x^i(1-x)^{r-i}}{i!(r-i)!} &= \alpha x^r + \beta(1-x)^r + (1 - x^r - (1-x)^r) \\ &= 1 - (1-\alpha)x^r - (1-\beta)(1-x)^r.\end{aligned}$$

Hence $\mathfrak{g}_{\alpha,\beta}(x) = 1 - (1-\alpha)x^r - (1-\beta)(1-x)^r$. If $\alpha = 1$ and $\beta = 1$ then $\sup_{x\in[0,1]} g(x) = 1 = \alpha \oplus_r \beta$ by definition. Hence we may assume that $\alpha < 1$ or $\beta < 1$. Then

$$\mathfrak{g}'_{\alpha,\beta}(x) = -r(1-\alpha)x^{r-1} + r(1-\beta)(1-x)^{r-1}$$

and

$$\mathfrak{g}''_{\alpha,\beta}(x) = -r(r-1)(1-\alpha)x^{r-2} - r(r-1)(1-\beta)(1-x)^{r-2}.$$

Thus $\mathfrak{g}''_{\alpha,\beta} < 0$ on $(0,1)$, showing that $\mathfrak{g}_{\alpha,\beta}$ is strictly concave. Furthermore $\mathfrak{g}'_{\alpha,\beta}(x) = 0$ has a unique solution

$$x_{\alpha,\beta} = \frac{\sqrt[r-1]{1-\beta}}{\sqrt[r-1]{1-\alpha} + \sqrt[r-1]{1-\beta}}.$$

Hence $\mathfrak{g}_{\alpha,\beta}(x)$ has the global maximum

$$\mathfrak{g}_{\alpha,\beta}(x_{\alpha,\beta}) = 1 - \frac{(1-\alpha)(1-\beta)}{(\sqrt[r-1]{1-\alpha} + \sqrt[r-1]{1-\beta})^{r-1}},$$

which is the same as $\alpha \oplus_r \beta$. □

Lemma 68. *Let G and H be two r-graphs. Then*

$$\lambda(G \oplus_r H) = \lambda(G) \oplus_r \lambda(H).$$

Proof. We shall assume that G has vertex set $[n]$ and H has vertex set $\{n+1, \ldots, n+m\}$. Then $G \oplus_r H$ has vertex set $[n+m]$.

Let $\mathbf{x} \in \Delta_{n+m}$ arbitrary. Set $S_x := \sum_{i=1}^{n} x_i$ and note that

$$\begin{aligned} p_{G\oplus_r H}(x_1, \ldots, x_{n+m}) = p_G(\frac{x_1}{S_x}, \ldots, \frac{x_n}{S_x})S_x^r + p_H(\frac{x_{n+1}}{1-S_x}, \ldots, \frac{x_{n+m}}{1-S_x})(1-S_x)^r \\ + r! \sum_{i=1}^{r-1} \frac{S_x^i(1-S_x)^{r-i}}{i!(r-i)!}. \end{aligned}$$

Moreover $(\frac{a_1}{S_x}, \ldots, \frac{a_n}{S_x}) \in \Delta_n$ and $(\frac{a_{n+1}}{1-S_x}, \ldots, \frac{a_{n+m}}{1-S_x}) \in \Delta_m$. Hence

$$\begin{aligned} \lambda(G \oplus_r H) &= \sup_{x \in [0,1]} \left\{ \lambda(G)x^r + \lambda(H)(1-x)^r + r! \sum_{i=1}^{r-1} \frac{x^i(1-x)^{r-i}}{i!(r-i)!} \right\} \\ &= \lambda(G) \oplus_r \lambda(H), \end{aligned}$$

by Lemma 67. This proves the lemma. □

We can now prove the following.

Lemma 69. $(\Pi_\infty^{(r)}, \oplus_r)$ *is a commutative topological semigroup.*

Proof. Commutativity and associativity are simple exercises left to the reader. Continuity of $\oplus_r$ is clear everywhere except at $(1,1)$.

Let $\{x_n\}_{n\geq 1}$ and $\{y_n\}_{n\geq 1}$ be arbitrary sequences of real numbers from $[0,1)$ converging to 1. Then by definition $x_n \oplus_r y_n \leq 1$, and we want to show that equality holds in the limit. Let $\varepsilon > 0$ be arbitrary and define $\delta := 2^{r-1}\varepsilon$. Then there exists an $n_0 \geq 1$ such that $1 - x_n < \delta$ and $1 - y_n < \delta$. Let $n \geq n_0$ and assume w.l.o.g. that $x_n \geq y_n$. Then

$$1 - \frac{(1-x_n)(1-y_n)}{(\sqrt[r-1]{1-x_n} + \sqrt[r-1]{1-y_n})^{r-1}} \geq 1 - \frac{1-y_n}{2^{r-1}} > 1 - \frac{\delta}{2^{r-1}} = 1 - \varepsilon.$$

Hence $\oplus_r$ is continuous at $(1,1)$ as well.

Thus we only need to prove that $\Pi_\infty^{(r)}$ is closed with respect to $\oplus_r$.

Let $\alpha, \beta \in \Pi_\infty^{(r)}$. By Lemma 60, we can choose a sequence of r-graphs G_n with $\lambda(G_n) \to \alpha$, and a sequence of r-graphs H_n with $\lambda(H_n) \to \beta$.

Consider the sequence $G_n \oplus_r H_n$. By Lemma 68 and continuity of $\oplus_r$,

$$\begin{aligned} \lim_{n\to\infty} \lambda(G_n \oplus_r H_n) &= \lim_{n\to\infty} \lambda(G_n) \oplus_r \lambda(H_n) \\ &= \alpha \oplus_r \beta. \end{aligned}$$

By Lemma 61, $\lambda(G_n \oplus_r H_n) \in \Pi_\infty^{(r)}$. As this is a closed set, $\alpha \oplus_r \beta \in \Pi_\infty^{(r)}$, proving the lemma. □

We are now left to prove that $\Pi_{\text{fin}}^{(r)}$ is closed under $\oplus_r$. This task will be substantially more difficult.

3.4.2 The λ function

For the remaining part of the proof we shall need a number of additional statements, which we gather in the following 3 subsections.

We start with several observations on the Lagrangian function.

Lemma 70. *For any $r \geq 2$ the following holds.*

(i). If $H \subseteq G$ are r-graphs then $\lambda(H) \leq \lambda(G)$.

(ii). If $f : G \to H$ is a homomorphism of r-graphs then $\lambda(H) \geq \lambda(G)$.

(iii). If G and H are r-graphs then $\lambda(G \dot\cup H) = \max\{\lambda(G), \lambda(H)\}$.

Proof. Statement (i) is clear.

We prove (ii). Assume G has vertex set $[n]$ and H has vertex set $[m]$. Let $\mathbf{a} \in \Delta_n$ be an optimal vector for G. Define $\mathbf{b} \in \Delta_m$ by setting $b_i = \sum_{j\in f^{-1}(i)} a_j$. As f maps edges to edges, it follows that

$$\begin{aligned}\lambda(H) \geq p_H(\mathbf{b}) &= r! \sum_{e'\in E(H)} \prod_{i\in e'} (\sum_{j\in f^{-1}(i)} a_j) \\ &\geq r! \sum_{e'\in E(H)} \sum_{e\in f^{-1}(e')} \prod_{j\in e} a_j = r! \sum_{e\in E(G)} \prod_{j\in e} a_j = p_G(\mathbf{a}) = \lambda(G).\end{aligned}$$

This proves (ii).

We prove (iii). Assume G has vertex set $[n]$ and H has vertex set $\{n+1, \dots, n+m\}$. For any $\mathbf{x} \in \Delta_{n+m}$ let $S_x := \sum_{i=1}^{n} x_i$. Then

$$\begin{aligned}p_{G\dot\cup H}(\mathbf{x}) &= S_x^r p_G(\frac{x_1}{S_x}, \dots, \frac{x_n}{S_x}) + (1 - S_x)^r p_H(\frac{x_{n+1}}{1 - S_x}, \dots, \frac{x_{n+m}}{1 - S_x}) \\ &\leq S_x^r \lambda(G) + (1 - S_x)^r \lambda(H).\end{aligned}$$

Consider the function $f : [0, 1] \to [0, 1]$ given by $f(x) = x^r \lambda(G) + (1 - x)^r \lambda(H)$. Then the second derivative $f''(x) \geq 0$, hence f is convex. So the maximum is achieved at one of the endpoints of the interval. This implies that $\lambda(G\dot\cup H) = \max\{\lambda(G), \lambda(H)\}$, proving (iii). $\square$

3.4.3 The π function

We gather in this subsection several results about the π function.

Theorem 71 (Theorem 2, [17])**.** *For any $\varepsilon > 0$ and any family $\mathcal{F}$ of r-graphs, there are $\delta > 0$ and n_S such that the following holds. Any r-graph G on $n \geq n_S$ vertices and with more than $(\pi(\mathcal{F}) + \varepsilon)\binom{n}{r}$ edges contains at least $\delta n^{v(F)}$ copies of some $F \in \mathcal{F}$.*

Theorem 71 is a generalization of the supersaturation theorem of Erdős and Simonovits. It has the following consequence (see the proof of Theorem 2.2 in [73]).

Lemma 72. *For any $t \geq 1$, $\eta > 0$ and any r-graph F there exist $\rho > 0$ and n_L such that the following holds. If G is an r-graph on $n \geq n_L$ vertices containing at least $\eta n^{v(F)}$ copies of F, then G contains at least $\rho n^{v(F)t}$ copies of $F(t)$.*

Lemma 73. *For any $\eta > 0$, $a \geq 1$ and $k \geq 1$ there is an n_C such that the following holds. If G is an r-graph on $n \geq n_C$ vertices containing at least $\eta n^{v(F)}$ copies of some r-graph F on at most k vertices and $A \subset V(G)$ is a set of size a, then $G \setminus A$ contains at least $\frac{\eta}{2} n^{v(F)}$ copies of F.*

Proof. Set $n_C := \frac{2ak}{\eta}$. We show that the claim holds for n_C.

The number of copies of F intersecting A is at most $v(F)an^{v(F)-1} \leq kan^{v(F)-1}$. Hence the number of copies of F disjoint from A is at least

$$\eta n^{v(F)} - kan^{v(F)-1} \geq \frac{\eta n^{v(F)}}{2},$$

which holds by our choice of n_C. $\square$

The π function has several properties, which we list below.

Lemma 74. *For any $r \geq 2$ the following holds.*

(i). If $\mathcal{F}$ is a family of r-graphs and $\mathcal{F} \subseteq \mathcal{F}'$ then $\pi(\mathcal{F}') \leq \pi(\mathcal{F})$.

(ii). If $H \subseteq G$ are two r-graphs and $\mathcal{F}$ is a family of r-graphs then $\pi(\mathcal{F} \cup \{H\}) \leq \pi(\mathcal{F} \cup \{G\})$.

(iii). If G and H are r-graphs and $\mathcal{F}$ is a family of r-graphs then

$$\pi(\mathcal{F} \cup \{G \dot\cup H\}) = \max\{\pi(\mathcal{F} \cup \{H\}), \pi(\mathcal{F} \cup \{G\})\}. \tag{3.4}$$

(iv). If F is any r-graph, $t \geq 1$ and $\mathcal{F}$ is a family of r-graphs then $\pi(\mathcal{F} \cup \{F\}) = \pi(\mathcal{F} \cup \{F(t)\})$.

(v). If $\mathcal{F}$ is a family of r-graphs then $\pi(\mathcal{F}) = 1$ if and only if $\mathcal{F}$ is empty.

Proof. The statements (i) and (ii) are clear.

We prove (iii). Assume w.l.o.g. that $\pi(\mathcal{F} \cup \{G\}) \geq \pi(\mathcal{F} \cup \{H\})$. By (ii) we have that

$$\pi(\mathcal{F} \cup \{G \dot\cup H\}) \geq \pi(\mathcal{F} \cup \{G\}).$$

Assume for a contradiction that this inequality is strict. Then there exists $\varepsilon > 0$ such that $\pi(\mathcal{F} \cup \{G \dot\cup H\}) > \pi(\mathcal{F} \cup \{G\}) + \varepsilon$. Let δ and n_S be given by Theorem 71 on input ε and $\mathcal{F} \cup \{G\}$. Let n_C be given by Lemma 73 on input δ, $a := v(H)$ and $k := v(G)$. Furthermore, let n_1 be large enough so that for any $n \geq n_1$, $\mathrm{ex}(n, \mathcal{F} \cup \{G \dot\cup H\}) > (\pi(\mathcal{F} \cup \{G\}) + \varepsilon)\binom{n}{r} > \mathrm{ex}(n, \mathcal{F} \cup \{H\})$.

Now let $n \geq \max\{n_S, n_C, n_1, \frac{2}{\delta}\}$. By assumption there must exist an r-graph G' on n vertices with $e(G') > (\pi(\mathcal{F} \cup \{G\}) + \varepsilon)\binom{n}{r}$, which is $(\mathcal{F} \cup \{G \dot\cup H\})$-free. As $n \geq n_1$, there must exist a copy of H in G', which we now fix, and also denote by H. However, by our choice of n and Theorem 71, there are at least $\delta n^{v(G)}$ copies of G in G'. Consequently by Lemma 73, at least $\frac{\delta}{2}n^{v(G)}$ of them are disjoint from H, in particular we can find a copy of $G\dot{\cup}H$ in G', a contradiction. This proves (iii).

We prove (iv). By (ii) we have that $\pi(\mathcal{F}\cup\{F\}) \leq \pi(\mathcal{F}\cup\{F(t)\})$. Assume for a contradiction that this inequality is strict. Then there exists $\varepsilon > 0$ such that $\pi(\mathcal{F}\cup\{F(t)\}) > \pi(\mathcal{F}\cup\{F\}) + \varepsilon$. Let δ and n_S be given by Theorem 71 on input ε and $\mathcal{F}\cup\{F\}$. Let ρ and n_L be given by Lemma 72 on input t, δ and F. Furthermore, let n_1 be large enough so that for any $n \geq n_1, \mathrm{ex}(n, \mathcal{F}\cup\{F(t)\}) > (\pi(\mathcal{F}\cup\{F\}) + \varepsilon)\binom{n}{r}$.

Now let $n \geq \max\{n_S, n_L, n_1, \frac{1}{\rho}\}$. By assumption there must exist an r-graph G on n vertices with $e(G) > (\pi(\mathcal{F}\cup\{F\}) + \varepsilon)\binom{n}{r}$, which is $(\mathcal{F}\cup\{F(t)\})$-free. By Theorem 71, there are at least $\delta n^{v(F)}$ copies of F in G. Consequently by Lemma 72, there is at least one copy of $F(t)$ in G, a contradiction. This proves (iv).

We prove (v). Clearly $\pi(\emptyset) = 1$, so assume $\mathcal{F}$ is a non-empty family of r-graphs. Let $F \in \mathcal{F}$. Then $\pi(\mathcal{F}) \leq \pi(F)$ by (i), and we claim that $\pi(F) < 1$. This intuitive claim can be proved in several ways, for example by using the following inequality of Sidorenko [111]: if F is an r-graph with $f \geq 2$ edges then $\pi(F) \leq \frac{f-2}{f-1}$. This proves (v). $\square$

Finally, the following two lemmas will provide a better understanding of the structure of extremal r-graphs.

Lemma 75. *Let $r \geq 2$ and $\mathcal{F}$ be a family of r-graphs weakly closed under homomorphisms. Set $\alpha := \pi(\mathcal{F})$. For any $\delta > 0$ there exists an $n_D \geq 1$ such that any maximum $\mathcal{F}$-free r-graph on $n \geq n_D$ vertices has minimum degree at least $(\alpha - \delta)\binom{n-1}{r-1}$.*

Proof. Choose $n_D \geq 1$ so that for any $n \geq n_D$ we have $\mathrm{ex}(n, \mathcal{F}) \geq (\alpha - \frac{\delta}{2})\binom{n}{r}$, and furthermore $n_D > 1 + \frac{2(r-1)}{\delta}$.

Let G be any maximum $\mathcal{F}$-free r-graph on $n \geq n_D$ vertices, and assume for a contradiction that G contains a vertex x of degree $d(x) < (\alpha - \delta)\binom{n-1}{r-1}$.

As $e(G) = \sum_{x\in V(G)} \frac{d(x)}{r}$, there must exist a vertex $v \in V(G)$ of degree $d(v) \geq (\alpha - \frac{\delta}{2})\binom{n-1}{r-1}$. Then replace x by a new vertex v' and add all edges $\{v'\}\cup e$, where $x \notin e$ and $\{v\}\cup e \in E(G)$. In other words, we replace x by a copy of v, but we duplicate only the edges incident with v and not with x. Let G' be the resulting r-graph. As $\mathcal{F}$ is weakly closed under homomorphisms, G' is still $\mathcal{F}$-free.

However,

$$\begin{aligned}
e(G') - e(G) &\geq (\alpha - \frac{\delta}{2})\binom{n-1}{r-1} - (\alpha - \delta)\binom{n-1}{r-1} - \binom{n-2}{r-2} \\
&= \left(\frac{\delta(n-1)}{2(r-1)} - 1\right)\binom{n-2}{r-2} \\
&> 0,
\end{aligned}$$

as $n \geq n_D$. This contradicts the maximality of G. $\square$

One can strengthen the proof of Lemma 75 to show that in a maximum $\mathcal{F}$-free r-graph all degrees are roughly the same $(\alpha + o(1))\binom{n-1}{r-1}$. We shall not need this stronger statement, but rather a variation of it.

Lemma 76. *Let $r \geq 2$ and $\mathcal{F}$ be a family of r-graphs weakly closed under homomorphisms. Set $\alpha := \pi(\mathcal{F})$. For any $\varepsilon > 0$, there exist a $\tau > 0$ and an $n_V \geq 1$ such that the following holds. If G is any r-graph on $n \geq n_V$ vertices, density at least $\alpha - \tau$ and having a vertex v of degree at least $(\alpha + \varepsilon)\binom{n-1}{r-1}$, then G is not $\mathcal{F}$-free.*

Proof. We assume w.l.o.g. that $\varepsilon < 1$.

Set $\tau := \frac{\varepsilon^2 r}{72(r-1)}$ and choose $n_V \geq 2$ so that for any $n \geq n_V$, $\mathrm{ex}(n, \mathcal{F}) < (\alpha + \tau)\binom{n}{r}$.

We shall assume for a contradiction that G is $\mathcal{F}$-free. Set $n := v(G)$. Then

$$(\alpha + \frac{\varepsilon}{4})\binom{n}{r} > e(G) = \sum_{x \in V(G)} \frac{d(x)}{r}.$$

Let $S := \{x \in V(G) : d(x) \leq (\alpha + \frac{\varepsilon}{2})\binom{n-1}{r-1}\}$. Then $e(G) \geq \frac{n-|S|}{r}(\alpha + \frac{\varepsilon}{2})\binom{n-1}{r-1}$, so

$$|S| \geq (1 - \frac{\alpha + \varepsilon/4}{\alpha + \varepsilon/2})n = \frac{\varepsilon}{4\alpha + 2\varepsilon}n.$$

Note that $\frac{\varepsilon}{4\alpha+2\varepsilon} > \frac{\varepsilon}{6(r-1)}$, hence we can fix $S' \subseteq S$ of size $\frac{\varepsilon n}{6(r-1)}$ (here and in what follows we ignore upper and lower integer parts; this does not affect our arguments).

We construct a new r-graph G' from G by deleting all edges incident to S' and adding all edges $\{\{x\} \cup e : \{v\} \cup e \in E(G \setminus S'), x \in S'\}$. Then $S' \cup \{v\}$ is an independent set in G'.

We claim G' is $\mathcal{F}$-free. Indeed, if $f : V(F) \to V(G')$ is any embedding of a graph $F \in \mathcal{F}$ into G', then composing f with the map $g : V(G') \to V(G)$ that sends S' to v and is the identity otherwise, gives a homomorphism of F into G. Thus there exists a surjective homomorphism $f' : F \to F'$ with $F' \subseteq G$. As $\mathcal{F}$ is weakly closed under homomorphisms, F' and hence G contains an element of $\mathcal{F}$ as a subgraph, a contradiction.

Consequently $e(G') < (\alpha + \tau)\binom{n}{r}$. But

$$\begin{aligned}
e(G') - e(G) &\geq |S'|\left(d_G(v) - |S'|\binom{n-2}{r-2}\right) - |S'|(\alpha + \frac{\varepsilon}{2})\binom{n-1}{r-1} \\
&\geq |S'|\left((\alpha+\varepsilon)\binom{n-1}{r-1} - |S'|\binom{n-2}{r-2}\right) - |S'|(\alpha + \frac{\varepsilon}{2})\binom{n-1}{r-1} \\
&\geq |S'|\binom{n-2}{r-2}((\alpha+\varepsilon)\frac{n-1}{r-1} - |S'|) - |S'|(\alpha + \frac{\varepsilon}{2})\binom{n-1}{r-1} \\
&\geq |S'|(\alpha + \frac{2\varepsilon}{3})\binom{n-1}{r-1} - |S'|(\alpha + \frac{\varepsilon}{2})\binom{n-1}{r-1}, \text{ as } |S'| \leq \frac{\varepsilon n}{6(r-1)} \text{ and } n_V \geq 2, \\
&= \frac{\varepsilon}{6}|S'|\binom{n-1}{r-1} \\
&= \frac{\varepsilon^2 r}{36(r-1)}\binom{n}{r},
\end{aligned}$$

which is at least $2\tau\binom{n}{r}$. Hence $e(G') \geq e(G) + 2\tau\binom{n}{r} \geq (\alpha + \tau)\binom{n}{r}$, a contradiction. □

3.4.4 The $\oplus_r$ function

We now study the map $\oplus_r : [0,1] \times [0,1] \to [0,1]$.

Lemma 77. *The $\oplus_r$ function is nondecreasing in each of its arguments on $[0,1] \times [0,1]$. In fact, for any $\alpha, \beta \in [0,1)$ and $0 \le \varepsilon \le \alpha$ we have*

$$\alpha \oplus_r \beta \ge (\alpha - \varepsilon) \oplus_r \beta + \varepsilon \left(\frac{\sqrt[r-1]{1-\beta}}{1 + \sqrt[r-1]{1-\beta}} \right)^r \tag{3.5}$$

Proof. The first statement follows immediately from the definition of $\oplus_r$.

To prove the second, define $h_\beta : [0,1] \to \mathbb{R}$ by

$$h_\beta(x) = 1 - \frac{(1-x)(1-\beta)}{(\sqrt[r-1]{1-x} + \sqrt[r-1]{1-\beta})^{r-1}}.$$

Then the first order derivative of $h_\beta(x)$ exists and it is equal to

$$h'_\beta(x) = \left(\frac{\sqrt[r-1]{1-\beta}}{\sqrt[r-1]{1-x} + \sqrt[r-1]{1-\beta}} \right)^r.$$

Thus for any $x \in [0,1)$ we have

$$h'_\beta(x) \ge h'_\beta(0) = \left(\frac{\sqrt[r-1]{1-\beta}}{1 + \sqrt[r-1]{1-\beta}} \right)^r.$$

Hence

$$\begin{aligned}
\alpha \oplus_r \beta &= (\alpha - \varepsilon) \oplus_r \beta + \int_{\alpha-\varepsilon}^{\alpha} h'_\beta(x)\, dx \\
&\ge (\alpha - \varepsilon) \oplus_r \beta + \varepsilon h'_\beta(0) \\
&= (\alpha - \varepsilon) \oplus_r \beta + \varepsilon \left(\frac{\sqrt[r-1]{1-\beta}}{1 + \sqrt[r-1]{1-\beta}} \right)^r,
\end{aligned}$$

proving the lemma. □

3.4.5 The Rigidity Lemma

From this point on we adopt the strategy developed by Pikhurko in [96] (which in turn follows the Stability Method pioneered by Simonovits). The first step is to prove a *rigidity lemma*: we construct some graphs which can embed only in a prescribed way in a graph of the form $G \times H$, where $\times$ is a special type of product of hypergraphs which we now define.

Let G and H be two r-graphs on disjoint vertex sets. We define $G \times H$ as the r-graph with vertex set $V(G) \dot\cup V(H)$ and edge set $E(G) \dot\cup E(H) \dot\cup \{e \in \binom{V(G) \dot\cup V(H)}{r} : e \text{ intersects both } V(G) \text{ and } V(H)\}$. We then extend this definition to r-graphs with intersecting vertex sets in the same manner as before.

If $\mathcal{F}$ is family of r-graphs and $M \ge 1$ an integer, we let $\mathcal{F}(M)$ be a maximal family of r-graphs with the following properties:

(i). $\mathcal{F} \subseteq \mathcal{F}(M)$.

(ii). If $F \in \mathcal{F}(M) \setminus \mathcal{F}$ then F has at least one edge and $v(F) \leq M$.

(iii). $\pi(\mathcal{F}(M)) = \pi(\mathcal{F})$.

We call $\mathcal{F}(M)$ *an M-closure of $\mathcal{F}$*. Clearly there could be several distinct M-closures for a fixed $\mathcal{F}$, and moreover an M-closure always exists for any $M \geq 1$.

Suppose now that $\pi(\mathcal{F}) > 0$. By maximality, for any $F \notin \mathcal{F}(M)$ on at most M vertices, we have $\pi(\mathcal{F}(M) \cup \{F\}) < \pi(\mathcal{F})$. Thus we can define the *threshold of* $\mathcal{F}(M)$ as

$$\theta(\mathcal{F}(M)) := \max\{\pi(\mathcal{F}(M) \cup \{F\}) : F \notin \mathcal{F}(M) \text{ and } v(F) \leq M\}, \tag{3.6}$$

and this number is well-defined and strictly less than $\pi(\mathcal{F})$. For any $0 < \varepsilon < \pi(\mathcal{F}) - \theta(\mathcal{F}(M))$ and any $F \notin \mathcal{F}(M)$ on at most M vertices, Theorem 71 applied to ε and $\mathcal{F}(M) \cup \{F\}$ gives us a $\delta(\varepsilon, F) > 0$ and an $n_S(\varepsilon, F) \geq 1$. We set

$$\delta(\varepsilon, \mathcal{F}(M)) := \min\{\delta(\varepsilon, F) : F \notin \mathcal{F}(M) \text{ and } v(F) \leq M\},$$
$$n^*(\varepsilon, \mathcal{F}(M)) := \max\{n_S(\varepsilon, F) : F \notin \mathcal{F}(M) \text{ and } v(F) \leq M\}.$$

If $\pi(\mathcal{F}) = 0$ then $\mathcal{F}(M)$ contains all r-graphs on at most M vertices and with at least one edge. For technical reasons we define $\theta(\mathcal{F}(M)) = -1$, and for any $0 < \varepsilon < 1$, we set $\delta(\varepsilon, \mathcal{F}(M)) = 1$ and $n^*(\varepsilon, \mathcal{F}(M)) = 1$.

If $\mathcal{F}$ is family of r-graphs and F is any r-graph, we say F is *valid with respect to (w.r.t)* $\mathcal{F}$ if $\pi(\mathcal{F} \cup \{F\}) < \pi(\mathcal{F})$ or $\lambda(F) = \pi(\mathcal{F}) = 0$ (the point of this last condition is that we want the 1-vertex graph to be valid w.r.t. any family of r-graphs). Otherwise we call F *invalid w.r.t.* $\mathcal{F}$. We say F is *minimal invalid w.r.t.* $\mathcal{F}$ if F is not valid, but for any $x \in V(F)$, the r-graph $F \setminus x$ is valid. Note that any invalid r-graph contains at least one edge. Moreover the empty graph is valid w.r.t. $\mathcal{F}$ for any family of r-graphs $\mathcal{F}$.

One particular example the reader should keep in mind is the family $\mathcal{F} = \{I_{r-1}\}$, where I_{r-1} is the r-graph consisting of $r-1$ isolated vertices. By our definition I_{r-1} is valid w.r.t. $\mathcal{F}$. To complicate matters further, any I_{r-1}-free r-graph has a bounded number of vertices and hence there is no sequence of extremal graphs with size tending to infinity. Later on we will show that we can avoid working with such families, but we will allow them for now.

Now suppose two families of r-graphs $\mathcal{F}_\alpha$ and $\mathcal{F}_\beta$ are given. Let F be any r-graph. A partition of $V(F)$ into C_1 and C_2 is denoted by (C_1, C_2), and we identify C_1 with the r-graph $F[C_1]$, and similarly C_2 with the r-graph $F[C_2]$. We allow C_1 or C_2 to be empty. A partition (C_1, C_2) is called *valid w.r.t $\mathcal{F}_\alpha$ and $\mathcal{F}_\beta$* if C_1 is valid w.r.t. $\mathcal{F}_\alpha$ and C_2 is valid w.r.t. $\mathcal{F}_\beta$.

Let us now record several simple observations concerning valid graphs.

Lemma 78. *Let $\mathcal{F}$ be any family of r-graphs, $M \geq 1$ arbitrary and $\mathcal{F}(M)$ an arbitrary M-closure of $\mathcal{F}$. Then the following holds.*

(i). If $H \subseteq G$ and G is valid w.r.t. $\mathcal{F}$ then so is H.

(ii). If $H \subseteq G$ and H is invalid w.r.t. $\mathcal{F}$ then so is G.

(iii). If G and H are both valid w.r.t. $\mathcal{F}$ then so is $G\dot{\cup}H$.

(iv). If G is valid w.r.t. $\mathcal{F}$ then G is also valid w.r.t. $\mathcal{F}(M)$.

(v). If $G \in \mathcal{F}$ is invalid w.r.t. $\mathcal{F}$ then G is also invalid w.r.t. any family of r-graphs $\mathcal{G}$ containing $\mathcal{F}$, in particular G is invalid w.r.t. $\mathcal{F}(M)$.

Proof. Set $\alpha := \pi(\mathcal{F})$. We first prove (i).

If $\alpha > 0$ then by Lemma 74, (ii), we have that $\pi(\mathcal{F} \cup \{H\}) \leq \pi(\mathcal{F} \cup \{G\}) < \alpha$, and hence H is valid. If $\alpha = 0$, by Lemma 70, (i), we have that $\lambda(H) \leq \lambda(G) = 0$, and hence H is again valid. Thus (i) holds.

We prove (ii).

By Lemma 74, (i) and (ii), $\alpha = \pi(\mathcal{F} \cup \{H\}) \leq \pi(\mathcal{F} \cup \{G\}) \leq \alpha$ and so $\pi(\mathcal{F} \cup \{G\}) = \alpha$. Furthermore as H is invalid we have $\lambda(G) \geq \lambda(H) > 0$. Hence (ii) holds as well.

We now prove (iii).

If $\alpha = 0$ then by Lemma 70, (iii),

$$\lambda(G\dot{\cup}H) = \max\{\lambda(G), \lambda(H)\} = 0,$$

as G and H are both valid w.r.t. $\mathcal{F}$. Hence $G\dot{\cup}H$ is also valid.

If $\alpha > 0$, then by Lemma 74, (iii),

$$\pi(\mathcal{F} \cup \{G\dot{\cup}H\}) = \max\{\pi(\mathcal{F} \cup \{G\}), \pi(\mathcal{F} \cup \{H\})\} < \alpha,$$

again as G and H are both valid w.r.t. $\mathcal{F}$. Thus $G\dot{\cup}H$ is also valid, showing (iii).

We prove (iv).

Let G be any r-graph valid w.r.t. $\mathcal{F}$. If $\alpha > 0$, then by Lemma 74, (i),

$$\pi(\mathcal{F}(M) \cup \{G\}) \leq \pi(\mathcal{F} \cup \{G\}) < \pi(\mathcal{F}) = \pi(\mathcal{F}(M))$$

and hence G is valid w.r.t. $\mathcal{F}(M)$.

If $\alpha = 0$ then $\pi(\mathcal{F}(M)) = 0$ as well, and again G is valid w.r.t. $\mathcal{F}(M)$. This shows (iv).

We prove (v).

Assume $G \in \mathcal{F}$ and G is invalid w.r.t. $\mathcal{F}$. Then $\lambda(G) > 0$. If $\mathcal{G} \supseteq \mathcal{F}$ is any family of r-graphs containing $\mathcal{F}$, then $G \in \mathcal{G}$. Hence $\pi(\mathcal{G} \cup \{G\}) = \pi(\mathcal{G})$. Thus G is invalid w.r.t. $\mathcal{G}$, showing (v). □

Note that in Lemma 78, (v), the assumption $G \in \mathcal{F}$ played a crucial role.

Lemma 79 (The Rigidity Lemma)**.** *Let $\mathcal{F}_\alpha$ and $\mathcal{F}_\beta$ be two non-empty families of r-graphs with $\pi(\mathcal{F}_\alpha) = \alpha$ and $\pi(\mathcal{F}_\beta) = \beta$.*

Let P be any r-graph valid or minimal invalid w.r.t. $\mathcal{F}_\alpha$ such that if P is minimal invalid, then $P \in \mathcal{F}_\alpha$. Let Q be any r-graph minimal invalid w.r.t. $\mathcal{F}_\beta$ such that $Q \in \mathcal{F}_\beta$.

For any choice of $v \in V(P)$ and $w \in V(Q)$, there exists an $M_{P,Q,v,w} > 0$ such that for any $M \geq M_{P,Q,v,w}$ and any M-closures $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$ the following holds.

There exists an r-graph $C(P, Q, v, w)$ with the following properties:

(A). $v(C(P, Q, v, w)) \leq M_{P,Q,v,w}$.

(B). *Let K be the graph obtained from P and Q by identifying v with w. Then $C(P, Q, v, w)$ contains an induced copy K' of K such that $C(P, Q, v, w) \setminus K'$ has a valid partition (C_1, C_2) w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$, and any edge intersecting K' is either contained in K', or intersects $P \setminus v$ in one vertex and C_2 in $r - 1$ vertices, or intersects $Q \setminus w$ in one vertex and C_1 in $r - 1$ vertices.*

(C). *If P is valid w.r.t. $\mathcal{F}_\alpha$ then for any valid partition (C_1, C_2) of $C(P, Q, v, w)$ w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$ we have $P \subseteq C_1$.*

If P is minimal invalid w.r.t. $\mathcal{F}_\alpha$ then $C(P, Q, v, w)$ has no valid partition w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$.

Before we proceed to the proof of the Rigidity Lemma we note the following consequence, which we will also use in the proof.

Lemma 80 (Addendum to the Rigidity Lemma)**.** *Under the hypotheses of Lemma 79, the following holds.*

(a). *If P is valid w.r.t. $\mathcal{F}_\alpha$ then $C(P, Q, v, w)$ has a valid partition (C_1, C_2) w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$, with C_1 containing an induced copy P' of P. Furthermore any edge intersecting P' is either contained in P' or intersects P' in one vertex and C_2 in $r - 1$ vertices.*

(b). *If P is minimal invalid w.r.t. $\mathcal{F}_\alpha$ then $C(P, Q, v, w)$ has a partition (C_1, C_2) with C_1 containing an induced copy P' of P, such that $(C_1 \setminus P', C_2)$ is a valid partition of $C(P, Q, v, w) \setminus P'$ w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$. Furthermore any edge intersecting P' is either contained in P' or intersects P' in one vertex and C_2 in $r - 1$ vertices.*

Proof. By (B) of the Rigidity Lemma, $C(P, Q, v, w)$ contains an induced copy K' of K such that $C(P, Q, v, w) \setminus K'$ has a valid partition (C'_1, C'_2) w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$. Let P_1 and Q_1 be the copies of $P \setminus v$, respectively $Q \setminus w$, in K', and z the unique vertex in $K' \setminus (P_1 \cup Q_1)$.

Define $C_1 := C'_1 \dot\cup (P_1 \cup \{z\})$ and $C_2 := C'_2 \dot\cup Q_1$. Note that $C_1 \simeq C'_1 \dot\cup P$, as P_1 and z form an induced copy P' of P in C_1.

By (B), any edge intersecting P' is either contained in P' or intersects P' in one vertex and C_2 in $r - 1$ vertices.

As Q_1 is valid w.r.t. $\mathcal{F}_\beta$, C_2 is always valid w.r.t. $\mathcal{F}_\beta(M)$ by Lemma 78, (iii) and (iv).

Moreover if P is valid w.r.t. $\mathcal{F}_\alpha$ then C_1 is also valid w.r.t. $\mathcal{F}_\alpha(M)$ by Lemma 78, (iii) and (iv). This proves (a).

Finally, if P is minimal invalid w.r.t. $\mathcal{F}_\alpha$ then $C_1 \setminus P' = C'_1$ is valid w.r.t. $\mathcal{F}_\alpha(M)$. This proves (b). $\square$

Proof of the Rigidity Lemma. Define $M_1 := v(Q)$ and by induction on $k \geq 2$ define the positive integer

$$M_k := M_{k-1} + 2^{M_{k-1}}(2v(Q) - 2 + k). \tag{3.7}$$

We shall show that the lemma holds for $M_{P,Q,v,w} := M_{v(P)}$. Let $M \geq M_{P,Q,v,w}$ be arbitrary and consider any M-closures $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$.

We prove by induction on $v(P) \geq 1$ that an r-graph $C(P, Q, v, w)$ with the desired properties exists.

First assume that $v(P) = 1$. Then P is just a vertex v and necessarily P is valid w.r.t. $\mathcal{F}_\alpha$.

Define $C(P, Q, v, w) := Q$. Then $v(C(P, Q, v, w)) = M_1 = M_{P,Q,v,w}$, proving (A).

Clearly Q is isomorphic to the r-graph K' prescribed by (B), and the empty graph has always a valid partition $(\emptyset, \emptyset)$ w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$. This proves (B).

Finally, let (C_1, C_2) be any valid partition of $C(P, Q, v, w)$ w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$. If $C_1 = \emptyset$ then $C_2 = Q$, contradicting the validity of C_2 w.r.t. $\mathcal{F}_\beta(M)$ by Lemma 78, (v). Therefore $C_1 \neq \emptyset$, and hence $P \subseteq C_1$, proving (C).

Now assume that $v(P) > 1$ and the induction hypothesis holds for all r-graphs P' on fewer vertices, such that P' is either valid w.r.t. $\mathcal{F}_\alpha$, or belongs to $\mathcal{F}_\alpha$ and is minimal invalid with respect to it.

If P is valid, by Lemma 78, (i), so is $P \setminus v$. If P is minimal invalid, then $P \setminus v$ is valid by definition. Thus in any case $P \setminus v$ is valid. Fix $v' \in V(P \setminus v)$ arbitrary. As $M \geq M_{P\setminus v, Q, v', w}$, the induction hypothesis gives us an r-graph $C(P \setminus v) := C(P \setminus v, Q, v', w)$ satisfying (A)-(C).

We define a sequence of r-graphs $F_0, F_1, \ldots$ as follows.

Let (C_1^0, C_2^0) be the valid partition of $C(P \setminus v)$ w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$, guaranteed by the Addendum to the Rigidity Lemma, (a). Let Q_1^0 and Q_2^0 be vertex disjoint copies of $Q \setminus w$, and P_1^0 a copy of $P \setminus v$. To the r-graph $C_1^0 \times (C_2^0 \dot\cup Q_1^0)$ add P_1^0 and all edges intersecting P_1^0 in one vertex and Q_1^0 in $r-1$ vertices. Then add Q_2^0 and all edges intersecting Q_2^0 in one vertex and C_1^0 in $r-1$ vertices. Finally, add a vertex z_0 and edges in such a way that $P_1^0 \cup \{z_0\}$ induces a copy of P, and $Q_2^0 \cup \{z_0\}$ induces a copy of Q. No other edges incident with z_0 are added. This defines F_0.

Now suppose $i \geq 0$ and we have constructed F_i.

First assume there exists a partition (D_1^{i+1}, D_2^{i+1}) of F_i valid w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$ such that $P \not\subseteq D_1^{i+1}$. We construct F_{i+1} in a similar manner as above. Let Q_1^{i+1} and Q_2^{i+1} be vertex disjoint copies of $Q \setminus w$, and let P_1^{i+1} be a copy of $P \setminus v$. To the r-graph $D_1^{i+1} \times (D_2^{i+1} \dot\cup Q_1^{i+1})$ add P_1^{i+1} and all edges intersecting P_1^{i+1} in one vertex and Q_1^{i+1} in $r-1$ vertices. Then add Q_2^{i+1} and all edges intersecting Q_2^{i+1} in one vertex and D_1^{i+1} in $r-1$ vertices. Finally, add a vertex z_{i+1} and edges in such a way that $P_1^{i+1} \cup \{z_{i+1}\}$ induces a copy of P, and $Q_2^{i+1} \cup \{z_{i+1}\}$ induces a copy of Q. No other edges incident with z_{i+1} are added. This defines F_{i+1}.

If no partition (D_1^{i+1}, D_2^{i+1}) with the desired properties exists, we set $C(P, Q, v, w) := F_i$ and we stop.

Claim 1. *The sequence has at most $2^{v(C(P\setminus v))}$ terms.*

Proof. Note that

$$C(P \setminus v) \subseteq F_0 \subseteq F_1 \subseteq F_2 \ldots.$$

We identify the vertices of these graphs in such a way that all the above inclusion maps are given by the identity. Therefore we can speak about "the" subgraph $C(P \setminus v)$ of F_i, although F_i may possibly contain other copies of $C(P \setminus v)$.

For each $i \geq 1$, set $C_1^i := D_1^i \cap V(C(P \setminus v))$ and $C_2^i := D_2^i \cap V(C(P \setminus v))$. Then by Lemma 78, (i), (C_1^i, C_2^i) is a valid partition of $C(P \setminus v)$ w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$. Note that (C_1^0, C_2^0) was already defined.

We claim that all the partitions (C_1^i, C_2^i) must be distinct.

Indeed, suppose for a contradiction that there are $0 \leq i < j$ such that $(C_1^i, C_2^i) = (C_1^j, C_2^j)$.

By assumption, $D_1^j \supseteq C_1^j = C_1^i$ contains no copy of P. As (C_1^i, C_2^i) is valid, it follows by (C) that C_1^i and hence D_1^j contains a copy P' of $P \setminus v$. But $F_i \subseteq F_{j-1} = D_1^j \cup D_2^j$ contains the copy Q_1^i of $Q \setminus w$. As any vertex x of Q_1^i together with P' would induce the subgraph $P' \times x \supseteq P$, it follows that all vertices of Q_1^i are part of D_2^j. A similar argument shows that $V(Q_2^i) \subseteq V(D_2^j)$. But F_i also contains the copy P_1^i of $P \setminus v$, and $P_1^i \cup \{z_i\}$ forms a copy of P. As $P \not\subseteq D_1^j$, for some $x \in V(P_1^i) \cup \{z_i\}$ we have $x \in D_2^j$.

If $x \neq z_i$, then $Q \subseteq x \times Q_1^i \subseteq D_2^j$.

If $x = z_i$ then x and Q_2^i induce a copy of Q in D_2^j.

Thus in any case $Q \subseteq D_2^j$. But this is a contradiction with the validity of the partition (D_1^j, D_2^j).

Consequently all partitions (C_1^i, C_2^i) of $C(P \setminus v)$ must be distinct. There are at most $2^{v(C(P \setminus v))}$ such partitions, completing the proof. □

Let s be the length of the sequence. By Claim 1, $1 \leq s \leq 2^{v(C(P \setminus v))} \leq 2^{M_{v(P)-1}}$.

By induction on $0 \leq i \leq s-1$, we see that $v(F_i) \leq v(C(P \setminus v)) + (i+1)(2v(Q) - 2 + v(P))$. Therefore $v(C(P, Q, v, w)) \leq M_{v(P)-1} + 2^{M_{v(P)-1}}(2v(Q) - 2 + v(P)) = M_{v(P)} = M_{P,Q,v,w}$. This proves (A).

Define $D_1^0 := C_1^0$ and $D_2^0 := C_2^0$. Note that $C(P, Q, v, w) = F_{s-1}$ contains a copy K' of K given by P_1^{s-1}, Q_2^{s-1} and z_{s-1}, and $C(P, Q, v, w) \setminus K'$ has the partition $(D_1^{s-1}, D_2^{s-1} \dot\cup Q_1^{s-1})$.

It is easy to see that this partition is valid w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$, proving (B).

As the sequence stopped, any valid partition (D_1, D_2) of $C(P, Q, v, w)$ w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$ must satisfy $P \subseteq D_1$. If P is minimal invalid, this is a contradiction with Lemma 78, (ii) and (v), and therefore in this case no valid partition (D_1, D_2) can exist. This proves (C), finishing the proof of the lemma. □

3.4.6 The Collapsing Lemma

We continue to follow the strategy of Pikhurko from [96]. The next step is to prove a *collapsing lemma.* The general idea is to show that under certain circumstances two r-graphs which are extremal and "close" to one another must in fact be isomorphic (thus the extremal structure "collapses" onto a predefined pattern). This technique goes back to the work of Simonovits in the 60s [114] and was later developed as a tool for the exact determination of Turán densities.

To the best of our knowledge, all the previous applications of this method considered that one of the two r-graphs is a blow-up or an iterated blow-up structure. In such a situation many nice properties are available, most importantly, the addition of any new edge to this graph creates $\Omega(n^{v(F)-r})$ copies of some forbidden graph F. Thus even a small, local modification requires the deletion of many edges to maintain the property of being F-free, in particular the resulting graph can not be "close" to the initial one, and hence if it is "too close", it must be equal.

However, this is too much to expect in our present situation; for an arbitrary Turán density there is no structure to use, and there is no evidence that adding an edge to an extremal r-graph would create many copies of some forbidden subgraph.

Nevertheless in some cases local modifications create many forbidden r-graphs; this can be read out of the π function, and it was our goal in the previous section to extract this information. We shall use it to show that the number of edges in an extremal r-graph is suitably bounded by the function $\oplus_r$.

First we need a couple of definitions.

Let G and H be two r-graphs with the same number n of vertices. For $\varepsilon > 0$, we say that G and H are ε-close if G is isomorphic to an r-graph G' on $V(H)$ such that $|E(G')\Delta E(H)| \leq \varepsilon\binom{n}{r}$. In other words, we can obtain H from G by adding or deleting at most $\varepsilon\binom{n}{r}$ edges.

A family of r-graphs $\mathcal{F}$ is called *minimal* if it is weakly closed under homomorphisms and any r-graph $F \in \mathcal{F}$ is minimal invalid w.r.t. $\mathcal{F}$.

Lemma 81. *For any finite family of r-graphs $\mathcal{F}$ there exists a finite minimal family of r-graphs $\mathcal{F}'$ with $\pi(\mathcal{F}') = \pi(\mathcal{F})$.*

Proof. Clearly we may assume that each r-graph in $\mathcal{F}$ has at least one edge.

We shall repeatedly apply one of the following operations to $\mathcal{F}$.

(O1) If $F \in \mathcal{F}, f : F \to F'$ is a surjective homomorphism and $F' \notin \mathcal{F}$, then add F' to $\mathcal{F}$.

(O2) If $F \in \mathcal{F}$, $F' \notin \mathcal{F}$ is a proper subgraph of F with at least one edge and $\pi(\mathcal{F} \cup \{F'\}) = \pi(\mathcal{F})$, then add F' to $\mathcal{F}$.

(O3) If $F' \subsetneq F$ and $F', F \in \mathcal{F}$ then remove F from $\mathcal{F}$.

We start by applying (O1) and (O2) until none of these operations can be applied anymore, and then we apply (O3) as many times as possible. As (O1), (O2) and (O3) can each be applied only a finite number of times, we obtain a finite family $\mathcal{F}'$ of r-graphs.

We first claim that $\pi(\mathcal{F}') = \pi(\mathcal{F})$ and that any element of $\mathcal{F}'$ has at least one edge. To prove this we examine each operation separatedly.

Consider (O1). Let $F \in \mathcal{F}$ and suppose that there exists F' and a surjective homomorphism from F to F' such that $F' \notin \mathcal{F}$. Then for some $t \geq 1$, $F \subseteq F'(t)$. Hence by Lemma 74, (i), (ii) and (iv),

$$\pi(\mathcal{F}) \geq \pi(\mathcal{F} \cup \{F'\}) = \pi(\mathcal{F} \cup \{F'(t)\}) \geq \pi(\mathcal{F} \cup \{F\}) = \pi(\mathcal{F}),$$

proving that $\pi(\mathcal{F} \cup \{F'\}) = \pi(\mathcal{F})$. Thus we can add F' to $\mathcal{F}$ without changing the Turán density. Moreover F' has at least one edge, as it contains a homomorphic image of F.

Now consider (O2). Let $F \in \mathcal{F}$ and $F' \notin \mathcal{F}$ be a proper subgraph of F with at least one edge, such that $\pi(\mathcal{F} \cup \{F'\}) = \pi(\mathcal{F})$. Then we can add F' to $\mathcal{F}$, and this does not change the Turán density.

Finally, by Lemma 74, (i) and (ii), whenever $F' \subsetneq F$ and $F', F \in \mathcal{F}$, we can remove F from $\mathcal{F}$ without changing the Turán density.

Consequently it follows by induction that $\pi(\mathcal{F}') = \pi(\mathcal{F})$ and any element of $\mathcal{F}'$ has at least one edge. In particular, any $F \in \mathcal{F}'$ is invalid w.r.t. $\mathcal{F}'$.

We now claim that any r-graph in $\mathcal{F}'$ is also minimal invalid with respect to it. Indeed, suppose for a contradiction that there exists $F \in \mathcal{F}'$ and $x \in V(F)$ such that $F \setminus x$ is invalid w.r.t. $\mathcal{F}'$. Then $F \setminus x$ has at least one edge and $\pi(\mathcal{F}' \cup \{F \setminus x\}) = \pi(\mathcal{F}') = \pi(\mathcal{F})$. By (O2), $F \setminus x \in \mathcal{F}$ before the first application of (O3). But then (O3) forces the removal of F from $\mathcal{F}$, a contradiction.

Consequently any r-graph $F \in \mathcal{F}'$ is minimal invalid w.r.t. $\mathcal{F}'$.

Finally, we claim $\mathcal{F}'$ is weakly closed under homomorphisms. Indeed, let $F \in \mathcal{F}'$ and assume $f : F \to F'$ is a surjective homomorphism. Then F' was added to $\mathcal{F}$ by (O1) if it was not already present in $\mathcal{F}$. If $F' \notin \mathcal{F}'$, then it must have been removed by (O3). Consequently there exists $F'' \subsetneq F'$ such that $F'' \in \mathcal{F}'$. This shows that $\mathcal{F}'$ is weakly closed under homomorphisms. □

Note that if $\mathcal{F}$ is minimal and $\pi(\mathcal{F}) = 0$, then $\mathcal{F} = \{e\}$, where e is the r-edge.

Lemma 82 (The Collapsing Lemma). *Let $\mathcal{F}_\alpha$ and $\mathcal{F}_\beta$ be two non-empty finite minimal families of r-graphs with $\pi(\mathcal{F}_\alpha) = \alpha$ and $\pi(\mathcal{F}_\beta) = \beta$.*

Then there exists an $M_{\alpha,\beta} > 0$ such that for any $M \geq M_{\alpha,\beta}$ and any M-closures $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$ the following holds.

There exist a finite family of r-graphs $\mathcal{F}_{\alpha,\beta}$ and an $\varepsilon > 0$ with the following properties.

- *If H_1 is an $\mathcal{F}_\alpha(M)$-free r-graph and H_2 is an $\mathcal{F}_\beta(M)$-free r-graph then $H_1 \times H_2$ is $\mathcal{F}_{\alpha,\beta}$-free.*
- *Furthermore for any $\zeta > 0$ there exists an $n_0 \geq 1$ such that any maximum $\mathcal{F}_{\alpha,\beta}$-free r-graph G on $n \geq n_0$ vertices which is ε-close to an r-graph of the form $H_1 \times H_2$, with H_1 an $\mathcal{F}_\alpha(M)$-free r-graph, and H_2 an $\mathcal{F}_\beta(M)$-free r-graph, has at most*
$$(\alpha \oplus_r \beta + \zeta)\binom{n}{r} \tag{3.8}$$
edges.

Proof. For any pair $(F_\alpha, F_\beta) \in \mathcal{F}_\alpha \times \mathcal{F}_\beta$, and any choice of $v \in V(F_\alpha), w \in V(F_\beta)$, the Rigidity Lemma gives us a positive integer $M_{F_\alpha,F_\beta,v,w}$. We set $M_{\alpha,\beta}$ to be the maximum of these values.

Now let $M \geq M_{\alpha,\beta}$ be arbitrary and fix arbitrary M-closures $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$.

For any pair $(F_\alpha, F_\beta) \in \mathcal{F}_\alpha \times \mathcal{F}_\beta$, and any choice of $v \in V(F_\alpha), w \in V(F_\beta)$, the Rigidity Lemma now gives us an r-graph $C(F_\alpha, F_\beta, v, w)$. We set
$$\mathcal{F}^*_{\alpha,\beta} := \{C(F_\alpha, F_\beta, v, w) : F_\alpha \in \mathcal{F}_\alpha, F_\beta \in \mathcal{F}_\beta, v \in V(F_\alpha), w \in V(F_\beta)\}$$
and let $\mathcal{F}_{\alpha,\beta}$ be the closure of $\mathcal{F}^*_{\alpha,\beta}$ under homomorphisms.

For each $F_\alpha \in \mathcal{F}_\alpha$, choose an arbitrary $F_\beta \in \mathcal{F}_\beta$ along with an arbitrary $v \in V(F_\alpha)$ and $w \in V(F_\beta)$. With these choices, define the r-graph $C(F_\alpha) := C(F_\alpha, F_\beta, v, w)$. The exact choices we make are irrelevant for the argument to follow. Similarly for any $F_\beta \in \mathcal{F}_\beta$, choose $F_\alpha \in \mathcal{F}_\alpha, v \in V(F_\alpha)$ and $w \in V(F_\beta)$ arbitrary, and define the r-graph $C(F_\beta) := C(F_\alpha, F_\beta, v, w)$.

Before we go any further we establish the first part of the lemma.

Claim 2. *If H_1 is an $\mathcal{F}_\alpha(M)$-free r-graph and H_2 is an $\mathcal{F}_\beta(M)$-free r-graph then $H := H_1 \times H_2$ is $\mathcal{F}_{\alpha,\beta}$-free.*

Proof. Let $P \in \mathcal{F}_\alpha, Q \in \mathcal{F}_\beta$ and $v \in V(P), w \in V(Q)$. Let $f : C(P,Q,v,w) \to F$ be any surjective homomorphism. We show H is F-free.

Suppose for a contradiction that F embeds into $H = H_1 \times H_2$. Then there exists a partition (T_1, T_2) of $C(P,Q,v,w)$ such that $f(T_1) \subseteq H_1$ and $f(T_2) \subseteq H_2$. By Lemma 79, (C), this partition is not valid. Thus w.l.o.g. we may assume that T_1 is not valid w.r.t. $\mathcal{F}_\alpha(M)$.

Let $t \geq 1$ such that $T_1 \subseteq f(T_1)(t)$.

If $\alpha > 0$ then by Lemma 74, (iv) and (ii),

$$\pi(\mathcal{F}_\alpha(M) \cup \{f(T_1)\}) = \pi(\mathcal{F}_\alpha(M) \cup \{f(T_1)(t)\}) \geq \pi(\mathcal{F}_\alpha(M) \cup \{T_1\}) = \pi(\mathcal{F}_\alpha(M)).$$

Consequently $\pi(\mathcal{F}_\alpha(M) \cup \{f(T_1)\}) = \pi(\mathcal{F}_\alpha(M))$. However, $v(f(T_1)) \leq v(T_1) \leq v(C(P,Q,v,w)) \leq M_{\alpha,\beta} \leq M$ by definition. As $\alpha > 0$, T_1 and hence $f(T_1)$ certainly contains at least one edge. Thus by maximality, $f(T_1) \in \mathcal{F}_\alpha(M)$. This contradicts the fact that H_1 is $\mathcal{F}_\alpha(M)$-free.

Hence $\alpha = 0$. But then $\lambda(H_1) \geq \lambda(f(T_1)) \geq \lambda(T_1) > 0$ by Lemma 70, (ii). Consequently H_1 has at least one edge. But $\mathcal{F}_\alpha$ contains the one edge r-graph, again contradicting the fact that H_1 is $\mathcal{F}_\alpha$-free. This proves the claim. $\square$

We now prove the second part of the Collapsing Lemma. Define

$$\varepsilon_0 := \frac{1}{2} \min\{\alpha - \theta(\mathcal{F}_\alpha(M)), \beta - \theta(\mathcal{F}_\beta(M))\},$$

where recall that $\theta(\mathcal{F}(M))$ is the threshold of the M-closure $\mathcal{F}(M)$. By definition, $\theta(\mathcal{F}_\alpha(M)) = -1$ if $\alpha = 0$, and similarly $\theta(\mathcal{F}_\beta(M)) = -1$ if $\beta = 0$. Hence we always have $\varepsilon_0 > 0$.

Then define

$$\delta := \min\{\delta(\varepsilon_0, \mathcal{F}_\alpha(M)), \delta(\varepsilon_0, \mathcal{F}_\beta(M))\},$$
$$n^* := \max\{n^*(\varepsilon_0, \mathcal{F}_\alpha(M)), n^*(\varepsilon_0, \mathcal{F}_\beta(M))\}.$$

We are now going to define several other constants. Rather than giving a precise definition, we shall list a (admitedly long) list of inequalities they have to satisfy, and it will be obvious from this list that a choice satisfying all the given inequalities can be made. The reader may safely skip this part of the proof and return later when needed.

Note that $\alpha < 1$ and $\beta < 1$ by Lemma 74, (v), as $\mathcal{F}_\alpha$ and $\mathcal{F}_\beta$ are both non-empty.

Moreover, recall that in Lemma 67 we defined

$$\mathfrak{g}_{\alpha,\beta}(x) = \alpha x^r + \beta(1-x)^r + r! \sum_{j=1}^{r-1} \frac{x^j (1-x)^{r-j}}{j!(r-j)!}$$

and $x_{\alpha,\beta} = \frac{\sqrt[r-1]{1-\beta}}{\sqrt[r-1]{1-\alpha} + \sqrt[r-1]{1-\beta}}$.

Recall also the definition of τ in Lemma 76.

Now choose constants

$$0 < \varepsilon \ll c_1 \ll c_2 \ll c_3 \ll c_4 \ll c_5 \ll c_6$$

such that the following conditions hold:

$$\varepsilon < \min\left\{\frac{1-\beta}{3}\left(1-\frac{1}{(1+\sqrt[r-1]{1-\beta})^{r-1}}\right), \frac{1-\alpha}{3}\left(1-\frac{1}{(1+\sqrt[r-1]{1-\alpha})^{r-1}}\right)\right\}, \tag{3.9}$$

$$c_1 < \min\{x_{\alpha,\beta}, 1-x_{\alpha,\beta}\}, \tag{3.10}$$

$$\varepsilon < \frac{1}{4}\left(\alpha \oplus_r \beta - \max\{\mathfrak{g}_{\alpha,\beta}(x_{\alpha,\beta}-c_1), \mathfrak{g}_{\alpha,\beta}(x_{\alpha,\beta}+c_1)\}\right), \tag{3.11}$$

$$\varepsilon < \frac{c_2}{3}\min\left\{\left(\frac{\sqrt[r-1]{1-\alpha}}{1+\sqrt[r-1]{1-\alpha}}\right)^r, \left(\frac{\sqrt[r-1]{1-\beta}}{1+\sqrt[r-1]{1-\beta}}\right)^r\right\}, \tag{3.12}$$

$$\varepsilon^{1/2} < c_3 r \min\{(1-x_{\alpha,\beta}-2c_1)^{r-1}(x_{\alpha,\beta}-c_1), (x_{\alpha,\beta}-2c_1)^{r-1}(1-x_{\alpha,\beta}-c_1)\}, \tag{3.13}$$

$$\varepsilon < r!\frac{\delta^2}{8(r+1)}(1-c_3)^{M_{\alpha,\beta}}(x_{\alpha,\beta}-c_1)^r(1-x_{\alpha,\beta}-c_1)^r, \tag{3.14}$$

$$c_6 < (r-1)!\frac{\delta^2}{16M_{\alpha,\beta}}(1-c_3)^{M_{\alpha,\beta}}, \tag{3.15}$$

$$\varepsilon^{1/2} + c_3 r < c_6, \tag{3.16}$$

$$c_5 > c_3 r + c_4 + (4^r + 2^{r-1} + 2)c_1, \tag{3.17}$$

$$c_6 > \frac{(x_{\alpha,\beta}+c_1)^{r-1} - x_{\alpha,\beta}^{r-1} + c_5}{(x_{\alpha,\beta}+c_1)^{r-1}}, \tag{3.18}$$

$$c_2 + c_3 r < \min\{\alpha - \theta(\mathcal{F}_\alpha(M)) - \varepsilon_0, \beta - \theta(\mathcal{F}_\beta(M)) - \varepsilon_0\}, \tag{3.19}$$

$$c_2 + c_3 r < \min\{\tau(\mathcal{F}_\alpha, c_4), \tau(\mathcal{F}_\beta, c_4)\}. \tag{3.20}$$

In order for this system of inequalities to have a solution, it is enough if the following condition holds. For each inequality, the smaller quantity tends to zero when all the unknowns appearing in it tend to zero, and any unknown c_i appearing in the greater quantity has index strictly larger than any unknown c_j appearing in the smaller quantity. The only problematic inequality is (3.15); however after a rearrangement it can be seen that it also satisfies this condition.

Furthermore recall that if $\alpha = 0$ then $\theta(\mathcal{F}_\alpha(M)) = -1$, and hence if $\alpha = \beta = 0$ then $\varepsilon_0 = \frac{1}{2}$, and so the right-hand side of (3.19) is in this case equal to $\frac{1}{2}$.

Now let $\zeta > 0$ be arbitrary. We also choose constants

$$n^* \ll n_1 \ll n_2 \ll n_0$$

in the following way.

Recall the definition of n_C in Lemma 73. We require that

$$n_1 \geq n_C(\delta, \max\{v(F_\alpha), v(F_\beta)\}, v(C(F_\alpha, F_\beta, v, w))), \tag{3.21}$$

for any choice of $F_\alpha \in \mathcal{F}_\alpha, F_\beta \in \mathcal{F}_\beta, v \in V(F_\alpha)$ and $w \in V(F_\beta)$.

Recall the definition of n_V in Lemma 76. We require

$$n_1 \geq \max\{n_V(\mathcal{F}_\alpha, c_4), n_V(\mathcal{F}_\beta, c_4)\}, \tag{3.22}$$

and that for any $n \geq n_1$, $\mathrm{ex}(n, \mathcal{F}_\alpha) < (\alpha + \frac{\zeta}{4})\binom{n}{r}$ and $\mathrm{ex}(n, \mathcal{F}_\beta) < (\beta + \frac{\zeta}{4})\binom{n}{r}$.

Once n_1 is fixed, we choose n_2 such that

$$n_2 \geq \frac{n_1}{1 - c_3}, \tag{3.23}$$

and such that for any $n \geq n_2$, $\mathrm{ex}(n, \mathcal{F}_\alpha) < (\alpha + \varepsilon)\binom{n}{r}$ and $\mathrm{ex}(n, \mathcal{F}_\beta) < (\beta + \varepsilon)\binom{n}{r}$.

Finally, we choose n_0 such that

$$n_0 \geq \max\left\{\frac{n_2 r}{\varepsilon}, \frac{r}{c_1}, n_D(\mathcal{F}_{\alpha,\beta}, c_1)\right\}, \tag{3.24}$$

where n_D is given by Lemma 75 (the assumptions of the lemma are satisfied because $\mathcal{F}_{\alpha,\beta}$ is closed under homomorphisms), and such that

$$\prod_{i=1}^{r-1}\left(1 - \frac{i}{n_0}\right) \geq \max\left\{\frac{\alpha \oplus_r \beta - 2c_1}{\alpha \oplus_r \beta - c_1}, \frac{\alpha \oplus_r \beta - \varepsilon}{\alpha \oplus_r \beta}, \frac{\alpha \oplus_r \beta + \frac{\zeta}{2}}{\alpha \oplus_r \beta + \zeta}\right\}. \tag{3.25}$$

Note the dependency of n_0 and n_1 on ζ.

For the rest of the proof we shall assume that $\pi(\mathcal{F}_{\alpha,\beta}) \geq \alpha \oplus_r \beta$, otherwise we can choose n_0 large enough so that the lemma trivially holds.

Finally, let G be any maximum $\mathcal{F}_{\alpha,\beta}$-free r-graph G on $n \geq n_0$ vertices and suppose G is ε-close to an r-graph of the form $H := H_1 \times H_2$, with H_1 an $\mathcal{F}_\alpha(M)$-free r-graph, and H_2 an $\mathcal{F}_\beta(M)$-free r-graph. We shall assume for a contradiction that G has more than $(\alpha \oplus_r \beta + \zeta)\binom{n}{r}$ edges.

We identify the vertex set of G with that of H in such a way that $|E(G)\Delta E(H)| \leq \varepsilon\binom{n}{r}$. Let (G_1, G_2) be the partition of G so that $V(G_1) = V(H_1)$ and $V(G_2) = V(H_2)$.

The proof now begins in earnest. We start with the following claim.

Claim 3. *G_1 and G_2 have each at least n_2 vertices.*

Proof. We only prove that $v(G_1) \geq n_2$, as the other statement is proved similarly.

Assume for a contradiction that this is not the case. By (3.24), $n_0 \geq 2n_2$, in particular $v(G_2) \geq n_2$.

Hence

$$\begin{aligned}
e(G) &\leq e(G_2) + n_2\binom{n-1}{r-1} \\
&\leq e(H_2) + n_2\binom{n-1}{r-1} + \varepsilon\binom{n}{r} \\
&\leq (\beta + 3\varepsilon)\binom{n}{r}, \text{ by our choice of } n_2 \text{ and (3.24)}, \\
&< (\alpha \oplus_r \beta)\binom{n}{r},
\end{aligned}$$

where the last inequality follows from (3.9) and the inequality $\alpha \oplus_r \beta \geq 0 \oplus_r \beta = 1 - \frac{1-\beta}{(1+\sqrt[r-1]{1-\beta})^{r-1}}$.

This is a contradiction, proving the claim. □

Now set $a := v(G_1)$ and $b := v(G_2)$. Note that $a + b = n$.

Claim 4. *$\frac{a}{n} \in (x_{\alpha,\beta} - c_1, x_{\alpha,\beta} + c_1)$ and $\frac{b}{n} \in (1 - x_{\alpha,\beta} - c_1, 1 - x_{\alpha,\beta} + c_1)$.*

Proof. We only prove that $\frac{a}{n} \in (x_{\alpha,\beta} - c_1, x_{\alpha,\beta} + c_1)$, as the other statement follows from this one.

Suppose for a contradiction that this is not the case.

Note that

$$\begin{aligned}
e(G) &\leq e(H_1) + e(H_2) + \sum_{i=1}^{r-1} \binom{a}{i}\binom{b}{r-i} + \varepsilon\binom{n}{r} \\
&\leq (\alpha+\varepsilon)\binom{a}{r} + (\beta+\varepsilon)\binom{b}{r} + \sum_{i=1}^{r-1} \binom{a}{i}\binom{b}{r-i} + \varepsilon\binom{n}{r}, \text{by Claim 3 and our choice of } n_2, \\
&\leq \alpha\frac{a^r}{r!} + \beta\frac{b^r}{r!} + \sum_{i=1}^{r-1} \frac{a^i b^{r-i}}{i!(r-i)!} + 3\varepsilon\frac{n^r}{r!} \\
&= (\mathfrak{g}_{\alpha,\beta}(\frac{a}{n}) + 3\varepsilon)\frac{n^r}{r!}.
\end{aligned}$$

As $\mathfrak{g}_{\alpha,\beta}$ is strictly concave, it follows by (3.11) that

$$e(G) \leq (\alpha \oplus_r \beta - \varepsilon)\frac{n^r}{r!} \overset{(3.25)}{\leq} (\alpha \oplus_r \beta)\binom{n}{r},$$

a contradiction. □

Claim 5. *G_1 has density at least $\alpha - c_2$ and G_2 has density at least $\beta - c_2$.*

Proof. We only prove that $d(G_1) \geq \alpha - c_2$, as the other statement is proved similarly.

Assume for a contradiction that this is not the case.

As in the previous claim, note that

$$\begin{aligned}
e(G) &\leq d(G_1)\frac{a^r}{r!} + \beta\frac{b^r}{r!} + \sum_{i=1}^{r-1} \frac{a!b!}{i!(r-i)!} + 2\varepsilon\frac{n^r}{r!} \\
&= (\mathfrak{g}_{d(G_1),\beta}(\frac{a}{n}) + 2\varepsilon)\frac{n^r}{r!} \\
&\leq (d(G_1) \oplus_r \beta + 2\varepsilon)\frac{n^r}{r!}, \text{by Lemma 67,} \\
&\leq ((\alpha - c_2) \oplus_r \beta + 2\varepsilon)\frac{n^r}{r!} \\
&\leq \left(\alpha \oplus_r \beta + 2\varepsilon - c_2\left(\frac{\sqrt[r-1]{1-\beta}}{1+\sqrt[r-1]{1-\beta}}\right)^r\right)\frac{n^r}{r!}, \text{by Lemma 77,} \\
&\overset{(3.12)}{\leq} (\alpha \oplus_r \beta - \varepsilon)\frac{n^r}{r!} \\
&\overset{(3.25)}{\leq} (\alpha \oplus_r \beta)\binom{n}{r},
\end{aligned}$$

a contradiction. □

In fact, the same proof shows that $d(H_1) \geq \alpha - c_2$ and $d(H_2) \geq \beta - c_2$.

Define $S_1 := \{x \in V(G_1) : d_{G_2}(x) \leq (1-\varepsilon^{1/2})\binom{b}{r-1}\}$ and $S_2 := \{x \in V(G_2) : d_{G_1}(x) \leq (1-\varepsilon^{1/2})\binom{a}{r-1}\}$.

Claim 6. *$|S_1| \leq c_3 a$ and $|S_2| \leq c_3 b$.*

Proof. We only prove that $|S_1| \leq c_3 a$, as the other statement is proved similarly.

Note that $\varepsilon\binom{n}{r} \geq |E(H) \setminus E(G)| \geq |S_1|\varepsilon^{1/2}\binom{b}{r-1}$, hence

$$\begin{aligned}
|S_1| &\leq \varepsilon^{1/2}\binom{n}{r}\binom{b}{r-1}^{-1} \\
&\leq \varepsilon^{1/2}\frac{n^r}{r(b-r)^{r-1}} \\
&\leq \varepsilon^{1/2}\frac{n}{r(1-x_{\alpha,\beta}-c_1-\frac{r}{n})^{r-1}}, \text{ by Claim 4,} \\
&\overset{(3.24)}{\leq} \varepsilon^{1/2}\frac{n}{r(1-x_{\alpha,\beta}-2c_1)^{r-1}} \\
&\overset{(3.13)}{\leq} c_3(x_{\alpha,\beta}-c_1)n \\
&\leq c_3 a, \text{ by Claim 4.}
\end{aligned}$$

□

Define $G_1' := G_1 \setminus S_1$ and $G_2' := G_2 \setminus S_2$. Furthermore set $H_1' := H_1 \setminus S_1$ and $H_2' := H_2 \setminus S_2$. By (3.23), G_1' and G_2' have each at least n_1 vertices. Moreover,

$$d(H_1') \geq d(H_1) - c_3 r \geq \alpha - c_2 - c_3 r \overset{(3.19)}{>} \theta(\mathcal{F}_\alpha(M)) + \varepsilon_0, \tag{3.26}$$

and similarly

$$d(H_2') \geq \beta - c_2 - c_3 r \overset{(3.19)}{>} \theta(\mathcal{F}_\beta(M)) + \varepsilon_0. \tag{3.27}$$

Claim 7. *Let $V(G_1') \subseteq U$ and $V(G_2') \subseteq W$ be disjoint sets of vertices in G with the following property: $d_{G_2'}(x) \geq (1-c_6)\binom{b}{r-1}$ for any $x \in U$, and $d_{G_1'}(x) \geq (1-c_6)\binom{a}{r-1}$ for any $x \in W$. Then $G[U]$ is $\mathcal{F}_\alpha$-free and $G[W]$ is $\mathcal{F}_\beta$-free.*

Proof. We only prove that $G[U]$ is $\mathcal{F}_\alpha$-free, as the other statement is proved in a similar way.

Suppose for a contradiction that for some $F \in \mathcal{F}_\alpha$, there is a copy of F inside $G[U]$, which we also denote by F.

Recall that by the Addendum to the Rigidity Lemma, (b), $C(F)$ contains a copy K of F such that $C(F) \setminus K$ has a partition (C_1, C_2) valid w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$. Either of C_1 or C_2 can be empty, but not both, as then $C(F) \simeq F$, contradicting the fact that G is $\mathcal{F}_{\alpha,\beta}$-free. By our choice of δ, (3.26), (3.27) and the fact that $n_1 \geq n^*$, there are at least $\delta v(H_1')^{v(C_1)}$ copies of C_1 in H_1', and at least $\delta v(H_2')^{v(C_2)}$ copies of C_2 in H_2'. By Lemma 73 and (3.21), we can find $N \geq \frac{\delta^2}{2} v(H_1')^{v(C_1)} v(H_2')^{v(C_2)}$ embeddings of $C_1 \times C_2$ in $(H_1' \setminus V(F)) \times H_2'$, mapping C_1 into H_1' and C_2 into H_2'.

As $C(F)$ is not a subgraph of G, for each of the above embeddings $f : C_1 \times C_2 \to (H_1' \setminus V(F)) \times H_2'$, either one of the edges in $f(C_1 \times C_2)$ is in $E(H) \setminus E(G)$, or one of the edges intersecting $V(F)$ in one vertex and $f(C_2)$ in $r-1$ vertices is in $E(H) \setminus E(G)$.

Suppose first that in at least $\frac{N}{2}$ of the embeddings f, one of the edges in $f(C_1 \times C_2)$ is in $E(H) \setminus E(G)$.

Every edge e in $f(C_1 \times C_2)$ intersects H_1' in some $s_1(e) \geq 0$ vertices, and H_2' in some $s_2(e) \geq 0$ vertices, so that $s_1(e) + s_2(e) = r$. Thus there are s_1 and s_2 such that in at least $\frac{N}{2(r+1)}$ of the embeddings f, there is an edge in $f(C_1 \times C_2)$ intersecting H_1' in s_1 vertices, and H_2' in s_2 vertices, and this edge is in $E(H) \setminus E(G)$.

We count the number of such edges. It is at least

$$
\begin{aligned}
\frac{N}{2(r+1)a^{v(C_1)-s_1}b^{v(C_2)-s_2}} &\geq \frac{\delta^2}{4(r+1)}(1-c_3)^{v(C_1)+v(C_2)} \frac{a^{v(C_1)}b^{v(C_2)}}{a^{v(C_1)-s_1}b^{v(C_2)-s_2}} \\
&= \frac{\delta^2}{4(r+1)}(1-c_3)^{v(C_1)+v(C_2)} a^{s_1} b^{s_2} \\
&\geq \frac{\delta^2}{4(r+1)}(1-c_3)^{M_{\alpha,\beta}}(x_{\alpha,\beta}-c_1)^r(1-x_{\alpha,\beta}-c_1)^r n^r, \text{as } v(C(F)) \leq M_{\alpha,\beta}, \\
&\overset{(3.14)}{>} \varepsilon \frac{n^r}{r!}.
\end{aligned}
$$

Thus $|E(H) \setminus E(G)| > \varepsilon \binom{n}{r}$, a contradiction.

Consequently in at least $\frac{N}{2}$ of the embeddings f, one of the edges intersecting $V(F)$ in one vertex and $f(C_2)$ in $r-1$ vertices is in $E(H) \setminus E(G)$. Then for some $x \in V(F)$ we have

$$
\begin{aligned}
\binom{b}{r-1} - d_{G_2'}(x) &\geq \frac{N}{2v(F)a^{v(C_1)}b^{v(C_2)-(r-1)}} \\
&\geq \frac{\delta^2}{4v(F)}(1-c_3)^{v(C_1)+v(C_2)} \frac{a^{v(C_1)}b^{v(C_2)}}{a^{v(C_1)}b^{v(C_2)-(r-1)}} \\
&\geq \frac{\delta^2}{4M_{\alpha,\beta}}(1-c_3)^{M_{\alpha,\beta}} b^{r-1},
\end{aligned}
$$

a contradiction with (3.15) and our assumption that $d_{G_2'}(x) \geq (1-c_6)\binom{b}{r-1}$. $\square$

Every vertex $x \in V(G_1')$ has degree

$$
d_{G_2'}(x) \geq (1-\varepsilon^{1/2})\binom{b}{r-1} - |S_2|\binom{b-1}{r-2} \geq (1-\varepsilon^{1/2}-c_3 r)\binom{b}{r-1},
$$

and similarly for every $x \in V(G_2')$ we have $d_{G_1'}(x) \geq (1-\varepsilon^{1/2}-c_3 r)\binom{a}{r-1}$. As $\varepsilon^{1/2} + c_3 r < c_6$ by (3.16), it follows from Claim 7 that G_1' is $\mathcal{F}_\alpha$-free and G_2' is $\mathcal{F}_\beta$-free.

Claim 8. *For any $x \in S_1 \dot\cup S_2$, either $d_{G_1'}(x) \leq (\alpha + c_4)\binom{a}{r-1}$ or $d_{G_2'}(x) \leq (\beta + c_4)\binom{b}{r-1}$.*

Proof. Let $x \in S_1 \dot\cup S_2$ arbitrary and assume for a contradiction that $d_{G'_1}(x) > (\alpha + c_4)\binom{a}{r-1}$ and $d_{G'_2}(x) > (\beta + c_4)\binom{b}{r-1}$.

Then by Lemma 76, (3.22) and the fact that G'_1 is $\mathcal{F}_\alpha$-free, there exist $P \in \mathcal{F}_\alpha$, $v \in V(P)$ and a copy K_1 of $P \setminus v$ in G'_1 such that K_1 and x form a copy of P in G. Similarly, there exist $Q \in \mathcal{F}_\beta$, $w \in V(Q)$ and a copy K_2 of $Q \setminus w$ in G'_2 such that K_2 and x form a copy of Q in G.

Recall the description of $C(P, Q, v, w)$ given in Lemma 79, (B). According to this description, $C(P, Q, v, w)$ contains a copy of $P \setminus v$, which we also denote by K_1, and a copy of $Q \setminus w$, which we also denote by K_2, and a vertex z, such that $C(P, Q, v, w) \setminus (K_1 \cup K_2 \cup \{z\})$ has a partition (C_1, C_2) valid w.r.t. $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$. Again one of C_1 or C_2 can be empty, but not both, as then $C(P, Q, v, w) \subseteq G$, a contradiction with the fact that G is $\mathcal{F}_{\alpha,\beta}$-free.

By our choice of δ, (3.26), (3.27) and the fact that $n_1 \geq n^*$, there are at least $\delta v(H'_1)^{v(C_1)}$ copies of C_1 in H'_1, and at least $\delta v(H'_2)^{v(C_2)}$ copies of C_2 in H'_2. By Lemma 73 and (3.21), we can find $N \geq \frac{\delta^2}{4} v(H'_1)^{v(C_1)} v(H'_2)^{v(C_2)}$ embeddings $f : C_1 \times C_2 \to (H'_1 \setminus V(K_1)) \times (H'_2 \setminus V(K_2))$.

Together with K_1, K_2 and x, any such embedding would potentially form a copy of $C(P, Q, v, w)$ in G. Therefore as in the proof of Claim 7, we distinguish two cases.

First suppose there are s_1 and s_2 such that in at least $\frac{N}{2(r+1)}$ of the embeddings f, one of the edges in $f(C_1 \times C_2)$ is in $E(H) \setminus E(G)$, and intersects H'_1 in s_1 vertices, and H'_2 in s_2 vertices. A similar count to that in Claim 7 gives a contradiction.

Consequently in at least $\frac{N}{2}$ of the embeddings f, one of the edges intersecting K_1 in one vertex and $f(C_2)$ in $r-1$ vertices, or one of the edges intersecting K_2 in one vertex and $f(C_1)$ in $r-1$ vertices, is in $E(H) \setminus E(G)$.

Thus w.l.o.g. for some $y \in V(K_1)$ we obtain

$$\begin{aligned}\binom{b}{r-1} - d_{G'_2}(y) &\geq \frac{N}{4v(P)a^{v(C_1)}b^{v(C_2)-(r-1)}} \\ &\geq \frac{\delta^2}{16M_{\alpha,\beta}}(1-c_3)^{M_{\alpha,\beta}} b^{r-1} \\ &\overset{(3.15)}{>} c_6 \binom{b}{r-1},\end{aligned}$$

a contradiction with the fact that $y \notin S_1$ and hence $d_{G'_2}(y) \geq (1 - \varepsilon^{1/2} - c_3 r)\binom{b}{r-1} > (1 - c_6)\binom{b}{r-1}$. □

Now let $U := \{x \in S_1 \dot\cup S_2 : d_{G'_1}(x) \leq (\alpha + c_4)\binom{a}{r-1}\}$ and $W := S_1 \dot\cup S_2 - U$. By Claim 8, any $x \in W$ has $d_{G'_2}(x) \leq (\beta + c_4)\binom{b}{r-1}$.

Claim 9. *For any $x \in U, d_{G'_2}(x) \geq (1 - c_6)\binom{b}{r-1}$ and for any $x \in W, d_{G'_1}(x) \geq (1 - c_6)\binom{a}{r-1}$.*

Proof. We only prove the claim for $x \in W$, as the other statement is similar.

First note the following identity:

$$\alpha \oplus_r \beta - \beta(1 - x_{\alpha,\beta})^{r-1} - (r-1)! \sum_{j=1}^{r-2} \frac{x_{\alpha,\beta}^j (1 - x_{\alpha,\beta})^{r-1-j}}{j!(r-1-j)!} = x_{\alpha,\beta}^{r-1}. \tag{3.28}$$

Indeed, recall that

$$\alpha \oplus_r \beta = \mathfrak{g}_{\alpha,\beta}(x_{\alpha,\beta}) = 1 - (1-\alpha)x_{\alpha,\beta}^r - (1-\beta)(1-x_{\alpha,\beta})^r.$$

Also

$$(r-1)!\sum_{j=1}^{r-2} \frac{x_{\alpha,\beta}^j(1-x_{\alpha,\beta})^{r-1-j}}{j!(r-1-j)!} = 1 - x_{\alpha,\beta}^{r-1} - (1-x_{\alpha,\beta})^{r-1}.$$

Hence

$$\begin{aligned}
&\alpha \oplus_r \beta - \beta(1-x_{\alpha,\beta})^{r-1} - (r-1)!\sum_{j=1}^{r-2} \frac{x_{\alpha,\beta}^j(1-x_{\alpha,\beta})^{r-1-j}}{j!(r-1-j)!} \\
&\quad = 1 - (1-\alpha)x_{\alpha,\beta}^r - (1-\beta)(1-x_{\alpha,\beta})^r - \beta(1-x_{\alpha,\beta})^{r-1} - 1 + x_{\alpha,\beta}^{r-1} + (1-x_{\alpha,\beta})^{r-1} \\
&\quad = x_{\alpha,\beta}^{r-1} - (1-\alpha)x_{\alpha,\beta}^r + (1-\beta)(1-x_{\alpha,\beta})^{r-1}x_{\alpha,\beta} \\
&\quad = x_{\alpha,\beta}^{r-1} - \frac{(1-\alpha)(1-\beta)}{(\sqrt[r-1]{1-\alpha} + \sqrt[r-1]{1-\beta})^{r-1}}x_{\alpha,\beta} + \frac{(1-\alpha)(1-\beta)}{(\sqrt[r-1]{1-\alpha} + \sqrt[r-1]{1-\beta})^{r-1}}x_{\alpha,\beta} \\
&\quad = x_{\alpha,\beta}^{r-1},
\end{aligned}$$

proving (3.28).

By (3.24), Lemma 75 and our assumption that $\pi(\mathcal{F}_{\alpha,\beta}) \geq \alpha \oplus_r \beta$, we see that

$$d_G(x) \geq (\alpha \oplus_r \beta - c_1)\binom{n-1}{r-1} \tag{3.29}$$

for any $x \in V(G)$.

Now let $x \in W$ arbitrary. Using (3.29) and the fact that $|S_1 \dot\cup S_2| \le c_3 n$ we obtain

$$\begin{aligned}
d_{G_1'}(x) &\overset{(3.24)}{\geq} (\alpha \oplus_r \beta - c_1)\binom{n-1}{r-1} - (\beta + c_4)\binom{b}{r-1} - \sum_{j=1}^{r-2}\binom{a}{j}\binom{b}{r-1-j} - c_3 r\binom{n-1}{r-1} \\
&\overset{(3.25)}{\geq} (\alpha \oplus_r \beta - 2c_1 - (\beta + c_4)(1 - x_{\alpha,\beta} + c_1)^{r-1} \\
&\qquad - (r-1)!\sum_{j=1}^{r-2}\frac{(x_{\alpha,\beta}+c_1)^j(1-x_{\alpha,\beta}+c_1)^{r-1-j}}{j!(r-1-j)!} - c_3 r)\frac{n^{r-1}}{(r-1)!} \\
&\geq (\alpha \oplus_r \beta - \beta(1-x_{\alpha,\beta})^{r-1} - (2+2^{r-1})c_1 - c_4 - c_3 r \\
&\qquad - (r-1)!\sum_{j=1}^{r-2}\frac{x_{\alpha,\beta}^j(1-x_{\alpha,\beta})^{r-1-j} + 3\cdot 2^{r-1}c_1}{j!(r-1-j)!})\frac{n^{r-1}}{(r-1)!} \\
&\overset{(3.17)}{\geq} (\alpha \oplus_r \beta - \beta(1-x_{\alpha,\beta})^{r-1} - (r-1)!\sum_{j=1}^{r-2}\frac{x_{\alpha,\beta}^j(1-x_{\alpha,\beta})^{r-1-j}}{j!(r-1-j)!} - c_5)\frac{n^{r-1}}{(r-1)!} \\
&\overset{(3.28)}{=} (x_{\alpha,\beta}^{r-1} - c_5)\frac{n^{r-1}}{(r-1)!} \\
&\overset{(3.18)}{\geq} (1-c_6)(x_{\alpha,\beta}+c_1)^{r-1}\frac{n^{r-1}}{(r-1)!} \\
&\geq (1-c_6)\binom{a}{r-1}.
\end{aligned}$$

This proves the claim. □

Define $U' := V(G_1') \cup U$ and $W' := V(G_2') \cup W$. Then by Claims 7 and 9, $G[U']$ is $\mathcal{F}_\alpha$-free and $G[W']$ is $\mathcal{F}_\beta$-free. But (U', W') is a partition of G. Setting $a' := |U'|$ and $b' := |W'|$, we obtain

$$\begin{aligned}
e(G) &\leq e(G[U']) + e(G[V']) + \sum_{j=1}^{r-1}\binom{a'}{j}\binom{b'}{r-j} \\
&\leq \frac{1}{r!}(\alpha a'^r + \beta b'^r + r!\sum_{j=1}^{r-1}\frac{a'^j b'^{r-j}}{j!(r-j)!} + \frac{\zeta n^r}{2}), \quad \text{by our choice of } n_1, \\
&\leq (\alpha \oplus_r \beta + \frac{\zeta}{2})\frac{n^r}{r!} \\
&\overset{(3.25)}{\leq} (\alpha \oplus_r \beta + \zeta)\binom{n}{r}.
\end{aligned}$$

This finishes the proof of the Collapsing Lemma. □

3.4.7 End of the proof

We are now ready to finish the proof of Theorem 33. There is only one further ingredient that we need.

Theorem 83 (Strong Removal Lemma, [103]). *For every family $\mathcal{F}$ of r-graphs and any $\varepsilon > 0$ there exist δ, m and n_R such that the following holds. If G is any r-graph on $n \geq n_R$ vertices which contains at most $\delta n^{v(F)}$ copies of any r-graph $F \in \mathcal{F}$ with $v(F) \leq m$, then G can be made $\mathcal{F}$-free by removing at most $\varepsilon\binom{n}{k}$ edges.*

The Removal Lemma for hypergraphs is a deep and rather recent result. Its origins can be traced back to the 70s, in the (now famous) Triangle Removal Lemma of Ruzsa and Szemerédi [108], though it was only in the last decade that a suitable version for hypergraphs was obtained, independently by Gowers [52] and by Nagle, Rödl, Schacht and Skokan ([87], [104], [105]). Subsequently generalizations and other versions were proved. We remark that the original Removal Lemma is stated in terms of a single hypergraph; we crucially need here a version applicable to an *infinite* family of hypergraphs.

Proof of Theorem 33. Let $\alpha, \beta \in \Pi^{(r)}_{\text{fin}}$. In view of Lemma 69, to finish the proof we only need to show that $\alpha \oplus_r \beta \in \Pi^{(r)}_{\text{fin}}$. Clearly we may assume that $\alpha, \beta \neq 1$.

Choose a finite family of r-graphs $\mathcal{F}_\alpha$ with $\pi(\mathcal{F}_\alpha) = \alpha$. By Lemma 81, we may assume that $\mathcal{F}_\alpha$ is minimal. As $\alpha \neq 1$, $\mathcal{F}_\alpha$ is non-empty.

Similarly we can choose a finite non-empty minimal family of r-graphs $\mathcal{F}_\beta$ with $\pi(\mathcal{F}_\beta) = \beta$.

On input $\mathcal{F}_\alpha$ and $\mathcal{F}_\beta$, the Collapsing Lemma gives a positive integer $M := M_{\alpha,\beta}$. Choose and fix arbitrary M-closures $\mathcal{F}_\alpha(M)$ and $\mathcal{F}_\beta(M)$. The Collapsing Lemma now gives us a finite family of r-graphs $\mathcal{F}_{\alpha,\beta}$ and an $\varepsilon > 0$.

Let $\{G^1_n\}_{n\geq 1}$ be any sequence of $\mathcal{F}_\alpha(M)$-free r-graphs with $v(G^1_n) = n$ and $d(G^1_n) \to \alpha$. Such a sequence exists, even in the case $\alpha = 0$, as any r-graph in $\mathcal{F}_\alpha(M)$ has at least one edge. Similarly, let $\{G^2_n\}_{n\geq 1}$ be any sequence of $\mathcal{F}_\beta(M)$-free r-graphs with $v(G^2_n) = n$ and $d(G^2_n) \to \beta$.

Consider the sequence $\{G_n\}_{n\geq 1}$ with $G_n := G^1_{x_{\alpha,\beta}n} \times G^2_{(1-x_{\alpha,\beta})n}, n \geq 1$ (we disregard lower and upper integer parts here, as it does not affect our proof). Then $d(G_n)$ converges to $\alpha \oplus_r \beta$. Moreover G_n is $\mathcal{F}_{\alpha,\beta}$-free by the first part of the Collapsing Lemma, for any $n \geq 1$.

Define $\mathcal{F}_\infty := \{F : F \nsubseteq G_n, \forall n \geq 1\}$. Then $\mathcal{F}_{\alpha,\beta} \subseteq \mathcal{F}_\infty$. Apply the Strong Removal Lemma to $\mathcal{F}_\infty$ and $\frac{\varepsilon}{2}$ to obtain δ (which we disregard), m and n_R.

Finally, set $\mathcal{F}_m := \mathcal{F}_{\alpha,\beta} \cup \{F \in \mathcal{F}_\infty : v(F) \leq m\}$.

We claim $\pi(\mathcal{F}_m) = \alpha \oplus_r \beta$.

Clearly $\pi(\mathcal{F}_m) \geq \alpha \oplus_r \beta$, as the sequence $\{G_n\}_{n\geq 1}$ shows.

Let $\zeta > 0$ arbitrary. We show $\pi(\mathcal{F}_m) \leq \alpha \oplus_r \beta + \zeta$.

The Collapsing Lemma gives us an $n_0 \geq 1$. Let G be any maximum $\mathcal{F}_m$-free r-graph on $n \geq \max\{n_0, n_R\}$ vertices. Then G can be made $\mathcal{F}_\infty$-free by removing at most $\frac{\varepsilon}{2}\binom{n}{r}$ edges. Let G' be the resulting r-graph. As $G' \notin \mathcal{F}_\infty$, there exists $k \geq 1$ such that $G' \subseteq G_k$. Let H be the subgraph of G_k isomorphic with G'. Now $H' := G_k[V(H)]$ is also $\mathcal{F}_m$-free and hence $e(G) \geq e(H') \geq e(G')$. Thus $|E(H') \setminus E(H)| \leq \frac{\varepsilon}{2}\binom{n}{r}$. Consequently G is ε-close to H'. But H', being an induced subgraph, is of the form $H_1 \times H_2$, with H_1 an $\mathcal{F}_\alpha(M)$-free r-graph, and H_2 an $\mathcal{F}_\beta(M)$-free r-graph.

By the Collapsing Lemma, G has at most $(\alpha \oplus_r \beta + \zeta)\binom{n}{r}$ edges. Thus $\pi(\mathcal{F}_m) \leq \alpha \oplus_r \beta + \zeta$. As ζ was arbitrary, the proof is finished. $\square$

We briefly highlight some of the difficulties we had to overcome in the above proof.

The general strategy was taken from the proof of Theorem 3 in [96]. While the Removal Lemma allows us to force any maximum $\mathcal{F}_{\alpha,\beta}$-free r-graph to be close to the desired structure by adding some more forbidden r-graphs, it can not be made arbitrarily close. This requires the proof of a Collapsing Lemma.

The most serious obstacle appears when we pass from the sequence $\{G_n\}_{n\geq 1}$ to the induced subgraph H'. While H' is a subgraph of some G_k, the number of vertices of H' is not (in any way) bounded from below by the number of vertices of G_k. Thus properties of graphs in the sequence $\{G_n\}_{n\geq 1}$ need not pass to H'. In particular, if G_k contains $\Omega(v(G_k)^{v(F)})$ copies of some r-graph F, H' can very well have only a few such copies, or none at all. In order to overcome this obstacle we used information hidden in the function π. This is reflected in the Rigidity Lemma, which is not stated for a particular graph sequence, but more generally in terms of two families of r-graphs.

Remark. One can extract a proof of Theorem 28 from the above arguments. Let $r \geq 2$ and $t \geq r$. We take $\alpha = \pi(K_t^r), \beta = 0$ with $\mathcal{F}_\alpha = \{K_t^r\}, \mathcal{F}_\beta = \{e\}$ in the proof of Theorem 33. For this particular choice one can take $M = 1$ and $\mathcal{F}_{\alpha,\beta} = \{K_{t+r-1}^r\}$ in the Collapsing Lemma. Indeed, if H_1 is a K_t^r-free r-graph and H_2 is an $\mathcal{F}_\beta$-free r-graph (thus a graph with no edges), then by the pigeonhole principle $H_1 \times H_2$ is K_{t+r-1}^r-free. By taking $C(K_t^r) = C(e) = C(K_t^r, e, v, w) = K_{t+r-1}^r$, the rest of the proof of Lemma 82 goes through. The additional family $\mathcal{F}_{r,t}$ comes from the application of the Strong Removal Lemma.

3.5 Towards a semiring structure

As we have seen in the Introduction, the set $\mathfrak{h}_2(\Pi_\infty^{(2)} \setminus \{1\}) = \mathbb{N}$ has a natural semiring structure under addition and multiplication. We call pull back these two operations to $\Pi_\infty^{(2)} = \{1\} \cup \{1 - \frac{1}{k} : k \geq 1\}$, similar to what we did in Section 1.2.3. Formally we define for all $a, b \geq 1$ and $\alpha \in \Pi_\infty^{(2)}$,

$$(1 - \frac{1}{a}) \oplus_2 (1 - \frac{1}{b}) = 1 - \frac{1}{a+b},$$
$$\alpha \oplus_2 1 = 1 \oplus_2 \alpha = 1,$$

and

$$(1 - \frac{1}{a}) \otimes_2 (1 - \frac{1}{b}) = 1 - \frac{1}{ab},$$
$$\alpha \otimes_2 1 = 1 \otimes_2 \alpha = 1.$$

Algebraically this means $\oplus_2$ maps (α, β) to the real number $1 - \frac{(1-\alpha)(1-\beta)}{1-\alpha+1-\beta}$, while $\otimes_2$ maps (α, β) to the real number $\alpha + \beta - \alpha\beta$, for any $\alpha, \beta \in \Pi_\infty^{(2)} \setminus \{1\}$. Note that we have already successfully generalized the operation $\oplus_2$ to any $r \geq 2$. We shall try to do the same with $\otimes_2$. Unfortunately, this time the natural approach fails.

As in the case of $\oplus_r$, we must find a corresponding operation on r-graphs. There would be many advantages if $\oplus_2$ would hold unchanged for any r, in particular $\otimes_2$ is the pull-back of real multiplication under $\mathfrak{h}_r$, and we would obtain a semiring structure on $\Pi_\infty^{(r)}$ as desired. As far as we can see, there is only one natural construction associated to this operation (we keep in mind the concrete examples given by $r = 2$ and graph cliques).

Let G and H be two r-graphs. We define an r-graph $G \otimes H$ in the following way.

Assume w.l.o.g. that G has vertex set $[n]$. The vertex set of $G \otimes H$ consists of n disjoint copies $V_1, \ldots, V_n$ of the vertex set of H. If v is a vertex of H, we let $v_i \in V_i$ be its i-th copy. We add the following edges to $G \otimes H$.

For all $h = (v^1, \ldots, v^r) \in E(H)$, we add all edges f with $|f \cap \{v_1^t, \ldots, v_n^t\}| = 1, 1 \leq t \leq r$. Furthermore for all $e = (i_1, \ldots, i_r) \in E(G)$, we add all edges f with $|f \cap V_{i_j}| = 1, 1 \leq j \leq r$. No other edges are added.

This is a generalization of the strong product of graphs to uniform hypergraphs. Informally, a simple way to understand $G \otimes H$ is to think of it as a 2-dimensional construction: the axis are labelled with the vertices of G, respectively H, and an edge is part of $G \otimes H$ if either its projection on the G-axis, or its projection on the H-axis, is an edge.

Note that if $r = 2$, G is an n-clique[1] and H is an m-clique then $G \otimes H$ is an mn-clique. This corresponds to our objective and hence we would like to prove the following.

Target 1. *For any two r-graphs G and H we have $\lambda(G \otimes H) = \lambda(G) + \lambda(H) - \lambda(G)\lambda(H)$.*

Target 1 would imply via continuity that $\Pi_\infty^{(r)}$ is closed under $\otimes_2$. We further have the following.

[1] i.e. a complete 2-graph on n vertices.

Lemma 84. *For any two r-graphs G and H we have $\lambda(G \otimes H) \geq \lambda(G) + \lambda(H) - \lambda(G)\lambda(H)$. Furthermore, if $r = 2$ then equality holds here.*

Proof. Assume G has vertex set $[n]$ and H has vertex set $[m]$. We identify the vertex set of $G \otimes H$ with $[nm]$: the vertex set of the i-th copy of H runs from $(i-1)m+1$ to im.

If $\mathbf{a} \in \Delta_n$ is an optimal vector for G and $\mathbf{b}$ is an optimal vector for H, then considering the vector $\mathbf{c} \in \Delta_{nm}$ with $c_{(i-1)m+j} := a_i b_j, 1 \leq i \leq n, 1 \leq j \leq m$, we see that

$$\lambda(G \otimes H) \geq p_{G \otimes H}(\mathbf{c}) \geq \lambda(G) + \lambda(H) - \lambda(G)\lambda(H),$$

as claimed.

Now assume that $r = 2$. We know that $\lambda(G) = 1 - \frac{1}{a}$, where $a \geq 1$ is the size of the largest clique in G. Similarly, $\lambda(H) = 1 - \frac{1}{b}$, where $b \geq 1$ is the size of the largest clique in H.

Let S be a largest clique in $G \otimes H$. We claim $|S| \leq ab$. Indeed, if S contains two vertices $im + u$ and $im + v, 1 \leq u \neq v \leq m$, then u and v must be adjacent in H by construction. Define

$$P_H := \{u \in V(H) : im + u \in S \text{ for some } 0 \leq i < n\}.$$

Then P_H is a clique in H, and hence $|P_H| \leq b$. Furthermore, if S contains two vertices $im+u$ and $jm + v, i \neq j$, then i and j are adjacent in G. Hence

$$P_G := \{i \in V(G) : im + u \in S \text{ for some } 1 \leq u \leq m\}$$

is a clique in G, of maximum size a. Consequently $|S| \leq |P_H||P_G| \leq ab$.

But then

$$\lambda(G \otimes H) \leq 1 - \frac{1}{ab} = 1 - \frac{1}{a} + 1 - \frac{1}{b} - (1 - \frac{1}{a})(1 - \frac{1}{b}),$$

as desired. □

Despite Lemma 84, for $r \geq 3$, Target 1 is false.

3.5.1 A permanent détour

Let us look at the special case when both G and H are one edge r-graphs. Then $\lambda(G) = \lambda(H) = \frac{r!}{r^r}$ and we would like that $\lambda(G \otimes H) = \frac{2r!}{r^r} - \frac{(r!)^2}{r^{2r}}$. This claim is highly non-trivial; it is equivalent to a statement known as Dittert's conjecture.

First let us state the former van der Waerden conjecture, now a theorem.

Theorem 85 (Egorychev, Falikman, 1981)**.** *The minimum permanent among all $n \times n$ doubly stochastic matrices is $\frac{n!}{n^n}$, and is achieved only by the matrix with all entries equal to $1/n$.*

Theorem 85 was conjectured by van der Waerden in 1926. After attracting a lot of interest and a series of partial results, it was finally proved independently by Egorychev and Falikman in 1981, and it is still considered a milestone result in combinatorics, with many applications across the field. A different proof was recently found by Gurvits [59].

Several other similar conjectures have been made over time. The following conjecture is due to E. Dittert.

Conjecture 86 (Dittert, 1983, [84]). *Let A be a non-negative $n \times n$ matrix with row sums $r_1, \ldots, r_n$ and column sums $c_1, \ldots, c_n$. Suppose $\sum_{i=1}^n r_i = \sum_{i=1}^n c_i = 1$. Then the function*

$$\psi(A) := \prod_{i=1}^n r_i + \prod_{i=1}^n c_i - \operatorname{per}(A)$$

has the maximum $\frac{2}{n^n} - \frac{n!}{n^{2n}}$, and it is achieved only by the matrix J_n with all entries equal to $1/n^2$.

Conjecture 86 clearly implies Theorem 85. There is substantial evidence towards Conjecture 86. It was proved for $n = 2$ by Sinkhorn [115] and for $n = 3$ by Hwang [67]. Hwang further showed in [66] that if the ψ-maximising matrix is positive then it must equal J_n, and that J_n is a strict local maximum of ψ. Other partial results were obtained by Cheon and Yoon [25] and Cheon and Wanless [24]. Most importantly Cheon and Wanless [24] showed that the maximum value of ψ is exponentially close to the conjectured value.

Theorem 87 (Cheon-Wanless, [24]). *For any non-negative $n \times n$ matrix A with the sum of all elements equal to 1 we have $\psi(A) < \psi(J_n) + O(n^{4-n}e^{2n})$.*

Aside from the unicity of the maximum, it is easy to see that Conjecture 86 is equivalent (with $n = r$) to our claim $\lambda(e \otimes e) = \frac{2r!}{r^r} - \frac{(r!)^2}{r^{2r}}$, where e is an r-edge. Thus any proof of this claim must give a proof of the notoriously hard van der Waerden conjecture.

Let us now note that Target 1 implies a much stronger statement. Write $e^{\otimes k}$ for $e \otimes \ldots \otimes e$ (k times). Then we would like that

$$\lambda(e^{\otimes k}) = \binom{k}{1}\frac{r!}{r^r} - \binom{k}{2}\frac{(r!)^2}{r^{2r}} + \ldots + (-1)^{k-1}\binom{k}{k}\frac{(r!)^k}{r^{kr}}, \tag{3.30}$$

To see this, imagine the r-graph $e^{\otimes k}$ as an $r \times r \times \ldots \times r$ k-dimensional matrix. We want that $\lambda(e^{\otimes k})$ equals $p_{e^{\otimes k}}(\frac{1}{r^k}, \ldots, \frac{1}{r^k})$. We evaluate this polynomial by inclusion-exclusion: first project onto a single coordinate and sum up all the terms after that coordinate; then consider two coordinates and subtract the terms that were added twice etc. Equivalently one can expand the conjectured identity $1 - \lambda(e^{\otimes k}) = (1 - \lambda(e))^k$.

While studying hashing, Hajek made the following conjecture.

Conjecture 88 (Hajek, 1987, [60]). *Let k and n be positive integers and for $1 \leq j \leq k$, let S_j denote the collection of subsets L of $[n]^k$ such that L has cardinality n and no two distinct elements of L have the same j coordinate. Let $S := \cup S_j$ and define the multinomial $F_{n,k}(x) := \sum_{L \in S} \prod_{i \in L} x_i$. Then on Δ_{n^k}, $F_{n,k}$ attains its maximum only at $(\frac{1}{n^k}, \ldots, \frac{1}{n^k})$.*

Aside from the unicity of the maximum, Conjecture 88 is clearly equivalent to (3.30) with $n = r$.

Unfortunately Conjecture 88 was shown to be false for $n = 3$ and $k = 4$ by Körner and Marton [77]. This in particular disproves Target 1. The counterexample is a construction using the tetra-code, a self-dual code in $\mathbb{F}_3^4$. Körner and Marton also give an upper bound to (their equivalent notion of) $\lambda(e^{\otimes k})$, using graph entropy.

Theorem 89 (Körner-Marton, [77]). *For any r and k we have that*

$$\log_2 \frac{1}{1-\frac{r!}{r^r}} \leq \frac{1}{k}\log_2 \frac{1}{1-\lambda(e^{\otimes k})} \leq \frac{r!}{r^{r-1}}.$$

The lower bound follows from the construction given before. Interestingly, as shown in [77], any improvement on these bounds would most likely give an improvement on the best known bounds for the perfect hashing problem in a special case.

3.5.2 On Conjecture 32

Most of our interest in Target 1 stems from the following.

Proposition 90. *If $\Pi_\infty^{(r)}$ is closed under $\otimes_2$ for all $r \geq 2$ then $\overline{\cup_{r\geq 2}\Pi_{\text{fin}}^{(r)}} = \overline{\cup_{r\geq 2}\Pi_\infty^{(r)}} = [0,1]$.*

Proof. As $\Pi_{\text{fin}}^{(r)}$ is dense in $\Pi_\infty^{(r)}$ for all r, we only need to prove that $\overline{\cup_{r\geq 2}\Pi_\infty^{(r)}} = [0,1]$.

Fix $\alpha \in (0,1)$ and let $\varepsilon > 0$ arbitrary. We prove that there exists $\gamma \in \cup_{r\geq 2}\Pi_\infty^{(r)}$ with $|\gamma - \alpha| < \varepsilon$.

Write $\alpha = 1 - \frac{1}{\ell}$ for some real $\ell > 1$. Let $\delta > 0$ such $\delta\ell < \varepsilon$. As $\frac{r!}{r^r} \in \Pi_\infty^{(r)}$ (see the discussion after Theorem 59) and $\frac{r!}{r^r} \to 0$, there exists r and s with $1 \leq s < 1+\delta$ and $s < \ell$ such that $\beta := 1 - \frac{1}{s} \in \Pi_\infty^{(r)}$. Then for any $n \geq 1$, $\beta^{\otimes n} := \beta \otimes_2 \beta \otimes_2 \ldots \otimes_2 \beta$ (n times) is an element of $\Pi_\infty^{(r)}$, by assumption. However, $\beta^{\otimes n} = 1 - \frac{1}{s^n}$ by definition.

Choose n such that $s^n \leq \ell \leq s^{n+1}$. Then $n \geq 1$ and

$$\ell - s^n \leq s^{n+1} - s^n \leq \delta s^n \leq \delta\ell < \varepsilon.$$

Consequently

$$|\beta^{\otimes n} - \alpha| = \frac{\ell - s^n}{\ell s^n} < \varepsilon,$$

as $\ell s^n \geq 1$. This proves the claim. □

The proof could still be carried over if Target 1 would only hold for hypergraphs of the form $e^{\otimes k}$, where e is an r-edge. However, as we have seen, this is not the case. Nevertheless, this made us propose Conjecture 32, which was later proved by Pikhurko [97].

3.5.3 Some more results

One can lift $\otimes_2$ to Π_∞, and in this setting the law holds.

More precisely, define a binary operation $\circ$ on the set $\mathbb{R} \times \mathbb{N}$ (which contains Π_∞) as follows:

$$\begin{aligned} \circ\colon (\mathbb{R}\times\mathbb{N}) \times (\mathbb{R}\times\mathbb{N}) &\to \mathbb{R}\times\mathbb{N} \\ (\alpha, r) \times (\beta, s) &\mapsto ((\alpha + \beta - \alpha\beta)\binom{r+s}{r}\frac{r^r s^s}{(r+s)^{r+s}}, r+s). \end{aligned}$$

By using a similar trick as in Theorem 29 one can prove the following result.

Theorem 91. $(\Pi_\infty, \circ)$ *is a commutative cancellative semigroup.*

The associated construction is the following. If G is an r-graph and H is an s-graph on disjoint vertex sets, we define $G \circ H$ as the $(r+s)$-multigraph on vertex set $V(G) \cup V(H)$ and edge set

$$\{e \cup f : e \in E(G), f \in V(H)^{(s)}\} \cup \{e \cup f : e \in V(G)^{(r)}, f \in E(H)\}$$

The proof then proceeds similarly to that of Theorem 29, and so we shall not present it here. Unfortunately $\circ$ and $*$ do not define a ring structure on Π_∞.

In fact other relations concerning Turán densities can be obtained, though none seem to define any interesting algebraic structure. As an example, we have the following theorem.

Theorem 92. *For any* $r \geq 2$ *define the map* $\mathfrak{j} : [0,1] \to [0,1]$ *by* $\mathfrak{j}(x) = \left(\frac{r-1}{r-x}\right)^{r-1}$. *Then* $\mathfrak{j}(\Pi_\infty^{(r)}) \subsetneq \Pi_\infty^{(r)}$.

Again the construction is the only important step of the proof. For any r-graph G, define $\mathfrak{j}(G)$ as the r-multigraph on vertex set $\{v\} \cup V(G)$ (v is a vertex not belonging to G) and edge set $E(G) \cup \{\{v\} \cup e : e \in V(G)^{(r-1)}\}$. Then one can show that $\lambda(\mathfrak{j}(G)) = \mathfrak{j}(\lambda(G))$ and Theorem 92 follows by continuity.

3.6 Proof of Theorems 26 and 27

The proof of Theorem 26 requires a result from number theory.

Theorem 93 ([19])**.** *Let $r_i > 0$ be roots of rationals (i.e. $r_i^{n_i} \in \mathbb{Q}$ with $n_i \in \mathbb{N}$) for i in a finite indexing set I. Suppose*

$$\sum_{i \in I} q_i r_i = q \in \mathbb{Q}$$

for positive rationals q_i. Then each r_i is rational.

Proof of Theorem 26. Let $r \geq 3$. We already know that $\frac{r!}{r^r} \in \Pi_{\text{fin}}^{(r)}$. Hence by Theorem 33,

$$\frac{r!}{r} \oplus_r 0 = 1 - \frac{r^{r-1} - (r-1)!}{\left(r + \sqrt[r-1]{r^{r-1} - (r-1)!}\right)^{r-1}} \in \Pi_{\text{fin}}^{(r)}.$$

This number is rational if and only if $q := \left(r + \sqrt[r-1]{r^{r-1} - (r-1)!}\right)^{r-1}$ is rational.

Assume for a contradiction that q is rational. Let $r_i := (r^{r-1} - (r-1)!)^{\frac{i}{r-1}}$ for $0 \leq i \leq r-1$. Then

$$q = \sum_{i=0}^{r-1} \binom{r-1}{i} r^{r-1-i} r_i.$$

All r_i are positive roots of rationals. Hence by Theorem 93, we obtain that r_i is rational for all $0 \leq i \leq r-1$. In particular, r_1 is a natural number.

As $r \geq 3$, we have $r_1 > 1$. So we can find a prime divisor $p|r_1$. Then $p^{r-1}|r^{r-1} - (r-1)!$. Hence $1 < p < r$ and so $p|(r-1)!$.

Thus $p|r^{r-1}$. But then $p|r$ and so certainly $p^{r-1}|r^{r-1}$. As p^{r-1} divides $r^{r-1} - (r-1)!$, it must also divide $(r-1)!$. But it is well known that for any prime p, the power of p dividing $(r-1)!$ is

$$\left\lfloor \frac{r-1}{p} \right\rfloor + \left\lfloor \frac{r-1}{p^2} \right\rfloor + \ldots < \frac{r-1}{p-1} \leq r-1.$$

This is a contradiction, completing the proof. □

For the proof of Theorem 27 we shall need the following result of Sidorenko.

Theorem 94 (Sidorenko, [112])**.** $1 - \frac{1}{2^p} \in \Pi_{\text{fin}}^{(2k)}$ *for any $k, p \geq 1$.*

Proof of Theorem 27. Let $r \geq 4$ even. By Theorem 94, $\frac{1}{2} \in \Pi_{\text{fin}}^{(r)}$. Consequently by Theorem 33,

$$\frac{1}{2} \oplus_r 0 = 1 - \frac{1}{(1 + \sqrt[r-1]{2})^{r-1}} \in \Pi_{\text{fin}}^{(r)}.$$

This number is rational if and only if $(1 + \sqrt[r-1]{2})^{r-1}$ is rational.

However, $x_0 := \sqrt[r-1]{2}$ has the minimal polynomial $f(x) = x^{r-1} - 2$ ($f(x)$ is irreducible by Eisenstein's criterion). Consequently if $(1 + x_0)^{r-1}$ equals some rational $\frac{s}{t}$, then $f(x)$ must divide $t(1+x)^{r-1} - s$. Then $tf(x) = t(1+x)^{r-1} - s$, which is not possible as $r \geq 4$. Thus $(1 + x_0)^{r-1}$ is irrational, completing the proof. □

3.7 Proof of Theorems 30 and 31

Proof of Theorem 31. Let x_n be a sequence of elements in $\Pi_\infty^{(r)}$ converging to $c\frac{r!}{r^r}$ from above. Then $(1,1) * (x_n, r) \in \Pi_\infty^{(r+1)}$ by Theorem 29 (here we abuse the notation slightly, as the result of $*$ is a pair, and not a number). Hence

$$\lim_{n\to\infty} (1,1) * (x_n, r) = \lim_{n\to\infty} x_n \frac{r^r}{(r+1)^r} = c\frac{(r+1)!}{(r+1)^{r+1}}.$$

Thus $c\frac{(r+1)!}{(r+1)^{(r+1)}}$ is not a jump for $(r+1)$-graphs, and the claim follows by induction. □

For the proof of Theorem 30 we shall need the following result of Steinhaus ([119], *Théorème VII*).

Theorem 95 (Steinhaus, 1920)**.** *If $A, B \subset \mathbb{R}$ are sets of positive measure then $A+B$ contains an open interval.*

Proof of Theorem 30. We show that (i) implies (ii).

Set $A := \Pi_\infty^{(r)} \setminus \{1\}$ and note that A is a Borel set (as $\Pi_\infty^{(r)}$ is closed) and still has positive Lebesgue measure. Furthermore A is a semigroup under $\oplus_r$.

By definition, $\mathfrak{h}_r$ is a homeomorphism between $[0,1)$ and $[1,+\infty)$. The inverse of $\mathfrak{h}_r$ is $\mathfrak{h}_r^{-1}(x) = 1 - \frac{1}{x^{r-1}}$, which has first-order derivative $(\mathfrak{h}_r^{-1})'(x) = (r-1)\frac{1}{x^r}$. This is bounded on $[1,+\infty)$. As A is Borel, by Lemma B.11, $\mathfrak{h}_r(A)$ has positive Lebesgue measure.

Hence by Theorem 95, $\mathfrak{h}_r(A) + \mathfrak{h}_r(A)$ contains an open interval. As $(\mathfrak{h}_r(A), +)$ is a semigroup, $\mathfrak{h}_r(A)$ contains an open interval. Consequently A contains an open interval, proving (ii).

We show that (ii) implies (iv).

Suppose $\Pi_\infty^{(r)}$ contains an open interval. Multiplying with $(1,1)$ as in the proof of Theorem 31, we obtain an open interval in $\Pi_\infty^{(r+1)}$. Then (iv) follows by induction.

As any open interval has positive measure, (iv) implies (iii). Finally, (iii) implies (i) trivially. □

3.8 Some open problems

We end the chapter with several open problems, which we now discuss.

3.8.1 The set of all Turán densities

In view of our definition of Π_∞, one can ask what is the set $\cup_{r\geq 2}\Pi_\infty^{(r)}$. We could not even solve the following.

Problem 1. *Prove or disprove that* $\limsup_{r\to\infty}\Pi_\infty^{(r)} = \cup_{r\geq 2}\Pi_\infty^{(r)}$.

Here the limit is taken under the discrete metric, that is, an element belongs to $\limsup_{r\to\infty}\Pi_\infty^{(r)}$ if and only if it belongs to $\Pi_\infty^{(r)}$ for infinitely many r. By Theorem 94, $1-\frac{1}{2^p}\in\limsup_{r\to\infty}\Pi_\infty^{(r)}$ for any $p\geq 1$, and to the best of our knowledge no other values from this set have been determined. Moreover Sidorenko's proof of Theorem 94 does not generalize to other Turán densities ([74]).

3.8.2 Polynomials preserving Turán densities

By Theorem 29, the polynomial $\frac{1}{2^{2r}}\binom{2r}{r}x^2$ takes values in $\Pi_\infty^{(2r)}$ when evaluated at an element of $\Pi_\infty^{(r)}$. The following question remains open.

Problem 2. *For any $r\geq 3$ find a polynomial $f\in\mathbb{Q}[x]$ such that for any Turán density α for r-graphs, $f(\alpha)$ is also a Turán density for r-graphs.*

For $r=2$ one such polynomial is $2x-x^2$ (indeed, this is nothing else than our rule $\otimes_2$). Moreover an example of a rational function with the required properties is given by Theorem 92.

3.8.3 The algebraic degree of Turán densities

As the reader recalls, one of our objectives was Question 1. We have not been able to resolve it, though $\oplus_r$ prompts the following question.

Problem 3. *For some $r\geq 3$, find $\alpha\in\Pi_{\text{fin}}^{(r)}$ algebraic with minimal polynomial of degree greater than $r-1$, or show that none exists.*

3.8.4 Other finiteness theorems

In view of Theorem 33 it is natural to ask the following question.

Question 2. *Do Theorems 29, 91 and 92 have finite counterparts? That is, is Π_{fin} (respectively $\Pi_{\text{fin}}^{(r)}$) closed under the described operations?*

3.8.5 The Hausdorff dimension of $\Pi_\infty^{(r)}$

Recall the map $\mathfrak{h}_r : [0,1) \to [1,+\infty)$ defined by $\mathfrak{h}_r(x) = \left(\frac{1}{1-x}\right)^{1/(r-1)}$. It is an isomorphism between $A := \Pi_\infty^{(r)} \setminus \{1\}$ and a subsemigroup of $(\mathbb{R},+)$. As $\mathfrak{h}_r^{-1}$ is Lipschitz, if $\Pi_\infty^{(r)}$ has positive Hausdorff dimension then so does $\mathfrak{h}_r(A)$. What can we say about $\mathfrak{h}_r(A)$ in this case?

Recall that a subset of $\mathbb{R}$ is called analytic if it is the continous image of some Borel set in some Euclidean space $\mathbb{R}^n$.

Proposition 96. *Let $\mathbb{G}_r$ be the subgroup of $\mathbb{R}$ generated by $\mathfrak{h}_r(A)$ under addition. Then $\mathbb{G}_r$ is an analytic set and for any $r \geq 3$, it is dense in $\mathbb{R}$.*

Proof. $\mathbb{G}_r$ is generated by an analytic set and so it must be analytic too.

Furthermore for $r \geq 3$, $\mathbb{G}_r$ contains $\mathbb{Z}$ (as it contains 1) and $n\alpha, n \geq 1$, where α is some irrational number. By Diophantine approximation, $\mathbb{G}_r$ is dense in $[0,1]$, and hence dense in $\mathbb{R}$. □

It was a question of Erdős and Volkmann [39] if there exist subrings of $\mathbb{R}$ which are Borel sets and have Hausdorff dimension strictly between 0 and 1. This question was resolved by Edgar and Miller in 2003 (a discrete version was proved independently by Bourgain [13]).

Theorem 97 (Edgar-Miller, 2003, [30]). *If $E \subseteq \mathbb{R}$ is a subring and a Borel (or analytic) set then either E has Hausdorff dimension 0 or $E = \mathbb{R}$.*

This implies the following.

Proposition 98. *Suppose $\Pi_\infty^{(r)}$ is closed under $\otimes_2$. Then $\mathbb{G}_r$ is a subring of $\mathbb{R}$. If $\Pi_\infty^{(r)}$ has positive Hausdorff dimension then $\mathbb{G}_r = \mathbb{R}$.*

Proof. If $\Pi_\infty^{(r)}$ is closed under $\otimes_2$ then $\mathfrak{h}_r(A)$ is a semigroup under real multiplication. As any element of $\mathbb{G}_r$ is of the form $\alpha - \beta$, with $\alpha, \beta \in \mathfrak{h}_r(A)$, $\mathbb{G}_r$ must be closed under multiplication as well. As $1 \in \mathbb{G}_r$, $\mathbb{G}_r$ is a subring and an analytic set.

If $\Pi_\infty^{(r)}$ has positive Hausdorff dimension, $\mathbb{G}_r$ has too, and hence by Theorem 97, $\mathbb{G}_r = \mathbb{R}$. □

This might help in resolving the following two problems.

Problem 4. *Is $\Pi_\infty^{(r)}$ closed under $\otimes_2$ for $r \geq 3$?*

Problem 5. *Compute the Hausdorff dimension of $\Pi_\infty^{(r)}$ or at least determine if it is zero.*

3.8.6 Revisiting the case $r = 2$

It is a consequence of the Erdős-Stone-Simonovits theorem that

$$\Pi_{\text{fin}}^{(2)} = \Pi_\infty^{(2)} = \{1\} \cup \{1 - \frac{1}{k} : k \geq 1\}. \tag{3.31}$$

Consider the following problem.

Problem 6. *Find a proof of* (3.31) *without relying on the Erdős-Stone-Simonovits theorem, and generalize it as much as possible to* $r \geq 3$.

Here is a short proof. By Turán's theorem, $\{1 - \frac{1}{k} : k \geq 1\} \subset \Pi^{(2)}_{\text{fin}}$. Moreover, $\pi(\emptyset) = 1 \in \Pi^{(2)}_{\text{fin}}$. On the other hand, by the results of Brown and Simonovits, $\Pi^{(2)}_{\text{fin}} \subseteq \Pi^{(2)}_{\infty} \subseteq \overline{\Lambda}^{(2)}$. It is well-known that $\Lambda^{(2)} = \{1\} \cup \{1 - \frac{1}{k} : k \geq 1\}$, and so (3.31) holds.

This argument is in some sense unsatisfactory, as it relies on the exact computation of $\Lambda^{(2)}$, a feat which we can not hope to reproduce for $r \geq 3$. Nevertheless, we have the following.

Proposition 99. *Let* $r \geq 2$ *and suppose* $\Pi^{(r)}_{\text{fin}}$ *is closed under* $\oplus_r$, $\Pi^{(r)}_{\infty}$ *is closed under* $\otimes_2$, *and the subgroup* $\mathbb{G}_r$ *of* $(\mathbb{R}, +)$ *generated by* $\mathfrak{h}(\Pi^{(r)}_{\infty} \setminus \{1\})$ *is not dense in* $\mathbb{R}$. *Then* $\Pi^{(r)}_{\text{fin}} = \Pi^{(r)}_{\infty} = \{1\} \cup \{1 - \frac{1}{k^{r-1}} : k \geq 1\}$.

Of course, the conclusion of Proposition 99 does not hold for $r \geq 3$, which is not a contradiction, as $\mathbb{G}_r$ is dense in $\mathbb{R}$ for $r \geq 3$. The point of Proposition 99 is that rather weak algebraic and topological properties are enough to determine the full structure of $\Pi^{(r)}_{\text{fin}}$ and $\Pi^{(r)}_{\infty}$.

Proof. By continuity and the fact that $\Pi^{(r)}_{\infty} \subseteq \overline{\Pi}^{(r)}_{\text{fin}}$, $\Pi^{(r)}_{\infty}$ is also closed under $\oplus_r$. By Proposition 98, if $\Pi^{(r)}_{\infty}$ is closed under $\otimes_2$ then $\mathbb{G}_r$ is a subring of $\mathbb{R}$. It is well-known and easy to prove that a subgroup of $\mathbb{R}$ which is not dense must be cyclic. Hence $\mathbb{G}_r \cap (0, +\infty)$ has a smallest element a. As $a^2 \in \mathbb{G}_r$, it follows that $a = 1$. Thus $\mathbb{G}_r = \mathbb{Z}$, hence $\Pi^{(r)}_{\infty} = \{1\} \cup \{1 - \frac{1}{k^{r-1}} : k \geq 1\}$.

Now $0 \in \Pi^{(r)}_{\text{fin}}$, and therefore the subsemigroup generated by 0 under $\oplus_r$ belongs to $\Pi^{(r)}_{\text{fin}}$. This subsemigroup is exactly $\{1 - \frac{1}{k^{r-1}} : k \geq 1\}$, and consequently $\Pi^{(r)}_{\text{fin}} = \Pi^{(r)}_{\infty}$, proving the claim. $\square$

We know that $\Pi^{(2)}_{\text{fin}}$ is closed under $\oplus_2$ by Theorem 33, and $\Pi^{(2)}_{\infty}$ is closed under $\otimes_2$. Moreover $\mathbb{G}_r$ is dense in $\mathbb{R}$ for all $r \geq 3$ (Proposition 96), but I can not show that $\mathbb{G}_2$ is not dense without relying on the Erdős-Stone-Simonovits theorem. Furthermore my proof that $(\Pi^{(2)}_{\infty}, \otimes_2)$ is a semigroup (Lemma 84) relies on the ability to compute $\lambda(G)$, where G is any 2-graph. It would be interesting to find a different proof of these two claims, and possibly find other sets of hypotheses which imply (3.31).

Chapter 4

The Graceful Tree Conjecture

The results presented in this Chapter are joint work with Anna Adamaszek, Michał Adamaszek, Peter Allen and Jan Hladký.

4.1 Notation and auxiliary results

If G is any graph, the *order* of G is the number of vertices of G.

We write $a = b \pm \varepsilon$ when we have $a \in [b - \varepsilon, b + \varepsilon]$. Extending this, and in a slight abuse of notation, we write $a \pm \delta = b \pm \varepsilon$ for the inclusion $[a - \delta, a + \delta] \subseteq [b - \varepsilon, b + \varepsilon]$.

If (T, r) is a rooted tree, with root r, we let $V_{\mathrm{odd}}(T)$ denote the set of vertices of T at odd distance from r, and $V_{\mathrm{even}}(T)$ the set of vertices at even distance. Thus $V_{\mathrm{odd}}(T)$ and $V_{\mathrm{even}}(T)$ form a bipartition of T. A vertex from $V_{\mathrm{odd}}(T)$ is called *primary*, while a vertex from $V_{\mathrm{even}}(T)$ is called *secondary*.

We shall need an estimate on the tail of the distribution of a sum of independent random variables. A proof of the following result can be found in [86].

Theorem 100 (Chernoff's bound)**.** *Let $X_1, X_2, \dots, X_n$ be independent Bernoulli trials such that $\mathbf{P}[X_i = 1] = p_i$, $1 \le i \le n$. If $X := \sum_{i=1}^n X_i, \mu := \mathbf{E}[X] = \sum_{i=1}^n p_i$ and $0 < \delta < 1$ then*

$$\mathbf{P}[|X - \mu| > \delta\mu] < 2\exp(-\mu\delta^2/4) \ . \tag{4.1}$$

Let us consider a product probability space $\Omega = \prod_{i=1}^k \Omega_i$. Note that any element of Ω is a vector with coordinate i sampled from Ω_i. A measurable function $f : \Omega \to \mathbb{R}$ is said to be *C-Lipschitz* if $|f(\underline{x}) - f(\underline{x}')| \le C$ whenever the vectors $\underline{x}, \underline{x}' \in \Omega$ differ only in a single coordinate. The following result states that Lipschitz functions are strongly concentrated.

Lemma 101 (McDiarmid's Inequality, [83])**.** *Let $f : \Omega \to \mathbb{R}$ be a C-Lipschitz function defined on a product probability space $\Omega = \prod_{i=1}^k \Omega_i$. Then for each $t > 0$ we have*

$$\mathbf{P}\big[|f - \mathbf{E}[f]| > t\big] \le 2\exp\left(-\frac{2t^2}{C^2 k}\right) \ .$$

We now recall the notion of superdependency graphs as is needed for our last probabilistic tool, Suen's inequality. Let $\{\mathcal{B}_i\}_{i \in I}$ be a finite collection of events in an arbitrary probability space Ω. A *superdependency graph* for the events $\{\mathcal{B}_i\}_{i \in I}$ is an arbitrary graph whose vertices are I, and whose edges satisfy the following condition. Whenever $I_1, I_2 \subset I$ are two disjoint sets with no edge crossing from I_1 to I_2, then any Boolean combination of the events $\{\mathcal{B}_i\}_{i\in I_1}$ is independent of any Boolean combination of the events $\{\mathcal{B}_i\}_{i\in I_2}$. The fact that i and j form an edge in a superdependency graph (which is always clear from the context) is denoted by $i \sim j$. The following inequality states that if we have a sparse superdependency graph, then the probability of no $\mathcal{B}_i$ occurring is roughly as if the events were independent.

Lemma 102 (Suen's Inequality, [124], see also [3, p. 128])**.** *Using the above setting, define* $M = \prod \mathbf{P}[\overline{\mathcal{B}_i}]$*, and for* $i \sim j$*,*

$$\nu_{i,j} = \frac{\mathbf{P}[\mathcal{B}_i \wedge \mathcal{B}_j] + \mathbf{P}[\mathcal{B}_i]\mathbf{P}[\mathcal{B}_j]}{\prod_{\ell \,\sim\, i \text{ or } \ell \,\sim\, j}(1 - \mathbf{P}[\mathcal{B}_\ell])} \ . \tag{4.2}$$

Then we have

$$\left|\mathbf{P}\Big[\bigwedge \overline{\mathcal{B}_i}\Big] - M\right| \le M \cdot \left(\exp\big(\sum_{i\sim j} \nu_{i,j}\big) - 1\right) \ .$$

4.2 Proof of Theorem 38

Before we embark on the proof, we outline the main ideas. We wish to embed our tree in a random process, during which we will keep track of the sets of vertex and induced edge labels which remain available. In order to make our argument work, one property we will need is that these available label sets stay fairly uniformly distributed on $[(1+\varepsilon)n]$ throughout. Observe that if uv is an edge of T, if we choose $\psi(u)$ a uniformly random vertex label, and then choose $\psi(v)$ randomly 'close to' $n - \psi(u)$ then the distributions of each of $\psi(u)$, $\psi(v)$ and the induced edge label $\big|\psi(u) - \psi(v)\big|$ are close to the uniform distribution; thus we would like to embed our tree such that 'most' edges and vertices are embedded in about this way.

This suggests the following strategy for gracefully labelling a forest F whose components are small. For each tree in F, we label the bipartition classes as 'red' and 'blue'. We pick uniformly at random an integer $t \in [(1+\varepsilon)n]$, and give the 'red' vertices labels chosen independently at random 'close to' t, and the blue vertices labels chosen independently at random close to $(1+\varepsilon)n - t$. It is again easy to check that this strategy picks both vertex labels and induced edge labels fairly uniformly in $[(1+\varepsilon)n]$ (though there is substantial dependence between labels from a given tree of F) and since the choices for different trees are independent and there are many trees in F, it is not hard to show that we end up using both vertex and edge labels fairly uniformly. This strategy results in 'conflicts', that is, pairs of vertices or of edges receiving the same labels: but provided $v(F)$ is much less than n it is easy to check that the number of conflicts is much less than $v(F)$.

Our strategy to label T is a version of Rödl's celebrated nibble method. We first break T into a forest with small components by removing a few edges of T. We next choose a collection F_1 of these small trees, with $v(F_1)$ a small fraction of n, and embed it as described above. Now we choose a second collection F_2 of the remaining trees, and follow the same strategy. However this time we label the 'red' vertices' randomly with labels 'close to' t which weren't used in labelling F_1, and similarly the blue vertices randomly with labels which neither themselves conflict with F_1 nor induce conflicts with the edge labels used on F_1. We repeat this procedure for R rounds, at which time we have labelled all of $V(T)$.

In order to analyse this strategy for the second and subsequent rounds, we need to assume that the available edge and vertex labels maintain certain 'quasirandomness' properties (see Definition 105) which roughly state that the available edge and vertex labels remain close to uniformly distributed on $[(1+\varepsilon)n]$, and that when we label 'blue' vertices the set from which we choose randomly is never very small. We will see that our strategy indeed preserves these properties.

Once we complete this procedure, we have an 'almost-graceful' labelling of $V(T)$. There are conflicted edges and vertices, but there are $\ll n$ such conflicts by construction: when we label each F_i we create $\ll v(F_i)$ conflicts, plus we made no attempt to avoid conflicts on the edges we removed from T to obtain a forest. We would like to 'repair' these conflicts by re-labelling the $\ll n$ vertices which are either themselves conflicted or which are incident to a conflicted edge. Unfortunately, it is not clear how to do this; it could easily be that for some vertex which is conflicted, re-labelling it with any of the available vertex labels creates a conflict at an incident edge.

The solution to this problem is to modify our strategy slightly. Before we begin labelling T, we choose uniformly at random small sets of the vertex and edge labels, which we call a 'repair pair', and we mark these as unavailable throughout our labelling process (which otherwise does not change). It is not hard to show that removing small random sets of edge and vertex labels leaves available sets for the first round which are quasirandom, so that the labelling strategy duly completes. It is furthermore easy to check that, for any collection of at most Δ distinct integers $a_1, \dots, a_s$ in $[(1+\varepsilon)n]$, there are many labels a in the vertex repair set such that each $|a_i - a|$ is in the edge repair set. In fact, there are so many that we can use this property to greedily re-label vertices of $V(T)$ at which there is a conflict (either of the vertex label or an incident edge label) using distinct vertex and edge labels from the repair set. We then have by construction a graceful labelling of T, as desired.

4.2.1 General setup

We now describe the general setup which we use for proving Theorem 38 and Theorem 39.

Setup 103. *Given $n, \frac{1}{2} > \varepsilon > 0$ and $\Delta \geq 2$ we define*

$$m = 2\lfloor(1+\varepsilon)n/2\rfloor + 1, \tag{4.3}$$

$$d := \frac{\varepsilon}{16}, \quad \varepsilon' := \min\left\{\frac{3\varepsilon}{2(1+\varepsilon)}, \frac{d^{\Delta+1}}{16\Delta^2}\right\}, \quad \varepsilon'' := \min\left\{\frac{\varepsilon}{2}, \frac{\varepsilon'}{3}\right\},$$

$$\gamma := \frac{\varepsilon''}{960\Delta} \quad \text{and} \quad R := \max\left\{\frac{153600\Delta}{\varepsilon''^{\Delta+3}}, \frac{3}{\varepsilon'}\right\}.$$

Set $\ell := \lceil \gamma n \rceil$.

Let $\mathbb{A} = ([0, \frac{m}{2} - \ell] \cup [\frac{m}{2} + \ell, m]) \cap \mathbb{Z}$, $\mathbb{A}^+ = [-\ell, m+\ell] \cap \mathbb{Z}$. Let $\mathbb{C} = [2\ell, m] \cap \mathbb{Z}$ and $\mathbb{C}^+ = [1, m+2\ell+1] \cap \mathbb{Z}$.

Note that $|\mathbb{A}| = |\mathbb{C}| = m - 2\ell + 1$ and $|\mathbb{A}^+| = |\mathbb{C}^+| = m + 2\ell + 1$, as $\frac{m}{2} \notin \mathbb{Z}$. Furthermore the absolute difference of any two distinct elements in $\mathbb{A}^+$ is in $\mathbb{C}^+$.

For $a \in \mathbb{A}^+$ and $c \in \mathbb{C}^+$ arbitrary we define

$$a \oplus c = \begin{cases} a + c, & \text{if } a < \frac{m}{2}, \\ a - c, & \text{if } a > \frac{m}{2}. \end{cases}$$

Let $\mathcal{I}$ denote the collection of integer intervals of length ℓ which have non-empty intersection with $\mathbb{A}$. The purpose of the careful rounding when setting $m \approx (1+\varepsilon)n$ in (4.3) is that each $I \in \mathcal{I}$ intersects only one of the sets $[0, \frac{m}{2} - \ell] \cap \mathbb{Z}$ and $[\frac{m}{2} + \ell, m] \cap \mathbb{Z}$, and the leftmost elements of all the intervals from $\mathcal{I}$ form the set $[-\ell+1, m] \cap \mathbb{Z}$. In particular,

$$|\mathcal{I}| = \big|[-\ell+1, m] \cap \mathbb{Z}\big| = m + \ell. \tag{4.4}$$

Let $x \in \mathbb{Z}$. For the interval $I = [x, x+\ell-1]$ denote by $\bar{I}$ the complementary interval *$[m+1-x-\ell, m-x]$. Note that due to our definition of $\mathbb{A}$ we have $I \cap \bar{I} = \emptyset$ whenever $I \in \mathcal{I}$. Further, $I \in \mathcal{I}$ if and only if $\bar{I} \in \mathcal{I}$.*

Finally, given a pair $(A^+, C^+) \subseteq (\mathbb{A}^+, \mathbb{C}^+)$ we always write $A := A^+ \cap \mathbb{A}$ and $C := C^+ \cap \mathbb{C}$.

We shall consider all these definitions fixed throughout the paper. Thus whenever ε, Δ and n are given, we can refer to m, ℓ, $\mathbb{A}$ or any of the other objects defined above.

Although rather technical, we state below the quasirandomness definition that we use throughout the paper.

Definition 104. *Suppose $(A^+, C^+) \subseteq (\mathbb{A}^+, \mathbb{C}^+)$ is a pair of sets, $I \in \mathfrak{I}$ and $a_1, \ldots, a_s \in I \cap A^+$. We say that a number $a \in \mathbb{Z}$ is a* special solution *to $a_1, \ldots, a_s$ with respect to (A^+, C^+, I) if and only if $a \in \bar{I} \cap A^+$ and $|a - a_i| \in C^+$ for all $i = 1, \ldots, s$.*

The set of special solutions to $a_1, \ldots, a_s$ with respect to (A^+, C^+, I) will be denoted $\mathrm{Sol}(A^+, C^+, I; a_1, \ldots, a_s)$*, or simply* $\mathrm{Sol}(I; a_1, \ldots, a_s)$ *if the pair (A^+, C^+) is fixed.*

Note that $\bar{\bar{I}} = I$ and therefore a is a special solution to $a_1, \ldots, a_s$ with respect to (A^+, C^+, I) if and only if each a_i is a special solution to a with respect to $(A^+, C^+, \bar{I})$.

Definition 105. *(Quasirandomness conditions) A pair of sets $(A^+, C^+) \subseteq (\mathbb{A}^+, \mathbb{C}^+)$ is (α, ℓ, Δ)-quasirandom if*

a) (density of vertex labels) For all $I \in \mathfrak{I}$ we have

$$|I \cap A^+| = (1 \pm \alpha) \cdot |I| \cdot \frac{|A|}{|\mathbb{A}|}.$$

b) (uniformity of stars) For all $I \in \mathfrak{I}, p \in [\Delta]$ and $a_1, \ldots, a_p \in I \cap A^+$ pairwise distinct,

$$|\mathrm{Sol}(I; a_1, \ldots, a_p)| = (1 \pm \alpha) \cdot |\bar{I}| \cdot \frac{|A|}{|\mathbb{A}|} \cdot \left(\frac{|C|}{|\mathbb{C}|}\right)^p.$$

c) (uniformity of edge labels) For all $c \in C$ we have

$$|\{(a, I) \mid I \in \mathfrak{I}, a \in A^+ \cap I, a \oplus c \in A^+ \cap \bar{I}\}| = (1 \pm \alpha) \cdot \ell^2 \cdot \frac{|A|^2}{|\mathbb{A}|^2}.$$

Note that condition b) also extends to not necessarily distinct labels $a_1, \ldots, a_p$. Indeed, if q is the number of pairwise distinct labels among $a_1, \ldots, a_p$, and without lack of generality $a_1, \ldots, a_q$ are all distinct, then by definition of Sol we have $\mathrm{Sol}(I; a_1, \ldots, a_p) = \mathrm{Sol}(I; a_1, \ldots, a_q)$.

4.2.2 Almost graceful labelings

As we shall see, it is enough to prove a slightly weaker version of Theorem 38, Theorem 107 below. To state this relaxation, we introduce *almost graceful labelings*, in which we allow a small number of conflicts (i.e., slight non-injectivity of the labelling map on vertices, and on edges), and also we allow a small number of vertices to receive a “joker label” $*$.

Definition 106. *Suppose that* $\psi : V(H) \to \mathbb{N} \cup \{*\}$ *is given. The* deficiency *of* ψ *is defined as*

$$(v(H) - |\mathrm{im}(\psi) \cap \mathbb{N}| + |\psi^{-1}(*)|) + (e(H) - |\mathrm{im}(\psi_*) \cap \mathbb{N}| + |\psi_*^{-1}(*)|) . \tag{4.5}$$

Here, the induced edge labelling $\psi_* : E(H) \to \mathbb{Z}$ *is defined as above,* $\psi_*(xy) = |\psi(x) - \psi(y)|$ *with the convention* $|a - *| = | * - a| = | * - * | = *$. *We say that* ψ *is a* k-almost graceful labelling *if its deficiency is at most* k.

We say that the codomain of ψ is (A, C), *if* $A \subset \mathbb{N}$ *and* $C \subset \mathbb{N}$ *and if* $\mathrm{im}(\psi) \subseteq A \cup \{*\}$ *and* $\mathrm{im}(\psi_*) \subseteq C \cup \{0, *\}$.

Theorem 107. *For every* $\Delta \in \mathbb{N}$ *and every* $\varepsilon > 0$ *there exists an* $\alpha > 0$ *and an* $n_0 \geq 1$ *such that the following holds for every* $n > n_0$ *(with further constants defined as in Setup 103). Let* $(A^+, C^+) \subseteq (\mathbb{A}^+, \mathbb{C}^+)$ *such that* $|A|, |C| \geq (1 + \frac{\varepsilon}{2})n$, *and* (A^+, C^+) *forms an* (α, ℓ, Δ)*-quasirandom pair.*

Suppose that T *is an* n*-vertex tree with* $\Delta(T) \leq \Delta$. *Then there exists an* $(\varepsilon' n)$*-almost graceful labelling* $\psi : V(T) \to \mathbb{N}$ *with codomain* (A, C).

The proof of Theorem 107 is in Section 4.3.

We reduce Theorem 38 to Theorem 107 in Section 4.2.3. This amounts to repairing the *conflicts* from the labelling produced by Theorem 107, i.e. pairs of vertices (or edges) that receive the same label from A (or from C, respectively). To this end, we need the following definition.

Definition 108 (solution, repair pair)**.** *Suppose that* $(A^+, C^+) \subseteq (\mathbb{A}^+, \mathbb{C}^+)$. *We say that a number* $a \in \mathbb{Z}$ *is a* solution to an s-tuple $(b_1, b_2, \ldots, b_s)$ *in* (A^+, C^+) *if* $a \in A$, *and for the set* $D = \{|b_1 - a|, |b_2 - a|, \ldots, |b_s - a|\}$ *we have that* $D \subseteq C$, *and* $|D| = s$.

We say that a set $(A^+_{\mathrm{rep}}, C^+_{\mathrm{rep}})$ *is a* (k, Δ)-repair pair *if for each* $s \in [\Delta]$ *we have that each* s*-tuple of distinct numbers* $b_1, \ldots, b_s \in \mathbb{A}$ *has at least* k *solutions in* $(A^+_{\mathrm{rep}}, C^+_{\mathrm{rep}})$.

Similarly, $(A^+_{\mathrm{rep}}, C^+_{\mathrm{rep}})$ *is a* (k, d, Δ)-repair pair *if for each* $s \in [\Delta]$, *each* s*-tuple of distinct numbers from* $\mathbb{A}$ *has at least* kd^{s+1} *solutions in* $(A^+_{\mathrm{rep}}, C^+_{\mathrm{rep}})$.

Lastly, the size *of a repair pair* $(A^+_{\mathrm{rep}}, C^+_{\mathrm{rep}})$ *is* $\max\{|A_{\mathrm{rep}}|, |C_{\mathrm{rep}}|\}$.

The next lemma shows the existence of sparse repair pairs.

Lemma 109. *Let* $\varepsilon > 0$ *be given. For any* $\alpha > 0$ *there exists an* $n_0 \geq 1$, *such that for any* $n \geq n_0$ *and any* $\Delta \leq \frac{\log n}{100 \log(1/\varepsilon)}$ *there exists an* $(m/4, d, \Delta)$*-repair pair* $(A^+_{\mathrm{rep}}, C^+_{\mathrm{rep}})$ *of size at most* $2dm$ *with the property that the pair* $(\mathbb{A}^+ \setminus A^+_{\mathrm{rep}}, \mathbb{C}^+ \setminus C^+_{\mathrm{rep}})$ *is* (α, ℓ, Δ)*-quasirandom.*

The proof of Lemma 109 is given in Section 4.2.4. As Lemma 109 will be used later for a random tree (which has maximum degree roughly $\frac{\log n}{\log \log n}$), we prove it in a stronger version than needed for Theorem 38, for which a version with Δ constant would be sufficient.

The next easy lemma shows how the parameters of a repair pair change after deletion of elements.

Lemma 110. *Suppose that* (B^+, D^+) *is a* (k, Δ)*-repair pair, and* $B'^+ \subseteq B^+$, $D'^+ \subseteq D^+$ *are such that* $|B'^+| \geq |B^+| - r$ *and* $|D'^+| \geq |D^+| - r\Delta$. *Then* (B'^+, D'^+) *is a* $(k - (\Delta^2 + 1)r, \Delta)$*-repair pair.*

Proof. Let $s \in [\Delta]$ and suppose that $(b_1, \dots, b_s)$ is an s-tuple that has t solutions a in (B^+, D^+). For all but at most $|B^+| - |B'^+|$ of these solutions a we have $a \in B'$. Similarly, for each fixed $i = 1, \dots, s$, we have $|b_i - a| \in D'$ for all but at most $|D^+| - |D'^+|$ solutions a. We conclude that $(b_1, \dots, b_s)$ has at least $t - (\Delta^2 + 1)r$ solutions in (B'^+, D'^+). □

4.2.3 Proof of Theorem 38 from Theorem 107 and Lemma 109

Let Δ and ε be given. Recall that Setup 103 defines further constants $m, d, \varepsilon', \varepsilon'', \gamma, R, \ell$ and sets $\mathbb{A}, \mathbb{A}^+, \mathbb{C}, \mathbb{C}^+$.

We first apply Theorem 107 to obtain the parameter α. Then let $(A^+_{\mathrm{rep}}, C^+_{\mathrm{rep}})$ be the output of Lemma 109 for the parameters $\Delta_{\mathrm{L109}} := \Delta$ and $\alpha_{\mathrm{L109}} := \alpha$. Note that $(A^+_{\mathrm{rep}}, C^+_{\mathrm{rep}})$ is an $(md^{\Delta+1}/4, \Delta)$-repair pair.

Let T be any n-vertex tree with $\Delta(T) \leq \Delta$ and $n > n_0$, for n_0 sufficiently large.

We apply Theorem 107 to T and $(\mathbb{A}^+ \setminus A^+_{\mathrm{rep}}, \mathbb{C}^+ \setminus C^+_{\mathrm{rep}})$, in order to obtain an $(\varepsilon' n)$-almost graceful labelling ψ of T with codomain $(\mathbb{A} \setminus A_{\mathrm{rep}}, \mathbb{C} \setminus C_{\mathrm{rep}})$.

We now repair the conflicts. Sequentially, we shall decrease the deficiency of ψ by (at least) 1, while we keep track of an evolving repair pair $(A^*_{\mathrm{rep}}, C^*_{\mathrm{rep}})$, starting with $A^*_{\mathrm{rep}} := A_{\mathrm{rep}}$ and $C^*_{\mathrm{rep}} := C_{\mathrm{rep}}$ which will have the property that the cardinality of A^*_{rep} decreases by at most 1 and the cardinality of C^*_{rep} decreases by at most Δ in each step.

In an r-th step let us consider an arbitrary vertex or an arbitrary edge that contribute to (4.5). Suppose first that we have the case of a vertex $v \in V(T)$. Then $\psi(v) = *$ or $\psi(v) = \psi(v')$ for some vertex $v' \neq v$. Let $b_1, \dots, b_s$ be the distinct labels appearing on vertices adjacent with v. Here, we do not include joker vertices, and labels that appear with higher multiplicities are included only once. Since $(A_{\mathrm{rep}}, C_{\mathrm{rep}})$ is a $(md^{\Delta+1}/4, \Delta)$-repair pair, Lemma 110 tells us that there are at least

$$\frac{md^{\Delta+1}}{4} - (\Delta^2 + 1)r$$

solutions to $b_1, \dots, b_s$ in $(A^*_{\mathrm{rep}}, C^*_{\mathrm{rep}})$. Since $r \leq \varepsilon' n$ and $\varepsilon' \leq d^{\Delta+1}/(16\Delta^2)$, there is at least one solution.

Let us take any such solution a, and let us reset $\psi(v) := a$. We now update the set A^*_{rep} by removing a, and the set C^*_{rep} by removing $|a - b_1|, \dots, |a - b_s|$. It may now happen that an edge incident with v and with label $*$ has both endpoints labelled with values different from $*$; in this case we relabel the edge with the absolute difference of these values (the label will be of the form $|a - b_i|$ by assumption). Observe that the deficiency of the modified labelling ψ went down by at least 1 as promised.

Secondly, let us consider that we have an edge $uv \in E(T)$ that contributes to (4.5). That is, either there exists an edge $e \neq uv$ such that $\psi_*(e) = \psi_*(uv)$, or $\psi_*(uv) = *$. By considering all the joker vertices first when repairing conflicts on the vertices we may assume without lack of generality that both u and v have labels different from $*$. Then we relabel an arbitrary endvertex of uv, say u, using the same procedure as above. This ensures that uv receives a new label from C^*_{rep}. Again, the deficiency goes down by at least 1, and the sets A^*_{rep} and C^*_{rep} decrease in cardinality by at most 1 and Δ, respectively.

Continuing in this manner for at most $\varepsilon' n$ steps, we repair all conflicts.

4.2.4 Proof of Lemma 109

Let $\varepsilon < \frac{1}{2}$ and α be given. We will suppose n_0 is sufficiently large and let $\Delta \leq \frac{\log n}{100\log(1/\varepsilon)}$ be arbitrary.

For $a \in \mathbb{A}$ and an s-tuple $\mathbf{b} = (b_1, \ldots, b_s) \in \mathbb{A}^s$ we define the set $X_a := \{|b_1 - a|, \ldots, |b_s - a|\}$ (here and in what follows tuples always consist of pairwise distinct values). Observe that for each $\mathbf{b}$, there are at least $|\mathbb{A}| - 2\ell s - \binom{s}{2}$ solutions in $(\mathbb{A}^+, \mathbb{C}^+)$.

A formalism that is useful for later is to regard sets $B \subseteq \mathbb{A}^+$, $D \subseteq \mathbb{C}^+$ as a single subset $B \sqcup D$ of the disjoint union $\mathbb{A}^+ \sqcup \mathbb{C}^+$.

Let

$$\mathcal{S}_{\mathbf{b}} := \big\{\{a\} \sqcup X_a \,:\, a \text{ is a solution for } b \text{ in } (\mathbb{A}^+, \mathbb{C}^+)\big\} .$$

Consider now a d-random subset $A^+ \sqcup C^+$ of $\mathbb{A}^+ \sqcup \mathbb{C}^+$. That is, each element of the first component of $\mathbb{A}^+ \sqcup \mathbb{C}^+$ is included with probability d into A^+, and each element of the second component of $\mathbb{A}^+ \sqcup \mathbb{C}^+$ is included with probability d into C^+. To prove the lemma, it suffices to show that (A^+, C^+) constructed this way satisfies the demanded properties (which are: size, quasirandomness, and the possibility to repair) with positive probability. In particular, by the union bound, it suffices to prove

$$\mathbf{P}\big[|A| \geq 2dm \text{ or } |C| \geq 2dm\big] < \frac{1}{3} , \tag{4.6}$$

$$\mathbf{P}\big[(\mathbb{A}^+ \setminus A^+, \mathbb{C}^+ \setminus C^+) \text{ is not } (\alpha, \ell, \Delta)\text{-quasirandom}\big] < \frac{1}{3} , \tag{4.7}$$

and that for each $s = 1, \ldots, \Delta$ and each $\mathbf{b} \in \mathbb{A}^s$ (there are at most $2|\mathbb{A}|^\Delta$ choices of such s and $\mathbf{b}$) we have

$$\mathbf{P}\left[\big|\big\{\{a\} \sqcup X_a \in \mathcal{S}_{\mathbf{b}} : \{a\} \sqcup X_a \subseteq A \sqcup C\big\}\big| \leq \tfrac{1}{4} m d^{s+1}\right] < \frac{1}{6|\mathbb{A}|^\Delta} . \tag{4.8}$$

Property (4.6) follows trivially by Chernoff's bound. In Lemmas 112 and 111 we prove properties (4.7) and (4.8).

Lemma 111. *Let $s \in [\Delta]$ and $\mathbf{b} \in \mathbb{A}^s$ be arbitrary. Then (4.8) holds.*

Proof. Let us write $\mathbf{b} = (b_1, b_2, \ldots, b_s)$. For $a \in \mathbb{A}$, we write $\mathbb{I}_a$ for indicator of the event that $\{a\} \sqcup X_a \subseteq A \sqcup C$. Here the set X_a is taken with respect to $\mathbf{b}$. Furthermore, throughout the proof, 'solution' always refers to $\mathbf{b}$ and with respect to $(\mathbb{A}^+, \mathbb{C}^+)$.

Let

$$\Xi = \sum_{a \text{ solution}} \mathbb{I}_a \quad \text{ and } \quad \Psi = \sum_{\substack{a, a' \text{ solution} \\ X_a \cap X'_a \neq \emptyset}} \mathbb{I}_a \mathbb{I}_{a'} .$$

We have $\mathbf{E}[\Xi] = |\mathcal{S}_{\mathbf{b}}| d^{s+1} \geq \frac{m}{2} d^{s+1}$. Thus, in order to prove (4.8) we need to show that Ξ is rarely substantially smaller than its expectation.

Note that

$$\mathbf{E}[\Psi] = \mathbf{E}[\Xi] + \sum_{\substack{a \neq a' \text{ solutions} \\ X_a \cap X'_a \neq \emptyset}} \mathbf{E}[\mathbb{I}_a \mathbb{I}_{a'}] . \tag{4.9}$$

Suppose now that a solution a is fixed. Note that if $a \neq a'$ then $\mathbf{E}[\mathbb{I}_a \mathbb{I}_{a'}] = \mathbf{E}[\mathbb{I}_a] d^{1+|X_{a'} \setminus X_a|}$ and so

$$\sum_{\substack{a \neq a' \text{ solutions} \\ X_a \cap X'_a \neq \emptyset}} \mathbf{E}[\mathbb{I}_a \mathbb{I}_{a'}] = \sum_{\substack{a \neq a' \text{ solutions} \\ X_a \cap X'_a \neq \emptyset}} \mathbf{E}[\mathbb{I}_a] d^{1+|X_{a'} \setminus X_a|} = \sum_{\substack{a \neq a' \text{ solutions} \\ X_a \cap X'_a \neq \emptyset}} \mathbf{E}[\mathbb{I}_a] d^{1+s-|X_a \cap X_{a'}|}.$$

We will now upper-bound this quantity.

Claim 10. *For any $1 \leq i \leq s$, there are at most $2i^2$ other solutions a' such that $|X_a \cap X_{a'}| = s - i + 1$.*

Proof. Let $c_1 < c_2 < \ldots < c_s$ be the values in X_a and $j \geq s - i + 1$ arbitrary. We show that there are at most $2i$ solutions a' such that $|X_a \cap X_{a'}| = s - i + 1$ and $c_j \in X_a \cap X_{a'}$, but $c_{j+1}, \ldots, c_s \notin X_a \cap X_{a'}$. This clearly implies the claim.

Let a' be any such solution. Then we can find an index $f(a')$ such that $c_j = |a' - b_{f(a')}|$. If $b_{f(a')} < a'$, we call a' a *right match*. Otherwise, we call a' a *left match*.

The important fact is that there are at most i right matches, and at most i left matches. For example, assume for a contradiction that $a_1 < \ldots < a_{i+1}$ are $i+1$ right matches. Then $a_{i+1} - b_{f(a_t)} > a_t - b_{f(a_t)} = c_j$ for all $t \leq i$. As $f(a_1), \ldots, f(a_i)$ are all distinct numbers, we see that $X_{a_{i+1}}$ contains i values $a_{i+1} - b_{f(a_1)}, a_{i+1} - b_{f(a_2)}, \ldots, a_{i+1} - b_{f(a_i)}$, all of them larger than c_j. None of these values is one of $c_{j+1}, \ldots, c_s$, and so $|X_a \cap X_{a'}| \leq s - i$, a contradiction. This proves the assertion. □

Hence

$$\sum_{\substack{a \neq a' \text{ solutions} \\ X_a \cap X'_a \neq \emptyset}} \mathbf{E}[\mathbb{I}_a] d^{1+|X_{a'} \setminus X_a|} \leq \sum_{a \text{ solution}} \sum_{i=1}^{s} 2i^2 \mathbf{E}[\mathbb{I}_a] d^i \leq \sum_{a \text{ solution}} \mathbf{E}[\mathbb{I}_a] \frac{4d + 12d^2}{(1-d)^3} \leq \frac{\mathbf{E}[\Xi]}{2},$$

where in the last inequality we used $d \leq \frac{1}{16}$. Plugging this into (4.9), we get

$$\mathbf{E}[\Psi] \leq \frac{3}{2} \mathbf{E}[\Xi].$$

Janson's Inequality (see [69, Theorem 2.14]) gives that

$$\mathbf{P}[\Xi \leq \tfrac{1}{4} m d^{s+1}] \leq \exp\left\{-\tfrac{d^{s+1} n}{24}\right\} \leq \exp\left\{-\frac{\sqrt{n}}{24}\right\},$$

where the last inequality follows from $d = \frac{\varepsilon}{16}$ and $s \leq \Delta \leq \frac{\log n}{100 \log(1/\varepsilon)}$. As $6|\mathbb{A}|^{\Delta} \leq \exp\{(\log n)^2\} \leq \exp\left\{\frac{\sqrt{n}}{24}\right\}$, (4.8) follows. □

Lemma 112. *The pair $(\mathbb{A}^+ \setminus A^+, \mathbb{C}^+ \setminus C^+)$ is (α, ℓ, Δ)-quasirandom with probability more than $2/3$.*

Proof. Define $\beta := \frac{\alpha}{2^{\Delta+4}}$. We will assume $n_0 := n_0(\alpha, \varepsilon)$ is sufficiently large. As Δ may depend on n, we will need to control more carefully the quantities involved in our estimations. Note for later use that

$$\gamma \geq \frac{\varepsilon^{\Delta+1}}{2^{4\Delta+20}\Delta^3}, \tag{4.10}$$

and $(1 \pm 2\beta)^{\Delta+1} = 1 \pm \frac{\alpha}{4}$.

Let $B^+ := \mathbb{A}^+ \setminus A^+$ and $D^+ := \mathbb{C}^+ \setminus C^+$. Note that $\mathbf{E}[|B|] = (1-d)|\mathbb{A}|$ and so by Chernoff's inequality (Theorem 100),

$$\mathbf{P}\left[\big||B| - \mathbf{E}[|B|]\big| > \beta(1-d)|\mathbb{A}|\right] < 2\exp\left\{-\frac{(1-d)|\mathbb{A}|\beta^2}{4}\right\} \leq 2\exp\left\{-\frac{\beta^2 n}{16}\right\}.$$

As $\Delta \leq \frac{\log n}{100}$, this is upper-bounded by $2\exp\left\{-\alpha^2\sqrt{n}\right\}$. Thus with probability $1 - o(1)$ we have $|B| = (1\pm\beta)(1-d)|\mathbb{A}|$. Similarly, with probability $1-o(1)$ we have $|D| = (1\pm\beta)(1-d)|\mathbb{C}|$. Hence

$$1 - d = \frac{|B|}{(1\pm\beta)|\mathbb{A}|} \quad \text{and} \quad 1 - d = \frac{|D|}{(1\pm\beta)|\mathbb{C}|} \text{ a.a.s.} \tag{4.11}$$

We now check that (B^+, D^+) is (α, ℓ, Δ)-quasirandom with high probability.

First we check condition a) of Definition 105. For an arbitrary $I \in \mathcal{I}$ we have $|I \cap B^+| = (1\pm\beta)(1-d)|I|$ with probability at least $1 - 2\exp\{-\frac{\beta^2\ell}{8}\}$. By (4.10), our definition of β and assuming n_0 is sufficiently large compared to α, this is at least $1 - 2\exp\{-\sqrt{n}\}$. Thus using the union bound and (4.11), we see that with probability $1 - o(1)$ we have for all $I \in \mathcal{I}$ that

$$|I \cap B^+| = (1\pm\beta)|I|\frac{|B|}{(1\pm\beta)|\mathbb{A}|} = (1\pm\alpha)|I|\frac{|B|}{|\mathbb{A}|},$$

as required.

We check condition b). Let $I \in \mathcal{I}, p \in [\Delta]$ and $a_1, \ldots, a_p \in I$ be distinct arbitrary. Consider the random variable $X = |\mathrm{Sol}(B^+, D^+, I; a_1, \ldots, a_p)|$. By construction, $\mathbf{E}[X] = (1-d)^{p+1}\ell$.

Thus $\mathbf{E}[X] \geq \left(\frac{7}{8}\right)^{p+1}\ell$.

Note that X can be thought of as defined on the product probability space $\{0,1\}^{|\mathbb{A}^+|} \times \{0,1\}^{|\mathbb{C}^+|}$, corresponding to our independent decisions of putting labels in A^+ and C^+. This product has $|\mathbb{A}^+| + |\mathbb{C}^+| \leq 4n$ components. According to this definition, X is p-Lipschitz: indeed, changing a label from A^+ to $\mathbb{A}^+ \setminus A^+$ can affect at most one element of $\mathrm{Sol}(B^+, D^+, I; a_1, \ldots, a_p)$, while changing a label from C^+ to $\mathbb{C}^+ \setminus C^+$ can affect at most p of them. Thus by Lemma 101 we get

$$\mathbf{P}\left[|X - \mathbf{E}[X]| > \beta\mathbf{E}[X]\right] \leq 2\exp\left\{-\frac{2(\beta\mathbf{E}[X])^2}{4p^2 n}\right\} \leq 2\exp\left\{-\frac{\beta^2\gamma^2}{2\Delta^2}\left(\frac{7}{8}\right)^{2(p+1)} n\right\}.$$

By (4.10), this is at most $2\exp\{-\sqrt{n}\}$. There are at most m choices for I, at most Δ choices for p and at most $(\gamma n)^\Delta$ choices for $a_1, \ldots, a_p$. Thus with probability $1 - o(1)$, for all such choices we have

$$X = (1\pm\beta)(1-d)^{p+1}\ell \overset{(4.11)}{=} (1\pm\beta)(1\pm\frac{\alpha}{4})\ell\frac{|B|}{|\mathbb{A}|}\left(\frac{|D|}{|\mathbb{C}|}\right)^p = (1\pm\alpha)\ell\frac{|B|}{|\mathbb{A}|}\left(\frac{|D|}{|\mathbb{C}|}\right)^p,$$

as desired.

Finally, we check condition c).

Let $c \in C$ be arbitrary and consider the set

$$T_c := \left\{(a, I) : I \in \mathcal{I}, a \in I, a \oplus c \in \bar{I}\right\}.$$

We will show that $|T_c| = \ell^2 + O(1)$. For this we shall assume without lack of generality that ℓ is odd.

Let $(a, I) \in T_c$ and assume that $I = [x, x + \ell - 1]$ with $x \leq \frac{m}{2} - \ell$. Then $c \in [m - 2(x + \ell - 1), m - 2x]$. Hence $x \in [\frac{m-c}{2} - \ell + 1, \frac{m-c}{2}]$. Note that $\frac{m-c}{2} \leq \frac{m}{2} - \ell$ as $c \geq 2\ell$ by definition of $\mathbb{C}$. Recall that m is odd by definition. Assume without lack of generality that c is odd and let $x = \frac{m-c}{2} - i, 0 \leq i < \ell$. We want to count how many times we can obtain c as a difference of labels from the interval $\bar{I} = [\frac{m+c}{2} - (\ell - 1) + i, \frac{m+c}{2} + i]$ and $I = [\frac{m-c}{2} - i, \frac{m-c}{2} - i + \ell - 1]$. If $a = \frac{m-c}{2} - i + j \in I$ and $a \oplus c \in \bar{I}$ then $2i - (\ell - 1) \leq j \leq 2i$. Furthermore $0 \leq j \leq \ell - 1$, so we get $2i + 1$ choices of a if $i \leq \frac{\ell-1}{2}$, and $2(\ell - 1 - i) + 1$ choices otherwise. Thus we have

$$\left(2 \sum_{i=0}^{(\ell-1)/2} 2i + 1\right) - \ell = \frac{(\ell - 1)(\ell + 1)}{2} + 1$$

choices for (a, I) with $x \leq \frac{m}{2} - \ell$. This gives a total of $\ell^2 + 1$ choices, showing $|T_c| = \ell^2 + O(1)$.

Now let $X_c := \{(a, I) \in T_c : a \in B^+, a \oplus c \in B^+\}$. Clearly $\mathbf{E}[X_c] = (1 - d)^2 |T_c|$. Hence $\mathbf{E}[X_c] \geq \frac{\ell^2}{2}$.

We will estimate X_c using McDiarmid's inequality. We can think of X_c as defined on the product probability space $\{0, 1\}^{|\mathbb{A}^+|}$. Then X_c is 2ℓ-Lipschitz and using $|\mathbb{A}^+| \leq 2n$ we get

$$\mathbf{P}\left[|X_c - \mathbf{E}[X_c]| > \beta \mathbf{E}[X_c]\right] \leq 2 \exp\left\{-\frac{2(\beta \mathbf{E}[X_c])^2}{8\ell^2 n}\right\} \leq 2 \exp\left\{-\frac{\beta^2 \gamma^2}{16} n\right\} \leq 2 \exp\left\{-\sqrt{n}\right\}.$$

Hence for n large enough, for all $c \in C$ we have $X_c = (1 \pm \beta)(1 - d)^2 |T_c| \overset{(4.11)}{=} (1 \pm \alpha) \frac{|B|^2}{|\mathbb{A}|^2} \ell^2$. This verifies the third condition of quasirandomness. $\square$

4.3 Proof of Theorem 107

In this section we prove Theorem 107. We first describe a randomized tree labelling algorithm and explain how this algorithm can be used to provide a labelling as needed for Theorem 107. A key subroutine of the algorithm called `LocalLabelling` is introduced in Section 4.3.2 and analyzed in Section 4.3.3.

4.3.1 A synopsis of the proof

We use the notation from Setup 103. In particular, we are given ε, Δ and n and we have constants $m, d, \varepsilon', \varepsilon'', \gamma, R, \ell$ and sets $\mathbb{A}, \mathbb{A}^+, \mathbb{C}, \mathbb{C}^+$.

Let T be an n-vertex tree with $\Delta(T) \leq \Delta$. We will generate a labelling where the vertices obtain labels from $\mathbb{A}$, and the labels induced on the edges are from $\mathbb{C}$. This is done as follows.

(s1) We first remove from the sets $\mathbb{A}$ and $\mathbb{C}$ *repair sets* A_{rep} and C_{rep}. This was already done in Lemma 109. Initially, set $A = \mathbb{A} \setminus A_{\mathrm{rep}}$ and $C = \mathbb{C} \setminus C_{\mathrm{rep}}$. The sets A and C will evolve in steps, and in each step A and C will record the available vertex and edge labels, respectively. The key property which we will maintain during the evolution is that

$$\text{at each step, } A \text{ and } C \text{ are random-like sets in the sense of Definition 105.} \tag{4.12}$$

(s2) We partition the input tree T into a collection of trees $\mathcal{T}$, where each tree has a bounded number of vertices. The trees will each have order between R and ΔR. We then group the subtrees from $\mathcal{T}$ into R groups of roughly equal size (see Lemma 125). Denote by $T_i^j \in \mathcal{T}$ the j-th tree from group i. We will process the trees from the i-th group in the i-th round.

(s3) In each round $i = 1, \ldots, R$ we have a collection $(A, C) \subseteq (\mathbb{A}, \mathbb{C})$ of available vertex labels and edge labels.

We first extend the set A into A^+ by adding (randomly) elements of $\mathbb{A}^+ \setminus \mathbb{A}$ to A such that A^+ has the same density on $\mathbb{A}^+ \setminus \mathbb{A}$ as on $\mathbb{A}$. That is, each element is added with probability $\frac{|A|}{|\mathbb{A}|}$. In the same way we extend C to C^+. Observe that this is a random extension of sets which we assumed to be random-like by (4.12). Hence A^+ and C^+ are (with high probability) random-like.

Then we label the vertices of all the trees T_i^j with labels in $A \cup \{*\}$, where $*$ denotes a *joker*. That will induce a labelling of the edges of T_i^j with labels in $C \cup \{*\}$. Within each round we will allow conflicts.

To do this, for each tree T_i^j we perform the randomized labelling algorithm `LocalLabelling` described below with available labels (A^+, C^+).

A key feature of this randomized algorithm is that with high probability, assuming that the sets A^+ and C^+ were initially random-like, the labelling of $\mathcal{T}_i^j$ will have only a small number of conflicts, will only use a small number of jokers * and the labels of A^+ and C^+ used will be random-like.

We fix one such typical outcome of `LocalLabelling`. We now update the sets A and C by removing from them the labels which have been used in the current round. Observe that the invariant (4.12) is maintained.

(s4) Since we can control the number of conflicts and jokers used in each round individually we get in total an $(\varepsilon' n)$-graceful labelling of T.

4.3.2 The algorithm `LocalLabelling`

The vertices of each small tree T_i^j will be labelled using the procedure `LocalLabelling`, which we formally define below.

For an arbitrary finite set Ω let uniform(Ω) denote an element of Ω sampled uniformly at random. If $\Omega = \emptyset$ then the call to uniform(Ω) fails. We shall think of uniform(Ω) as choosing a real number $x \in (0, 1]$ uniformly at random and then returning the i-th element of Ω (in some arbitrary but fixed enumeration of Ω) if $i - 1 < x|\Omega| \leq i$. This allows us to encode uniform(Ω) without actually knowing Ω.

The following steps define the algorithm `LocalLabelling`(F, A^+, C^+):

- *Input*: a tree F with a bipartition $V(F) = V_{\text{odd}} \sqcup V_{\text{even}}$.
- *Output*: a vertex labelling $\pi : V(F) \to A \cup \{*\}$ and an edge labelling $\pi_* : E(F) \to C \cup \{*\}$.
- Sample a random interval $I \in \mathcal{I}$ (each with probability $\frac{1}{m+\ell}$, cf. (4.4)).
- For $x \in V_{\text{odd}}$, set $\pi(x) := \text{uniform}(A^+ \cap I)$.
- For $y \in V_{\text{even}}$, let the neighbors of y be $x_1, \ldots, x_s \in V_{\text{odd}}$. Then set

$$\pi(y) := \text{uniform}(\text{Sol}(A^+, C^+, I; \pi(x_1), \ldots, \pi(x_s))).$$

- For $xy \in E(T)$ set $\pi_*(xy) = |\pi(x) - \pi(y)|$.
- Change all vertex-/edge-labels from $A^+ \setminus A$ and $C^+ \setminus C$ to $*$.

4.3.3 Quasirandomness

To prove that the algorithm produces a labelling with the properties described earlier, we will need to show some properties of the sets of available labels A and C after each round of the algorithm. These properties, which will be stated below, will ensure that in each step of the algorithm the probabilities of obtaining a particular vertex label / edge label are roughly the same for each label, and that the expected number of jokers used is small.

Lemma 113. *Let $(A^+, C^+) \subseteq (\mathbb{A}^+, \mathbb{C}^+)$ be (α, ℓ, Δ)-quasirandom and let T be any tree of maximum degree at most Δ with a bipartition $V(T) = V_{odd} \sqcup V_{even}$.*

Suppose $\ell \geq \frac{\Delta^2|\mathbb{A}||\mathbb{C}|}{\alpha|A||C|}$ and $\alpha \leq \frac{1}{2}$. If $\pi : V(T) \to A \cup \{*\}$ and $\pi_* : E(T) \to C \cup \{*\}$ is the labelling provided by `LocalLabelling`*(T, A^+, C^+) then*

(i) for all $x \in V_{odd}$ and $a \in A$,

$$\mathbf{P}[\pi(x) = a] = \frac{m - 2\ell + 1}{m + \ell} \cdot \frac{1}{|A|}(1 \pm 2\alpha).$$

(ii) for all $y \in V_{even}$ and $a \in A$,

$$\mathbf{P}[\pi(y) = a] = \frac{m - 2\ell + 1}{m + \ell} \cdot \frac{1}{|A|}(1 \pm 2\alpha)^{2\Delta+2}.$$

(iii) for all $xy \in E(T)$ and all $c \in C$,

$$\mathbf{P}[\pi_*(xy) = c] = \frac{m - 2\ell + 1}{m + \ell} \cdot \frac{1}{|C|}(1 \pm 2\alpha)^{2\Delta+2}.$$

Proof of (i). A vertex $x \in V_{\text{odd}}$ is assigned a label $a \in A$ if and only if we initially picked an interval I containing a (there are ℓ such intervals), and we further picked a at the uniformly random selection from $I \cap A^+$. By (4.4), the probability of choosing a certain interval I is $\frac{1}{m+\ell}$. Using the quasirandomness condition a) that yields

$$\begin{aligned} \mathbf{P}[\pi(x) = a] &= \sum_{\substack{I \in \mathcal{I} \\ a \in I}} \frac{1}{m + \ell} \cdot \frac{1}{|I \cap A^+|} = \sum_{\substack{I \in \mathcal{I} \\ a \in I}} \frac{1}{m + \ell} \cdot \frac{1}{\ell} \cdot \frac{|\mathbb{A}|}{|A|} \cdot (1 \pm \alpha)^{-1} \\ &= \ell \cdot \frac{1}{m + \ell} \cdot \frac{1}{\ell} \cdot \frac{m - 2\ell + 1}{|A|} \cdot (1 \pm 2\alpha) = \frac{m - 2\ell + 1}{m + \ell} \cdot \frac{1}{|A|} \cdot (1 \pm 2\alpha). \end{aligned}$$

□

Proof of (ii). Denote the neighbors of a given vertex $y \in V_{\text{even}}$ by $x_1, \ldots, x_s \in V_{\text{odd}}$.

Let $I \in \mathcal{I}$ be arbitrary such that $a \in \bar{I}$. Let $\mathcal{B}_I$ be the event that the interval I is chosen for T, $\pi(y) = a$ and all labels $\pi(x_1), \ldots, \pi(x_s)$ are pairwise distinct. Similarly, let $\mathcal{C}_I$ be the event that the interval I is chosen for T, $\pi(y) = a$, but not all labels $\pi(x_1), \ldots, \pi(x_s)$ are distinct.

We first estimate $\mathbf{P}[\mathcal{B}_I]$. We get

$$\begin{aligned} \mathbf{P}[\mathcal{B}_I] &= \sum_{\{a_1, \ldots, a_s\} \subseteq \text{Sol}(\bar{I}; a) \text{ distinct}} \mathbf{P}[\forall_i \pi(x_i) = a_i] \cdot \frac{1}{|\text{Sol}(I; a_1, \ldots, a_s)|} \\ &= \sum_{\{a_1, \ldots, a_s\} \subseteq \text{Sol}(\bar{I}; a) \text{ distinct}} \frac{1}{|I \cap A^+|^s} \cdot \frac{1}{|\text{Sol}(I; a_1, \ldots, a_s)|} \\ &= \binom{|\text{Sol}(\bar{I}; a)|}{s} s! \frac{1}{\left((1 \pm \alpha)\ell \frac{|A|}{|\mathbb{A}|}\right)^{s+1} \left(\frac{|C|}{|\mathbb{C}|}\right)^s}. \end{aligned}$$

But by the quasirandomness condition b), $|\mathrm{Sol}(\bar{I};a)|-s+1 \geq (1-2\alpha)|\bar{I}|\frac{|A|}{|\mathbb{A}|}\frac{|C|}{|\mathbb{C}|}$, as $\ell \geq \frac{s|\mathbb{A}||\mathbb{C}|}{\alpha|A||C|}$. Hence

$$\binom{|\mathrm{Sol}(\bar{I};a)|}{s}s! \geq (|\mathrm{Sol}(\bar{I};a)|-s+1)^s \geq (1-2\alpha)^s\left(\ell\frac{|A|}{|\mathbb{A}|}\frac{|C|}{|\mathbb{C}|}\right)^s.$$

On the other hand,

$$\binom{|\mathrm{Sol}(\bar{I};a)|}{s}s! \leq |\mathrm{Sol}(\bar{I};a)|^s \leq (1+2\alpha)^s\left(\ell\frac{|A|}{|\mathbb{A}|}\frac{|C|}{|\mathbb{C}|}\right)^s,$$

so that

$$\binom{|\mathrm{Sol}(\bar{I};a)|}{s}s! = (1\pm 2\alpha)^s\left(\ell\frac{|A|}{|\mathbb{A}|}\frac{|C|}{|\mathbb{C}|}\right)^s.$$

From this it follows that $\mathbf{P}[\mathcal{B}_I] = (1\pm 2\alpha)^{2s+1}\frac{|\mathbb{A}|}{\ell|A|}$.

To estimate $\mathbf{P}[\mathcal{C}_I]$ we proceed similarly, except that in this case the number of choices of labels for $x_1,\ldots,x_s$ is at most $\binom{s}{2}|\mathrm{Sol}(\bar{I};a)|^{s-1}$. Moreover for any such choice $a_1,\ldots,a_s$ of labels we have $|\mathrm{Sol}(I;a_1,\ldots,a_s)| \geq (1-\alpha)|I|\frac{|A|}{|\mathbb{A}|}\left(\frac{|C|}{|\mathbb{C}|}\right)^{s-1}$. Hence $\mathbf{P}[\mathcal{C}_I] \leq (1+2\alpha)^{2s}\frac{s^2}{2}\frac{|\mathbb{A}|^2}{\ell^2|A|^2}$.

It follows that

$$\begin{aligned}\mathbf{P}[\pi(y)=a] &= \sum_{\substack{I\in\mathcal{I}\\ a\in\bar{I}}}\frac{1}{m+\ell}\cdot(\mathbf{P}[\mathcal{B}_I]+\mathbf{P}[\mathcal{C}_I])\\ &= (1\pm 2\alpha)^{2\Delta+2}\frac{m-2\ell+1}{m+\ell}\frac{1}{|A|},\end{aligned}$$

where we used that $\ell \geq \frac{s^2|\mathbb{A}|}{\alpha|A|}$. □

Proof of (iii). Let $xy \in E(T)$ with $x \in V_{\mathrm{odd}}$ and $y \in V_{\mathrm{even}}$ and suppose that the neighbors of y are $x = x_1,\ldots,x_s$. The event $\pi_*(xy) = c$ can only occur if we choose $\pi(x) \in I$ so that $\pi(x)\oplus c \in A^+\cap\bar{I}$.

So let $I \in \mathcal{I}$ be arbitrary and $a_1 \in I\cap A^+$ such that $a_1\oplus c \in \bar{I}\cap A^+$. As in part (ii), we first condition on the event that I is chosen for T and $\pi(x_1) = a_1$. We then define two events $\mathcal{B}_I$ and $\mathcal{C}_I$: $\mathcal{B}_I$ is the event that $\pi(y) = a_1\oplus c$ and the labels $a_1,\pi(x_2),\ldots,\pi(x_s)$ are pairwise distinct, while $\mathcal{C}_I$ is the event that $\pi(y) = a_1\oplus c$ but not all of $a_1,\pi(x_2),\ldots,\pi(x_s)$ are distinct.

We get

$$\begin{aligned}\mathbf{P}[\mathcal{B}_I] &= \sum_{a_2,\ldots,a_s\in\mathrm{Sol}(\bar{I};a_1\oplus c)\setminus\{a_1\}\text{ distinct}}\mathbf{P}[\pi(x_i)=a_i,\forall i\geq 2]\cdot\frac{1}{|\mathrm{Sol}(I;a_1,\ldots,a_s)|}\\ &= \binom{|\mathrm{Sol}(\bar{I};a_1\oplus c)|-1}{s-1}(s-1)!\frac{1}{|I\cap A^+|^{s-1}}\frac{1}{(1\pm\alpha)\ell\frac{|A|}{|\mathbb{A}|}\left(\frac{|C|}{|\mathbb{C}|}\right)^s}\\ &= (1\pm 2\alpha)^{2s-1}\frac{|\mathbb{A}||\mathbb{C}|}{\ell|A||C|}.\end{aligned}$$

Also $\mathbf{P}[\mathcal{C}_I] \leq (1+2\alpha)^{2s-2}\binom{s}{2}\frac{|\mathbb{A}|^2|\mathbb{C}|}{\ell^2|A|^2|C|}$.
Hence using condition c) of quasirandomness we get

$$\begin{aligned}
\mathbf{P}[\pi_*(xy)=c] &= \sum_{I\in\mathcal{J}}\frac{1}{m+\ell}\cdot\sum_{\substack{a_1\\ a_1\in I\cap A^+\\ a_1\oplus c\in\bar{I}\cap A^+}}\frac{1}{|I\cap A^+|}(\mathbf{P}[\mathcal{B}_I]+\mathbf{P}[\mathcal{C}_I])\\
&= \frac{1}{m+\ell}(1\pm\alpha)\ell^2\frac{|A|^2}{|\mathbb{A}|^2}\frac{1}{(1\pm\alpha)\ell\frac{|A|}{|\mathbb{A}|}}(1\pm 2\alpha)^{2s}\frac{|\mathbb{A}||\mathbb{C}|}{\ell|A||C|}\\
&= (1\pm 2\alpha)^{2s+2}\frac{|\mathbb{C}|}{(m+\ell)|C|}\\
&= (1\pm 2\alpha)^{2\Delta+2}\frac{m-2\ell+1}{m+\ell}\frac{1}{|C|}.
\end{aligned}$$

□

4.3.4 The main lemma

The core of the proof is the following nibbling lemma.

Lemma 114. *Let ε and $\Delta \geq 2$ be given. For every $1 > \beta > 0$ there exist $\alpha > 0$ and $n_0 \geq 1$ such that for any $n \geq n_0$ the following holds (we assume the constants defined in Setup 103).*

Let F be any forest of order $\frac{n}{R} \pm \Delta R$, of maximum degree at most Δ and with all trees having between R and ΔR vertices. Let (A^+, C^+) be any (α, ℓ, Δ)-quasirandom pair with $|A| \geq \varepsilon''|\mathbb{A}|$ and $|C| \geq \varepsilon''|\mathbb{C}|$. Then there exists an $(\varepsilon'' v(F))$-almost graceful labelling $\psi : V(F) \to \mathbb{N} \cup \{*\}$ with $\mathrm{im}(\psi) \subseteq A \cup \{*\}$ and $\mathrm{im}(\psi_*) \subseteq C \cup \{0, *\}$. Moreover there exists a pair (B^+, D^+) with $B = A \setminus \mathrm{im}(\psi)$ and $D = C \setminus \mathrm{im}(\psi_*)$ such that (B^+, D^+) is (β, ℓ, Δ)-quasirandom.

Proof. We define

$$\alpha := \frac{\beta}{2^{5\Delta+15}}.$$

Recall that $\varepsilon'' \leq \frac{\varepsilon}{2}$ and $\ell = \lceil \gamma n \rceil$,

$$\gamma = \frac{\varepsilon''}{960\Delta} \quad \text{and} \quad R \geq \frac{153600\Delta}{\varepsilon''^{\Delta+3}}.$$

We will also assume that n_0 is sufficiently large (compared to $\alpha, \beta, \varepsilon'', \gamma$ and R). In particular, $n_0 \geq \frac{\Delta^2}{\gamma\alpha\varepsilon''^2}$, in order to satisfy the hypothesis of Lemma 113.

We first root each component T in F arbitrarily and then label the vertices of T by applying $\texttt{LocalLabelling}(T, A^+, C^+)$ to the bipartition of T given by $V_{\mathrm{odd}}(T) \sqcup V_{\mathrm{even}}(T)$. That means we first randomly choose an interval $I \in \mathcal{J}$, and then label the primary and secondary vertices of T with labels from I and $\bar{I}$.

This gives a random labelling $\psi : V(F) \to \mathbb{N} \cup \{*\}$. We shall show that with high probability ψ has all the required properties. Note that by definition $\mathrm{im}(\psi) \subseteq A \cup \{*\}$ and $\mathrm{im}(\psi_*) \subseteq C \cup \{0, *\}$.

Fix an enumeration $T_1, \dots, T_k$ of the trees in F. Denote by $I_i \in \mathfrak{I}$ the random interval chosen for T_i.

Define the following two sets of vertices, respectively edges, of F:

$$\begin{aligned} \mathrm{FV} &= \{x \in V(F) : \exists y \in V(F) \setminus \{x\} \text{ with } \psi(x) = \psi(y)\}, \quad \text{ and} \\ \mathrm{FE} &= \{xy \in E(F) : \exists x'y' \in E(F) \setminus \{xy\} \text{ with } \psi_*(xy) = \psi_*(x'y')\}. \end{aligned}$$

Thus $|\mathrm{FV}|$ counts the number of vertices in F that repeat a label, and hence $v(F) - |\mathrm{im}(\psi) \cap \mathbb{N}| + |\psi^{-1}(*)| \leq |\mathrm{FV}| + 1$ (we add 1 as the label $*$ may be used only once). Similarly $e(F) - |\mathrm{im}(\psi_*) \cap \mathbb{N}| + |\psi_*^{-1}(*)| \leq |\mathrm{FE}| + 1$. Hence $|FV|$ and $|FE|$ allow us to control (4.5).

Lemma 115. *For any $x \in V(F)$ we have $\mathbf{P}[x \in FV] \leq \frac{\varepsilon''}{40}$.*

Proof. Assume without lack of generality that $x \in V(T_1)$. We shall distinguish two cases, according to the fact that x is a primary vertex or not.

First assume that $x \in V_{\mathrm{odd}}(T_1)$. Let $I \in \mathfrak{I}$ arbitrary. Conditioning on the choice of the interval I for T_1, we expose ψ on the set of vertices $V' := V(F) \setminus (\{x\} \cup N(x))$. There are at most $|V'| \leq v(F)$ distinct labels in $\mathrm{im}(\psi|_{V'}) \cap I \cap \mathbb{A}$, and furthermore by definition, after we fix the label of x, all labels of $N(x)$ are chosen from $\bar{I}$, which is disjoint from I. Hence for $I \subseteq \mathbb{A}$ we have

$$\mathbf{P}[x \in FV | I_1 = I] \leq \frac{v(F)}{|I \cap A|}.$$

Note that $\mathbf{P}[I \not\subseteq \mathbb{A}] = \frac{4(\ell-1)}{m+\ell}$ (this in particular bounds the probability of x receiving label $*$). Consequently,

$$\begin{aligned} \mathbf{P}[x \in FV] &\leq \frac{4(\ell-1)}{m+\ell} + \sum_{I \in \mathfrak{I}, I \subseteq \mathbb{A}} \mathbf{P}[I_1 = I] \frac{v(F)}{|I \cap A|} \\ &\leq \frac{4\ell}{m+\ell} + \frac{m - 3\ell + 4}{m+\ell} \frac{v(F)}{(1 \pm \alpha)\ell \frac{|A|}{|\mathbb{A}|}}, \text{ by quasirandomness condition a)}, \\ &\leq \frac{\varepsilon''}{80} + \frac{2}{R\varepsilon''\gamma} \leq \frac{\varepsilon''}{40}. \end{aligned}$$

Now consider the case when $x \in V_{\mathrm{even}}(T_1)$. Let $I \in \mathfrak{I}$ with $I \subseteq \mathbb{A}$ arbitrary. Conditioning on the choice of the interval I for T_1, we expose ψ on the set of vertices $V' := V(F) \setminus \{x\}$. Let $x_1, \dots, x_p$ be the neighbors of x in T_1 and consider the set $S := \mathrm{Sol}(I; \psi(x_1), \dots, \psi(x_p))$. Note that $S \subseteq \bar{I} \subseteq \mathbb{A}$. There are at most $|V'| \leq v(F)$ distinct labels in $\mathrm{im}(\psi|_{V'}) \cap S$. Hence proceeding as before we have

$$\begin{aligned} \mathbf{P}[x \in FV] &\leq \frac{4(\ell-1)}{m+\ell} + \frac{m - 3\ell + 4}{m+\ell} \frac{v(F)}{(1 \pm \alpha)\ell \frac{|A|}{|\mathbb{A}|} \left(\frac{|C|}{|\mathbb{C}|}\right)^{\Delta}} \\ &\leq \frac{\varepsilon''}{80} + \frac{2}{R\gamma\varepsilon''^{\Delta+1}} \leq \frac{\varepsilon''}{40}. \end{aligned}$$

This proves the lemma. $\square$

Lemma 116. *Let $x, y \in V_{odd}(T_i), i \in [k]$, be two distinct vertices. Then* $\mathbf{P}[\psi(x) = \psi(y)] \leq \frac{4\ell}{m+\ell} + \frac{2}{\varepsilon''\ell}$.

Proof. Let $I \in \mathcal{I}$ arbitrary. If $I \subseteq \mathbb{A}$, then

$$\mathbf{P}[\psi(x) = \psi(y)|I_i = I] = \frac{1}{|I \cap A|} = \frac{1}{(1 \pm \alpha)\ell\frac{|A|}{|\mathbb{A}|}} \leq \frac{2}{\varepsilon''\ell}.$$

As $\mathbf{P}[I \not\subseteq \mathbb{A}] = \frac{4(\ell-1)}{m+\ell}$, the claim follows. □

Lemma 117. *For any $xy \in E(F)$ we have* $\mathbf{P}[xy \in FE] \leq \frac{\varepsilon''}{40}$.

Proof. We may assume without lack of generality that $x \in V_{\text{odd}}(T_1)$ and $y \in V_{\text{even}}(T_1)$. We expose ψ on the set of vertices $V' := V(F)\setminus\{y\}$. If x and some other neighbour u of y receive the same label then neccessarily xy will belong to FE. However by Lemma 116,

$$\mathbf{P}[\psi(x) = \psi(u) \text{ for some } u \in N(y) \setminus \{x\}] \leq \frac{4\ell\Delta}{m+\ell} + \frac{2\Delta}{\varepsilon''\ell}.$$

Also $\mathbf{P}[I_1 \not\subseteq \mathbb{A}] \leq \frac{4\ell}{m+\ell}$.

Let $\mathcal{B}$ be the event that $I_1 \subseteq \mathbb{A}$ and furthermore that all other neighbours of y receive different labels from x. We see that

$$\mathbf{P}[xy \in \text{FE}|\mathcal{B}] \leq \frac{v(F)}{(1 \pm \alpha)\ell\frac{|A|}{|\mathbb{A}|}\left(\frac{|C|}{|\mathbb{C}|}\right)^{\Delta}} \leq \frac{2}{R\gamma\varepsilon''^{\Delta+1}}.$$

Putting all these inequalities together we obtain

$$\mathbf{P}[xy \in \text{FE}] \leq \frac{8\ell\Delta}{m+\ell} + \frac{2\Delta}{\varepsilon''\gamma n} + \frac{2}{R\gamma\varepsilon''^{\Delta+1}} \leq \frac{\varepsilon''}{40},$$

as desired. □

Lemma 115 gives $\mathbf{E}[|\text{FV}|] \leq \frac{\varepsilon'' v(F)}{40}$ and hence by Markov's inequality, $\mathbf{P}[|\text{FV}| \geq \frac{\varepsilon'' v(F)}{4}] \leq 0.1$. Similarly, by Lemma 117, $\mathbf{E}[|\text{FE}|] \leq \frac{\varepsilon'' v(F)}{40}$ and hence by Markov's inequality, $\mathbf{P}[|\text{FE}| \geq \frac{\varepsilon'' v(F)}{4}] \leq 0.1$. Thus with probability at least 0.8 we have $|\text{FV}| + |\text{FE}| + 2 \leq \varepsilon'' v(F)$, that is, ψ is an $(\varepsilon'' v(F))$-almost graceful labelling of F.

Before we go further we first define:

$$M_A := \left(1 - \frac{m - 2\ell + 1}{m+\ell}\frac{1}{|A|}\right)^{v(F)}, \quad \text{and}$$
$$M_C := \left(1 - \frac{m - 2\ell + 1}{m+\ell}\frac{1}{|C|}\right)^{e(F)}.$$

We extend $B := A \setminus \text{im}(\psi)$ to B^+ by adding each element $a \in A^+ \setminus A$ to B^+ with probability M_A independently. Similarly we extend $D := C \setminus \text{im}(\psi_*)$ to D^+ by adding each element $c \in C^+ \setminus C$ to D^+ with probability M_C independently. We will now show that with high probability (B^+, D^+) is a (β, ℓ, Δ)-quasirandom pair.

The proof will be very similar to that in Lemma 112. However, the choices of labels for vertices are no longer independent. Nevertheless, dependencies are confined to vertices or edges belonging to the same tree of the forest. This will result in a small total number of dependencies, allowing us to use Suen's inequality, with almost the same effect as if all choices were pairwise independent. This is done in the next two lemmas.

The first of these lemmas allows us to bound the quantities $\mathbf{P}[\mathcal{B}_i \wedge \mathcal{B}_j]$ in (4.2).

Lemma 118. *Let $T \in F$ be any tree and z_1, z_2 arbitrary vertices or edges of T. For $i = 1, 2$, if z_i is a vertex, let $b_i \in A$ be an arbitrary label, while if z_i is an edge, let $b_i \in C$ be arbitrary. Let $\mathcal{B}_{z_i}$ be the event that z_i receives label b_i. If $z_1 \neq z_2$ or $b_1 \neq b_2$, then*

$$\mathbf{P}[\mathcal{B}_{z_1} \wedge \mathcal{B}_{z_2}] \leq \frac{(1+2\alpha)^{4\Delta+2}}{\gamma \varepsilon''^5 n^2}. \tag{4.13}$$

Proof. We may assume that $z_1 \neq z_2$, otherwise the two events are disjoint and $\mathbf{P}[\mathcal{B}_{z_1} \wedge \mathcal{B}_{z_2}] = 0$.

We first construct two sets of vertices X and Y. For each $i \in \{1, 2\}$ we do the following. If $z_i \in V_{\text{odd}}(T)$ then we add z_i to X. If $z_i \in V_{\text{even}}(T)$ then we add z_i to Y, and we add all the neighbors of z_i to X. Finally, if $z_i = x_i y_i$ with $y_i \in V_{\text{even}}(T)$ then we add y_i to Y, and all the neighbors of y_i to X. Thus $|X| \leq 2\Delta$ and $|Y| \leq 2$.

If $\mathcal{B}_{z_1} \wedge \mathcal{B}_{z_2}$ holds, then a set $X_0 \subset X$ of at most 2 vertices have prescribed labels. For example, if z_1 and z_2 are both vertices in X then $X_0 := \{z_1, z_2\}$. However, it may also happen that just $z_1 \in X$ (in which case $X_0 := \{z_1\}$), or $z_2 \in Y$ and z_1 is an edge incident with z_2 (in which case X_0 contains the other endpoint of z_1). Let A_1 be the set of prescribed labels for the vertices in X_0. Similarly, a set $Y_0 \subset Y$ of at most 2 vertices have prescribed labels A_2. Finally, a set E_0 of at most 2 edges between X and Y have prescribed labels C'.

As an example, if $\mathcal{B}_{z_1}$ is the event $\psi(x) = a, x \in V_{\text{odd}}(T)$, and $\mathcal{B}_{z_2}$ is the event $\psi_*(xy) = c$, then the label of y must be $a \oplus c$, and so $X_0 = \{x\}, Y_0 = \{y\}$ and $E_0 = \{xy\}$.

Let Z be the set of vertices $x \in X$ which are incident with an edge xy in E_0, with the property that $y \in Y \setminus Y_0$. Note that $Z \cap X_0 = \emptyset$ and $|Z| \leq 2$.

We shall now bound $\mathbf{P}[\mathcal{B}_{z_1} \wedge \mathcal{B}_{z_2}]$, by first exposing the vertices in Z and then those in $X \setminus (X_0 \cup Z)$. It should be clear that exposing the vertices in Z will uniquely determine the label of any vertex in Y.

To start, let $I \in \mathcal{I}$ be such that $A_1 \subseteq I, A_2 \subseteq \bar{I}$ and $C' \subseteq |I - \bar{I}|$ (that is, any label in C' can be realized as the absolute difference of a label in I and a label in $\bar{I}$). At least one of A_1, A_2 or C' is non-empty, and this restricts the number of choices of I to at most ℓ.

We now choose labels for all vertices in Z. This can be done in at most $|I \cap A^+|^{|Z|}$ ways, and the probability that these are the labels chosen by `LocalLabelling` is $\left(\frac{1}{|I \cap A^+|}\right)^{|Z|}$.

The labels of the vertices in Y are now uniquely defined. It may happen that these labels are not in $\bar{I}$ or they can not be chosen anymore as stated; in this case nothing else should be done. Otherwise, we choose labels for the vertices in $X \setminus (X_0 \cup Z)$. If $x \in X \setminus (X_0 \cup Z)$ is adjacent to $y \in Y$ (such an y must exist, otherwise $x \in X_0$), there are at most $|\mathrm{Sol}(\bar{I}; \psi(y))| =$

$(1 \pm \alpha)|I|\frac{|A|}{|\mathbb{A}|}\frac{|C|}{|\mathbb{C}|}$ choices for a label. Each choice comes at a probability of $\frac{1}{|I \cap A^+|}$, so that we get a contribution of

$$\left((1 \pm \alpha)|I|\frac{|A|}{|\mathbb{A}|}\frac{|C|}{|\mathbb{C}|}\frac{1}{|I \cap A^+|}\right)^{|X|-|X_0|-|Z|} = \left((1 \pm 2\alpha)^2\frac{|C|}{|\mathbb{C}|}\right)^{|X|-|X_0|-|Z|}.$$

The labels of all vertices in X_0 are fixed from the beginning. However, the probability that they are indeed chosen as such during `LocalLabelling` is $\left(\frac{1}{|I \cap A^+|}\right)^{|X_0|}$.

Finally, the probability that a vertex $y \in Y$ with neighbors $u_1, \ldots, u_p$ receives its prescribed label during `LocalLabelling` is

$$\frac{1}{|\mathrm{Sol}(I; u_1, \ldots, u_p)|} \leq \frac{1}{(1 \pm \alpha)|I|\frac{|A|}{|\mathbb{A}|}\left(\frac{|C|}{|\mathbb{C}|}\right)^p},$$

with the maximum attained only if all neighbors of y receive distinct labels.

Putting all these facts together we get

$$\mathbf{P}[\mathcal{B}_{z_1} \wedge \mathcal{B}_{z_2}] \leq \frac{\ell}{m+\ell}\left((1 \pm 2\alpha)^2\frac{|C|}{|\mathbb{C}|}\right)^{|X|-|X_0|-|Z|}\left(\frac{1}{(1 \pm \alpha)\ell\frac{|A|}{|\mathbb{A}|}}\right)^{|X_0|}\frac{1}{\left((1 \pm \alpha)\ell\frac{|A|}{|\mathbb{A}|}\right)^{|Y|}\left(\frac{|C|}{|\mathbb{C}|}\right)^{e(X,Y)}}.$$

Now note that $e(X, Y) \leq |X| + 1$, with equality if Y contains two vertices which share a neighbor. Furthermore $|X_0| + |Y| = 2$ and $|X_0| + |Z| \leq 2$. Using the fact that $\frac{|A|}{|\mathbb{A}|} \geq \varepsilon''$ and $\frac{|C|}{|\mathbb{C}|} \geq \varepsilon''$, we get (4.13), as desired. □

Lemma 119. *Let $p, q \in [\Delta]$ arbitrary. Let $a_1, \ldots, a_p \in A^+$ be pairwise distinct and similarly let $c_1, \ldots, c_q \in C^+$ be pairwise distinct arbitrary labels. Then*

$$\mathbf{P}[\forall_i a_i \in B^+ \wedge \forall_i c_i \in D^+] = (1 \pm \alpha)M_A^{p(1\pm 2\alpha)^{2\Delta+2}}M_C^{q(1\pm 2\alpha)^{2\Delta+2}}. \tag{4.14}$$

Proof. Below, we prove (4.14) only in the case when $a_1, \ldots, a_p \in A$, and $c_1, \ldots, c_q \in C$. The general case reduces to this. Indeed, introducing an additional label $a_i \in A^+$ or $c_i \in C^+$ changes the probability on the left-hand side of (4.14) exactly by M_A or M_C, respectively.

We now take p disjoint copies $V_1, \ldots, V_p$ of $V(F)$ and q disjoint copies $E_1, \ldots, E_q$ of $E(F)$. Thus we can denote an element of V_i by a pair (x, i) with $x \in V(F)$, and an element of E_i by a triple (x, y, i) with $xy \in E(F)$.

Set $V := \bigcup_{i=1}^p V_i \cup \bigcup_{i=1}^q E_i$.

Define for any $z \in V$ an event $\mathcal{B}_z$ as follows. If $z = (x, i) \in V_i$ then $\mathcal{B}_z$ is the event $\psi(x) = a_i$. If $z = (x, y, i) \in E_i$ then $\mathcal{B}_z$ is the event $\psi_*(xy) = c_i$. This gives a collection of events $\mathcal{C} = \{\mathcal{B}_z\}_{z \in V}$. We further define a graph H on V as follows. We join two vertices z_1 and z_2 of H if they represent vertices or edges of the same component of F. Then H is clearly a superdependency graph for $\mathcal{C}$.

Write $M = \prod_{z \in V} \mathbf{P}\left[\overline{\mathcal{B}}_z\right]$ and note that for n large enough, by Lemma 113,

$$\begin{aligned}
M &= \left(1 - \frac{m-2\ell+1}{m+\ell}\frac{(1\pm 2\alpha)^{2\Delta+2}}{|A|}\right)^{pv(F)} \left(1 - \frac{m-2\ell+1}{m+\ell}\frac{(1\pm 2\alpha)^{2\Delta+2}}{|C|}\right)^{qe(F)} \\
&= (1 \pm \frac{\alpha}{16}) \exp\left\{-\frac{m-2\ell+1}{m+\ell}\frac{(1\pm 2\alpha)^{2\Delta+2}}{|A|} pv(F) - \frac{m-2\ell+1}{m+\ell}\frac{(1\pm 2\alpha)^{2\Delta+2}}{|C|} qe(F)\right\} \\
&= (1 \pm \frac{\alpha}{4}) \left(1 - \frac{m-2\ell+1}{m+\ell}\frac{1}{|A|}\right)^{(1\pm 2\alpha)^{2\Delta+2} pv(F)} \left(1 - \frac{m-2\ell+1}{m+\ell}\frac{1}{|C|}\right)^{(1\pm 2\alpha)^{2\Delta+2} qe(F)} \\
&= (1 \pm \frac{\alpha}{4}) M_A^{p(1\pm 2\alpha)^{2\Delta+2}} M_C^{q(1\pm 2\alpha)^{2\Delta+2}}.
\end{aligned}$$

We wish to use Suen's inequality to approximate $\mathbf{P}[\forall_i a_i \in B^+ \wedge \forall_i c_i \in D^+] = \mathbf{P}[\bigwedge \overline{\mathcal{B}}_z]$ by M. To this end we define for any two adjacent vertices z_1 and z_2 in H the quantity

$$\nu_{z_1,z_2} = \frac{\mathbf{P}[\mathcal{B}_{z_1} \wedge \mathcal{B}_{z_2}] + \mathbf{P}[\mathcal{B}_{z_1}][\mathcal{B}_{z_2}]}{\prod_{\substack{z \sim z_1 \\ \text{or } z \sim z_2}} (1 - \mathbf{P}[\mathcal{B}_z])}.$$

Then z_1 and z_2 belong to the same component T of F. Note that $z \sim z_1$ if and only if $z \sim z_2$. Moreover by Lemma 113, for n large enough,

$$\begin{aligned}
\prod_{\substack{z \sim z_1 \\ \text{or } z \sim z_2}} (1 - \mathbf{P}[\mathcal{B}_z]) &= \prod_{j=1}^{p} \prod_{x \in V(T)} (1 - \mathbf{P}[\mathcal{B}_{(x,j)}]) \prod_{j=1}^{q} \prod_{xy \in E(T)} (1 - \mathbf{P}[\mathcal{B}_{(x,y,j)}]) \\
&= \left(1 - \frac{m-2\ell+1}{m+\ell}\frac{(1\pm 2\alpha)^{2\Delta+2}}{|A|}\right)^{pv(T)} \left(1 - \frac{m-2\ell+1}{m+\ell}\frac{(1\pm 2\alpha)^{2\Delta+2}}{|C|}\right)^{qe(T)} \\
&= (1 \pm \frac{\alpha}{2}) \exp\left\{-\frac{m-2\ell+1}{m+\ell}\frac{(1\pm 2\alpha)^{2\Delta+2}}{|A|} pv(T) - \frac{m-2\ell+1}{m+\ell}\frac{(1\pm 2\alpha)^{2\Delta+2}}{|C|} qe(T)\right\} \\
&= 1 \pm \alpha,
\end{aligned}$$

as $|A| \geq \varepsilon'' n$ and $|C| \geq \varepsilon'' n$, while $e(T) < v(T) \leq \Delta R$.

Furthermore by Lemma 118,

$$\mathbf{P}[\mathcal{B}_{z_1} \wedge \mathcal{B}_{z_2}] \leq \frac{(1+2\alpha)^{4\Delta+2}}{\gamma \varepsilon''^5 n^2}.$$

Also, by Lemma 113,

$$\mathbf{P}[\mathcal{B}_{z_1}]\mathbf{P}[\mathcal{B}_{z_2}] \leq \frac{(1+2\alpha)^{4\Delta+4}}{\varepsilon''^2 n^2}.$$

Consequently

$$\nu_{z_1,z_2} \leq \frac{2(1+2\alpha)^{4\Delta+4}}{(1-\alpha)\gamma \varepsilon''^5 n^2} \leq \frac{2(1+2\alpha)^{4\Delta+5}}{\gamma \varepsilon''^5 n^2}.$$

Hence

$$\begin{aligned}
\sum_{z_1 \sim z_2} \nu_{z_1,z_2} &\leq \sum_{i=1}^{k} \binom{pv(T_i)+qe(T_i)}{2} \frac{2(1+2\alpha)^{4\Delta+5}}{\gamma \varepsilon''^5 n^2} \\
&\leq \frac{n}{R^2}\left(1+\frac{\Delta R^2}{n}\right) \frac{(p+q)^2\Delta^2 R^2}{2} \frac{2(1+2\alpha)^{4\Delta+5}}{\gamma \varepsilon''^5 n^2}, \text{ as } k \leq \frac{n}{R^2}\left(1+\frac{\Delta R^2}{n}\right), \\
&\leq \frac{8\Delta^4(1+2\alpha)^{4\Delta+5}}{\gamma \varepsilon''^5 n}, \text{ as } p+q \leq 2\Delta.
\end{aligned}$$

Thus for n sufficiently large we have $\exp\{\sum_{z_1\sim z_2} \nu_{z_1,z_2}\} \leq 1+\frac{\alpha}{4}$, and so by Suen's inequality applied to H and $\mathcal{C}$ we get (4.14), as desired. □

Lemma 119 has several consequences, which we list below.

Corollary 120. *(i) For any $a \in A^+$ we have $\mathbf{P}[a \in B^+] = (1\pm\alpha)M_A^{(1\pm 2\alpha)^{2\Delta+2}}$.*

(ii) For any $c \in C^+$ we have $\mathbf{P}[c \in D^+] = (1\pm\alpha)M_C^{(1\pm 2\alpha)^{2\Delta+2}}$.

(iii) Let $I \in \mathcal{I}$ and $p \in [\Delta]$ be arbitrary. Let $a_1, \dots, a_p \in I \cap A^+$ be pairwise distinct arbitrary labels. Then for any $a \in \mathrm{Sol}(A^+, C^+, I; a_1, \dots, a_p)$ we have $\mathbf{P}[a \in B^+ \wedge \forall_i |a-a_i| \in D^+] = (1\pm\alpha)M_A^{(1\pm 2\alpha)^{2\Delta+2}} M_C^{p(1\pm 2\alpha)^{2\Delta+2}}$.

(iv) Let $a_1, a_2 \in A^+$ be arbitrary distinct labels. Then $\mathbf{P}[a_1, a_2 \in B^+] = (1\pm\alpha)M_A^{2(1\pm 2\alpha)^{2\Delta+2}}$.

We now record several inequalities for later use.

Lemma 121 (Lemma 39, [12])**.** *For any $0 < M \leq 1$ and any $\delta \in [-\frac{1}{2}, \frac{1}{2}]$ we have $M^{1\pm\delta} = M \pm \delta$.*

Set $\delta := |(1\pm 2\alpha)^{2\Delta+2} - 1|$. Note that

$$\delta = |\sum_{i=1}^{2\Delta+2} \binom{2\Delta+2}{i} (\pm 2\alpha)^i| \leq 2^{2\Delta+2}(2\alpha) = \alpha 2^{2\Delta+3} < \frac{1}{8}.$$

For n large enough,

$$M_A = \left(1 - \frac{m-2\ell+1}{m+\ell}\frac{1}{|A|}\right)^{v(F)} \geq \frac{3}{4}\exp\left\{-\frac{m-2\ell+1}{m+\ell}\frac{v(F)}{|A|}\right\} \geq \frac{3}{4}\exp\left\{-\frac{2}{\varepsilon'' R}\right\} \geq \frac{5}{8},$$

as $\frac{\varepsilon'' R}{2} \geq 6$ and $e \leq (6/5)^6$. In particular, $M_A - \delta \geq \frac{1}{2}$. In a similar manner one deduces that $M_C - \delta \geq \frac{1}{2}$.

Furthermore, for n large enough,

$$|B| \geq |A| - v(F) \geq \frac{|A|}{2}, \quad \text{and} \tag{4.15}$$

$$|D| \geq |C| - e(F) \geq \frac{|C|}{2}. \tag{4.16}$$

Lemma 122. *With probability at least* 0.9 *for every interval* $J \in \mathfrak{I}$ *we have* $|J \cap B^+| = (1 \pm \alpha)\mathbf{E}[|J \cap B^+|]$.

Proof. Let $J \in \mathfrak{I}$ arbitrary. By Corollary 120, (i), we have that

$$\begin{aligned}\mathbf{E}[|J \cap B^+|] &= (1 \pm \alpha) M_A^{(1\pm 2\alpha)^{2\Delta+2}} |J \cap A^+| \\ &\geq (1-\alpha)(M_A - \delta)(1-\alpha)|J|\frac{|A|}{|\mathbb{A}|}, \text{ by Lemma 121,} \\ &\geq \frac{\varepsilon'' \gamma n}{8}.\end{aligned}$$

We will now show that $|J \cap B^+|$ is strongly concentrated around its expectation. By the definition of `LocalLabelling` and the function uniform(), we can think of it as defined on the product probability space $[0,1]^{v(F)} \times \{0,1\}^{|A^+ \setminus A|}$. Here the first component corresponds to the choices made in the algorithm `LocalLabelling`, while the second component corresponds to the independent decisions of adding labels from $A^+ \setminus A$ to B^+. It is easy to see that in this instance, $|J \cap B^+|$ is $(\Delta+1)$-Lipschitz. As $v(F) + |A^+ \setminus A| \leq \frac{2n}{R} + 4\ell \leq 5\ell$, by Lemma 101, for any $t > 0$, we have that

$$\mathbf{P}\left[||J \cap B^+| - \mathbf{E}[|J \cap B^+|]| > t\right] \leq 2\exp\left\{-\frac{2t^2}{(\Delta+1)^2 5\ell}\right\}.$$

Choosing $t := \alpha \mathbf{E}[|J \cap B^+|] \geq \frac{\alpha\gamma\varepsilon'' n}{8}$, we see that

$$\mathbf{P}\left[||J \cap B^+| - \mathbf{E}[|J \cap B^+|]| > \alpha\mathbf{E}[|J \cap B^+|]\right] \leq 2\exp\left\{-\frac{\alpha^2\gamma\varepsilon''^2}{160(\Delta+1)^2} n\right\}.$$

As the number of choices for J is at most m, we see that for n large enough we have with probability at least 0.9, for each $J \in \mathfrak{I}$, $|J \cap B^+| = (1 \pm \alpha)\mathbf{E}[|J \cap B^+|]$. □

Lemma 123. *With probability at least* 0.9*, for every choice of* $I \in \mathfrak{I}$, $p \in [\Delta]$ *and* $a_1, \ldots, a_p \in I \cap A^+$ *we have that*

$$|\mathrm{Sol}(B^+, D^+, I; a_1, \ldots, a_p)| = (1 \pm \alpha)\mathbf{E}[|\mathrm{Sol}(B^+, D^+, I; a_1, \ldots, a_p)|].$$

Proof. We shall assume without lack of generality that $a_1, \ldots, a_q$ are pairwise distinct and q is the number of distinct labels among $a_1, \ldots, a_p$. Consider the random variable $|\mathrm{Sol}(B^+, D^+, I; a_1, \ldots, a_p)|$, and note that it is the same as $X := |\mathrm{Sol}(B^+, D^+, I; a_1, \ldots, a_q)|$. So we can in fact assume that $q = p$.

By Corollary 120, (iii) we have that

$$\begin{aligned}\mathbf{E}[X] &= (1 \pm \alpha) M_A^{(1\pm\alpha)^{2\Delta+2}} M_C^{p(1\pm\alpha)^{2\Delta+2}} |\mathrm{Sol}(A^+, C^+, I; a_1, \ldots, a_p)| \\ &\geq (1-\alpha)(M_A - \delta)(M_C - \delta)^p (1-\alpha)|\bar{I}|\frac{|A|}{|\mathbb{A}|}\left(\frac{|C|}{|\mathbb{C}|}\right)^p \\ &\geq \frac{1}{2}\left(\frac{\varepsilon''}{2}\right)^{p+1}\gamma n.\end{aligned}$$

We will now show that X is strongly concentrated around its expectation. Again X can be though of as defined on the product probability space $[0,1]^{v(F)} \times \{0,1\}^{|A^+ \setminus A|+|C^+ \setminus C|}$. Changing the label of any vertex can switch at most $1+\Delta+\Delta^2 \leq 2\Delta^2$ labels of vertices or edges from (B^+, D^+) to its complement. This in turn can affect at most $2p\Delta^2 \leq 2\Delta^3$ elements of $\mathrm{Sol}(B^+, D^+, I; a_1, \dots, a_p)$, so that X is $(2\Delta^3)$-Lipschitz. Then by Lemma 101, we have that

$$\mathbf{P}\left[|X - \mathbf{E}[X]| > \alpha \mathbf{E}[X]\right] \leq 2\exp\left\{-\frac{2(\alpha \mathbf{E}[X])^2}{4\Delta^6(9\ell)}\right\} \leq 2\exp\left\{-\frac{\alpha^2\gamma}{72\Delta^6}\left(\frac{\varepsilon''}{2}\right)^{2(p+1)} n\right\}.$$

There are at most m choices for I, at most Δ choices for p, and at most $(\gamma n)^\Delta$ choices for $a_1, \dots, a_p$. Thus for n large enough we have

$$(1+\varepsilon)n\Delta(\gamma n)^\Delta 2\exp\left\{-\frac{\alpha^2\gamma}{72\Delta^6}\left(\frac{\varepsilon''}{2}\right)^{2(\Delta+1)} n\right\} < 0.1.$$

It follows that with probability at least 0.9, for all choices of I, p and $a_1, \dots, a_p \in I \cap A^+$ we have $|\mathrm{Sol}(B^+, D^+, I; a_1, \dots, a_p)| = (1 \pm \alpha)\mathbf{E}[|\mathrm{Sol}(B^+, D^+, I; a_1, \dots, a_p)|]$.. □

Lemma 124. *For $c \in C$ define the random variable*

$$X_c = |\{(a, I) | I \in \mathcal{I}, a \in B^+ \cap I, a \oplus c \in B^+ \cap \bar{I}\}|.$$

With probability at least 0.9*, for every choice of $c \in C$ we have $|X_c| = (1 \pm \alpha)\mathbf{E}[|X_c|]$.*

Proof. By Corollary 120, (iv),

$$\begin{aligned}\mathbf{E}[X_c] &= (1 \pm \alpha)^2 M_A^{2(1 \pm 2\alpha)^{2\Delta+2}} |I|^2 \frac{|A|^2}{|\mathbb{A}|^2} \\ &\geq (1-\alpha)^2 (M_A - \delta)^2 |I|^2 \frac{|A|^2}{|\mathbb{A}|^2} \\ &\geq \left(\frac{\varepsilon'' \gamma n}{4}\right)^2.\end{aligned}$$

We show that X_c is strongly concentrated around its expectation. Again X_c is a random variable on $[0,1]^{v(F)} \times \{0,1\}^{|A^+ \setminus A|}$, and we see that it is $2(\Delta+1)\ell$-Lipschitz (indeed, changing the label of a vertex can affect at most Δ other labels, and each in turn can affect at most 2ℓ pairs (a, I)). Then by Lemma 101, we have that

$$\mathbf{P}\left[|X_c - \mathbf{E}[X_c]| > \alpha \mathbf{E}[X_c]\right] \leq 2\exp\left\{-\frac{2(\alpha \mathbf{E}[X_c])^2}{20(\Delta+1)^2\ell^3}\right\} \leq 2\exp\left\{-\frac{\alpha^2\gamma\varepsilon''^4}{160(\Delta+1)^2} n\right\}.$$

There are at most m choices for $c \in C$. So with probability at least 0.9 we see that for any $c \in C$, $|X_c| = (1 \pm \alpha)\mathbf{E}[|X_c|]$, as desired. □

A slight modification of the proof of Lemma 122 shows that with the same probability,

$$|B| = (1 \pm \alpha)\mathbf{E}[|B|], \quad \text{and} \tag{4.17}$$
$$|D| = (1 \pm \alpha)\mathbf{E}[|D|]. \tag{4.18}$$

We shall now assume the event that (4.17), (4.18), as well as Lemmas 122, 123 and 124 hold. The union bound shows that this event occurs with probability at least 0.7.

Recall that $\mathbf{E}[|B|] = (1 \pm \alpha)M_A^{(1\pm 2\alpha)^{2\Delta+2}}|A|$. As $M_A^{(1\pm 2\alpha)^{2\Delta+2}} = M_A \pm \delta$ by Lemma 121, we see that $\mathbf{E}[|B|] = (1 \pm \alpha)(M_A \pm \delta)|A|$. Hence $M_A = (1 \pm 2\alpha)\frac{\mathbf{E}[|B|]}{|A|} \pm \delta$.

As a consequence,

$$\begin{aligned}
M_A^{(1\pm 2\alpha)^{2\Delta+2}} = M_A \pm \delta &= (1 \pm 2\alpha)\frac{\mathbf{E}[|B|]}{|A|} \pm 2\delta \\
&= (1 \pm 2\alpha)\left(1 \pm 3\delta\frac{|A|}{\mathbf{E}[|B|]}\right)\frac{\mathbf{E}[|B|]}{|A|} \\
&= (1 \pm 2\alpha)(1 \pm 6\delta)\frac{\mathbf{E}[|B|]}{|A|}, \text{ as } \mathbf{E}[|B|] \geq \frac{|A|}{2}, \text{ by (4.15)}, \\
&= (1 \pm 2\alpha)^2(1 \pm 6\delta)\frac{|B|}{|A|}, \text{ from (4.17)}.
\end{aligned} \tag{4.19}$$

Similarly,

$$M_C^{(1\pm 2\alpha)^{2\Delta+2}} = (1 \pm 2\alpha)^2(1 \pm 6\delta)\frac{|D|}{|C|}. \tag{4.20}$$

We now check the first condition of quasirandomness. Let $J \in \mathcal{I}$ arbitrary. By Lemma 122,

$$\begin{aligned}
|J \cap B^+| &= (1 \pm \alpha)\mathbf{E}[|J \cap B^+|] \\
&= (1 \pm \alpha)^2 M_A^{(1\pm 2\alpha)^{2\Delta+2}}|J \cap A^+| \\
&\overset{(4.19)}{=} (1 \pm 2\alpha)^5(1 \pm 6\delta)\frac{|B|}{|A|}|J|\frac{|A|}{|\mathbb{A}|} \\
&= (1 \pm \beta)\frac{|B|}{|\mathbb{A}|}.
\end{aligned}$$

We check condition b) of quasirandomness. Let $I \in \mathcal{I}$, $p \in [\Delta]$ and $a_1, \ldots, a_p \in I \cap A^+$ distinct. Then by Lemma 123,

$$\begin{aligned}
|\mathrm{Sol}(B^+, D^+, I; a_1, \ldots, a_p)| &= (1 \pm \alpha)^2 M_A^{(1\pm\alpha)^{2\Delta+2}} M_C^{p(1\pm\alpha)^{2\Delta+2}}|\mathrm{Sol}(A^+, C^+, I; a_1, \ldots, a_p)| \\
&\overset{(4.19),\,(4.20)}{=} (1 \pm 2\alpha)^{5+2p}(1 \pm 6\delta)^{p+1}\frac{|B|}{|A|}\left(\frac{|D|}{|C|}\right)^p |\bar{I}|\frac{|A|}{|\mathbb{A}|}\left(\frac{|C|}{|\mathbb{C}|}\right)^p \\
&= (1 \pm \beta)|\bar{I}|\frac{|B|}{|\mathbb{A}|}\left(\frac{|D|}{|\mathbb{C}|}\right)^p.
\end{aligned}$$

Finally, we check condition c) of quasirandomness. Let $c \in C$ be arbitrary. Then by Lemma 124,

$$\begin{aligned}
|X_c| &= (1 \pm \alpha)^3 M_A^{2(1\pm 2\alpha)^{2\Delta+2}} |I|^2 \frac{|A|^2}{|\mathbb{A}|^2} \\
&\overset{(4.19)}{=} (1 \pm 2\alpha)^7 (1 \pm 6\delta)^2 \left(\frac{|B|}{|A|}\right)^2 |I|^2 \frac{|A|^2}{|\mathbb{A}|^2} \\
&= (1 \pm \beta) |I|^2 \frac{|B|^2}{|\mathbb{A}|^2}.
\end{aligned}$$

Thus with probability at least 0.7, (B^+, D^+) is (β, ℓ, Δ)-quasirandom. Hence with non-zero probability there exists ψ such that all the claims of the lemma hold. □

4.3.5 Finishing the proof of Theorem 107

In order to apply the nibble lemma we will partition the initial tree into several smaller subtrees of roughly equal order. We shall need the following well-known fact.

Fact 1. *Every tree T contains a vertex v such that each component of $T - v$ has order at most $v(T)/2$.*

Proof. Let us consider T with root v at an arbitrary leaf. Sequentially, if one component C of $T - v$ has order more than $v(T)/2$ then we move v to the unique neighbor of v in C. After this move, each vertex outside of C still lies in a component of $T - v$ of order at most $v(T)/2$. Each vertex of C now lies in a component of order strictly less than $v(C)$. Thus, the process must eventually terminate. □

Lemma 125. *Let T be any n-vertex tree of degree at most $\Delta \geq 2$, and let $n \geq R \geq 1$. Then by deleting at most $\frac{n}{R}$ edges, T can be partitioned into R forests $F_1, \dots, F_R$ each of order $\frac{n}{R} \pm \Delta R$. Moreover, each forest F_i contains trees of order between R and ΔR.*

Proof. We first prove by induction on n that T has a partition $\mathcal{P}$ into subtrees of order between R and ΔR.

If $n = R$, then the statement is clearly true. So assume $n > R$, and let $v \in V(T)$ be a leaf and let $\mathcal{P}'$ be a partition of $T - v$ into trees of order between R and ΔR, which exists by the induction hypothesis. Let u be the unique neighbor of v in T, and assume u belongs to the tree $T' \in \mathcal{P}'$. If $v(T') + 1 \leq \Delta R$ then we can add v to T' and we are done. Otherwise, by Fact 1 we can find a vertex $w \in V(T')$ such that any tree in $T' - w$ has at most $\frac{\Delta R}{2}$ vertices. As w has degree at most Δ, there exists a subtree in $T' - w$ that has at least R vertices. Thus we can detach this subtree from T', forming two new trees T'_1 and T'_2, and then adjoin v to T'_1, if $u \in T'_1$, or to T'_2, otherwise. Replacing T' by T'_1 and T'_2 gives the required partition $\mathcal{P}$ of T.

Note that we delete an edge only when a new tree is formed. As there are at most $\frac{n}{R}$ trees in $\mathcal{P}$, we have deleted in total at most $\frac{n}{R}$ edges.

We distribute the trees in $\mathcal{P}$ to R forests $F_1, \dots, F_R$ such that the number $M = \sum_{i=1}^{R} \left| v(F_i) - \frac{n}{R} \right|^2$ is minimized. We claim that for all i we have $|v(F_i) - \frac{n}{R}| \leq \Delta R$.

Indeed, assume for a contradiction that there exists F_i with $|v(F_i) - \frac{n}{R}| > \Delta R$. If $v(F_i) > \frac{n}{R} + \Delta R$, then we can choose any tree from F_i and move it to another forest F_j with at most $\frac{n}{R}$ vertices (such a forest always exist, by averaging). Similarly, if $v(F_i) < \frac{n}{R} - \Delta R$, we can choose a forest F_j with at least $\frac{n}{R}$ vertices and move any tree from it to F_i (again, averaging shows that such a forest always exist). As any tree has at most ΔR vertices, in both cases we obtain a smaller value of M, a contradiction. Consequently $v(F_i) = \frac{n}{R} \pm \Delta R$ for all i, completing the proof. □

We are now ready to prove Theorem 107.

Proof of Theorem 107. We define constants $n_0 \geq 1$ and

$$\alpha_R \ll \alpha_{R-1} \ll \ldots \ll \alpha_0 := \frac{1}{2},$$

in the following way.

We set $\alpha_0 = \frac{1}{2}$ and apply Lemma 114 repeatedly to obtain $\alpha_1, \alpha_2, \ldots, \alpha_R$ (we input ε, Δ and α_i to get α_{i+1}). This will also provide an $n_0 \geq 1$.

We now show that Theorem 107 holds for α_R and n_0.

Let $(A^+, C^+) \subseteq (\mathbb{A}^+, \mathbb{C}^+)$ such that $|A|, |C| \geq (1 + \frac{\varepsilon}{2})n$ and (A^+, C^+) is (α_R, ℓ, Δ)-quasirandom. Let $n \geq n_0$ and T be any n-vertex tree of maximum degree at most Δ. We use Lemma 125 to partition T into R forests $F_1, \ldots, F_R$ each of order $\frac{n}{R} \pm \Delta R$, such that any forest F_i has trees of order between R and ΔR. To obtain this partition we delete a set E of at most $\frac{n}{R}$ edges.

Set $B_R^+ := A^+$ and $D_R^+ := C^+$. Starting at $i = R$ and decreasing i up to 1, we apply Lemma 114 repeatedly to α_{i-1}, F_i and (B_i^+, D_i^+) to get an $(\varepsilon'' v(F_i))$-almost graceful labelling $\psi^i : V(F_i) \to \mathbb{N} \cup \{*\}$ with $\mathrm{im}(\psi^i) \subseteq B_i \cup \{*\}$ and $\mathrm{im}(\psi_*^i) \subseteq D_i \cup \{0, *\}$, and an $(\alpha_{i-1}, \ell, \Delta)$-quasirandom pair (B_{i-1}^+, D_{i-1}^+) with $B_{i-1} = B_i \setminus \mathrm{im}(\psi^i)$ and $D_{i-1} = D_i \setminus \mathrm{im}(\psi_*^i)$.

The hypotheses of Lemma 114 are always satisfied. Indeed,

$$|B_{i-1}| \geq |B_R| - \sum_{j=i}^{R} \mathrm{im}(\psi^j) \geq \frac{\varepsilon}{2} n \geq \frac{\varepsilon'}{3}(m - 2\ell + 1) = \frac{\varepsilon'}{3}|\mathbb{A}| \geq \varepsilon''|\mathbb{A}|,$$

as $\varepsilon' \leq \frac{3\varepsilon}{2(1+\varepsilon)}$, and similarly $|D_{i-1}| \geq \frac{\varepsilon'}{3}|\mathbb{C}| \geq \varepsilon''|\mathbb{C}|$ for all i.

Note further that $B_1 \subseteq B_2 \subseteq \ldots B_R \subseteq A$ and $D_1 \subseteq D_2 \subseteq \ldots D_R \subseteq C$.

We now take $\psi := \bigcup_{i=1}^R \psi^i$ and extend ψ to E by setting $\psi_*(e) = *, \forall e \in E$. As $\mathrm{im}(\psi^i), 1 \leq i \leq R$, respectively $\mathrm{im}(\psi_*^i), 1 \leq i \leq R$, are pairwise disjoint by construction, and $|E| \leq \frac{n}{R} \leq \frac{\varepsilon' n}{3}$, we see that by definition ψ is an $(\varepsilon' v(T))$-almost graceful labelling of T with codomain (A, C). This proves the theorem. □

4.4 Proof of Theorem 39

We first have a lemma.

Lemma 126. *For any $\Delta \geq 1$ and $\delta > \frac{2}{\Delta+1}$, with probability at least $1 - \frac{\Delta+1}{\Delta!\delta^{\Delta+1}}$ the following holds. For all $k \geq \Delta$, there are less than $\delta^{k+2}n$ vertices of degree greater than k in a random n-vertex tree.*

Proof. We consider the vertices of the tree labelled with numbers from $[n]$.

We use the following estimate (see [85]): for any $x \in [n]$, $\mathbf{P}[\deg(x) > k] \leq \frac{1}{k!}$. Hence by Markov's inequality, for a fixed $k \geq \Delta$,

$$\mathbf{P}[\text{there are at least } \delta^{k+2}n \text{ vertices of degree greater than } k] \leq \frac{1}{k!\delta^{k+2}}.$$

We now use the union bound:

$$\begin{aligned}
\sum_{k\geq\Delta} \frac{1}{k!\delta^{k+2}} &\leq \frac{1}{\Delta!\delta^{\Delta+2}} \sum_{k\geq\Delta} \frac{1}{((\Delta+1)\dots k)\delta^{k-\Delta}} \\
&\leq \frac{1}{\Delta!\delta^{\Delta+2}} \sum_{k\geq\Delta} \frac{1}{((\Delta+1)\delta)^{k-\Delta}} \\
&\leq \frac{\Delta+1}{\Delta!\delta^{\Delta+1}(\delta(\Delta+1)-1)} \leq \frac{\Delta+1}{\Delta!\delta^{\Delta+1}}.
\end{aligned}$$

□

We will mimic the proof of Theorem 38 to obtain a proof of Theorem 39.

Proof of Theorem 39. Let $\varepsilon > 0$ be given. Let p_n be the probability that a random n-vertex tree has an $(1+\varepsilon)n$-graceful labelling. We will show that $p_n = 1 - o(1)$.

Let $\{\Delta_n\}_{n\geq 1}$ be a sequence of numbers tending to infinity sufficiently slower and assume n is large enough.

Given Δ_n and ε, we define the numbers d, ε', γ, ℓ, and m and the sets $\mathbb{A}^+$ and $\mathbb{C}^+$ as in Setup 103. We apply Theorem 107 to obtain the parameter α. Then let $(A^+_{\text{rep}}, C^+_{\text{rep}})$ be the output of Lemma 109 for parameters $\varepsilon, \Delta_{\text{L109}} = \Psi := \frac{\log n}{100\log(1/\varepsilon)}$ and $\alpha_{\text{L109}} := \alpha$. Note that $(A^+_{\text{rep}}, C^+_{\text{rep}})$ is an $(\frac{m}{4}, d, \Psi)$-repair pair.

Let T be an n-vertex random tree. We assume that T has maximum degree $\Delta(T) \leq 2\frac{\log n}{\log\log n} < \Psi$. By Corollary 1, [85], this holds with probability $1 - o(1)$.

Let V_0 be the set of vertices in T of degree greater than Δ_n. Then $F := T - V_0$ is a forest of maximum degree at most Δ_n.

Set $\delta := \frac{d}{2} = \frac{\varepsilon}{32}$. By taking n sufficiently large we may assume that $\delta > \frac{2}{\Delta_n+1}$.

We will assume that for all $k \geq \Delta_n$, there are less than $\delta^{k+2}n$ vertices of degree greater than k in V_0. By Lemma 126, this happens with probability at least $1 - \frac{\Delta_n+1}{\Delta_n!\delta^{\Delta_n+1}}$.

We now rejoin the trees of F into a new tree T' such that $\Delta(T') \leq \Delta_n$ and $F \subseteq T'$. This is possible, for example, by joining two leaves from distinct trees at a time, until only one tree is left.

We then apply Theorem 107 to T' and $(\mathbb{A}^+ \setminus A^+_{\text{rep}}, \mathbb{C}^+ \setminus C^+_{\text{rep}})$, in order to obtain an $(\varepsilon' n)$-almost graceful labelling ψ of T' with codomain $(\mathbb{A} \setminus A_{\text{rep}}, \mathbb{C} \setminus C_{\text{rep}})$. Clearly this gives an $(\varepsilon' n)$-almost graceful labelling of F, which we also denote by ψ.

We will now repair the conflicts in F, as well as find suitable labels for the vertices in V_0. We will argue similarly to the proof of Lemma 110. However, we cannot apply Lemma 110 directly, as the degrees of the vertices in V_0 can be quite high. Furthermore, note that repairing the conflicts and jokers in F will amount to relabeling a set V_1 of at most $\varepsilon' n$ conflicting vertices.

We start by labelling the vertices in V_0. Let $v_1, \ldots, v_r$ be an ordering of the vertices in V_0, so that the degrees are monotone non-increasing. We label the vertices v_i in turn, using labels from A_{rep}. At each step, the size of A_{rep} decreases by 1, while the size of C_{rep} decreases by at most $\deg(v_i)$. When we process vertex v_i, we consider the distinct labels $b_1, \ldots, b_s$ appearing on neighbors of v_i in the graph induced by $V(F) \cup \{v_1, \ldots, v_i\}$. Assume v_i has degree $j > \Delta_n$ in T. Then $s \leq j$. As $j \leq \Phi$, there were initially at least $\frac{md^{j+1}}{4}$ solutions in $(A_{\text{rep}}, C_{\text{rep}})$ to the tuple $(b_1, \ldots, b_s)$. By assumption, for all $k \geq j$, there are at most $\delta^{k+1} n$ vertices of degree k in V_0. Consequently there are still at least

$$\frac{md^{j+1}}{4} - (j+1) \sum_{k \geq j} (k+1) \delta^{k+1} n$$

available solutions to $(b_1, \ldots, b_s)$ in $(A_{\text{rep}}, C_{\text{rep}})$ (the $j+1$ term comes from the fact that we have at most $j+1$ labels, $\psi(v_i), |\psi(v_i) - b_1|, \ldots, |\psi(v_i) - b_s|$, that need to be in A_{rep}, respectively C_{rep}). As $\delta = \frac{d}{2}$, this is at least

$$\frac{md^{j+1}}{4} - 8(j+1)^2 \delta^{j+1} n \geq nd^{j+1} \left(\frac{1}{4} - \frac{8(j+1)^2}{2^{j+1}} \right) \geq nd^{j+1} \left(\frac{1}{4} - \frac{8(\Delta_n+1)^2}{2^{\Delta_n+1}} \right),$$

which is positive for $\Delta_n \geq 12$, so that we have at least one solution a which was not used before. Let us set $\psi(v_i) := a$. We now update the set A_{rep} by deleting a, and the set C_{rep} by removing $|a - b_1|, \ldots, |a - b_s|$. Continuing in this manner we correctly label all vertices in V_0.

We now relabel the vertices in V_1. All of these vertices had degree at most Δ_n to start with, so that at any step we still have at least

$$\frac{md^{\Delta_n+1}}{4} - (8(\Delta_n+1)^2 \delta^{\Delta_n+1} + (\Delta_n^2+1)\varepsilon')n$$

available labels to use. This quantity is positive for $\varepsilon' \leq \frac{d^{\Delta_n+1}}{16\Delta_n^2}$ and $\Delta_n \geq 12$. Hence we can process all the vertices in V_1, obtaining a graceful labelling of the initial tree T.

Finally, we need to bound the probability of obtaining this labelling. By applying the union bound on the events that $\Delta(T) > 2\frac{\log n}{\log\log n}$ and that the conclusion of Lemma 126 does not hold, we get that $p_n \geq 1 - o(1) - \frac{\Delta_n+1}{\Delta_n! \delta^{\Delta_n+1}}$. Hence $p_n = 1 - o(1)$, proving the theorem.

□

4.5 Remarks on the results

Ringel's conjecture (Conjecture 34) was wide open until recently, when a more general result concerning bounded-degree trees was proved in [12]. The method of proof presented here was modeled after the approach taken in [12]: the nibble method was combined with the key idea of first independently labelling (respectively, packing) one color class of each small tree and then extending the labelling (respectively, packing) to the other color class.

Chapter 5

The Towers of Hanoi

5.1 Definitions and auxiliary results

For $n \in \mathbb{N}$ let $[n] = \{0, 1, \ldots, n-1\}$ denote the set of natural numbers smaller than n. Given p pegs and N disks, we label the disks using numbers from $[N]$, in increasing order according to their size: the smallest disk receives the label 0, the second smallest disk the label 1, and so on, with the largest disk receiving the label $N-1$. We also label the pegs using numbers from $[p]$.

We now give a more precise description of Φ as follows.

Definition 127 (The operators Δ_p and ∇_p)**.** *Let $p \geq 3$. We define for all $n \geq 0$ the values*

$$\Delta_p(n) := \binom{n+p-3}{p-2}$$

and

$$\nabla_p(n) := \max\{k \geq 0 : \Delta_p(k) \leq n\}.$$

Note that $\Delta_p(0) = 0$ and hence the maximum is not taken over an empty set. Then it can be shown (see [75], [98]) that for all $p \geq 3$ and $N \geq 1$,

$$\Phi(p, N) = 2^{\nabla_p 0} + 2^{\nabla_p 1} + \ldots + 2^{\nabla_p (N-1)}. \tag{5.1}$$

This formula (in a similar form) already appears in the articles of Frame [44] and Stewart [122], and has been rediscovered many times. In the case $p = 4$, it can be written more compactly as follows. Let $N - 1 = \Delta_4 m + t, 0 \leq t \leq m$. Then

$$\Phi(4, N) = 1 + (m+t)2^m. \tag{5.2}$$

Note for later use the following property of Δ_p:

$$\Delta_p n = \Delta_p(n-1) + \Delta_{p-1} n, \quad \forall p \geq 4, n \geq 1. \tag{5.3}$$

Let $p \geq 3$. We call an arrangement of disks on p pegs a *configuration* if no disk is placed on top of a larger one. Note that the set of configurations of N disks can be identified with the set $[p]^{[N]}$ of functions $[N] \to [p]$, in particular, given a configuration $\mathbf{u}$, we let $\mathbf{u}^{-1}(x)$ denote the set of disks placed on peg x. Furthermore, $\mathbf{u}|_S$ represents the configuration obtained from $\mathbf{u}$ by deleting all disks in $[N] \setminus S$.

We define the *Hanoi graph* $\mathcal{H}(p,N)$ as having vertex set $[p]^{[N]}$, and an edge between two vertices u and v if the corresponding configurations can be obtained from one another by a single disk move. We consider $\mathcal{H}(p,N)$ to be a metric space with the usual metric that has distance 1 between any two adjacent vertices.

A path in the Hanoi graph is any map $\gamma : [T] \to \mathcal{H}(p,N)$ with the property that $\gamma(t)$ and $\gamma(t+1)$ are adjacent vertices for all $0 \le t \le T-2$. Thus γ corresponds to a sequence of consecutive disk moves, starting at the configuration $\gamma(0)$. Note that we do not require the vertices of γ to be distinct.

If $\gamma : [T] \to \mathcal{H}(p,N)$ is any path, we let $\ell(\gamma) := T-1$ denote its length, in other words, the number of disk moves represented by γ. We sometimes write γ_t instead of $\gamma(t)$, to denote the configuration at time t. For any $0 \le t \le T-2$, we let $D_{\gamma,t}$ be the unique disk moved between $\gamma(t)$ and $\gamma(t+1)$. We say that *$D_{\gamma,t}$ is moved at time t.* For any $0 \le t_1 \le t_2 \le T-1$, we let $\gamma|_{[t_1,t_2]}$ denote the path going through configurations $\gamma(t_1), \gamma(t_1+1), \ldots, \gamma(t_2)$.

The path γ is called *essential* if any disk is moved by γ at least once. Note that in this case the path $\gamma^* : [T] \to \mathcal{H}(p,N)$ given by $\gamma^*(t) := \gamma(T-t-1)$ is also essential. By definition,

$$\Gamma(p,N) := \min\{\ell(\gamma) : \gamma \text{ is an essential path in } \mathcal{H}(p,N)\}.$$

The structure of shortest paths (geodesics) in the Hanoi graph has been studied before (see [4]). Note that an essential path need not be a geodesic.

We now introduce a crucial definition, due to Bousch. Let $E \subset \mathbb{N}$ be finite. For any $L \in \mathbb{N}$ we define

$$\Psi_L(E) := (1-L)2^L - 1 + \sum_{n \in E} 2^{\min\{\nabla_4 n, L\}}, \tag{5.4}$$

and further

$$\Psi(E) := \sup_{L \in \mathbb{N}} \Psi_L(E). \tag{5.5}$$

The value $\Psi(E)$ is always a natural number, as $\Psi_L(E)$ becomes negative for large L, and $\Psi_0(E) = |E|$. Bousch showed the following.

Theorem 128 (Theorem 2.9, [16])**.** *Let $a \in [4]$ arbitrary. Let $\mathbf{u}, \mathbf{v} \in \mathcal{H}(4,N)$ be two configurations such that in $\mathbf{v}$, peg a and some other peg b do not contain any disks. Then $d(\mathbf{u},\mathbf{v}) \ge \Psi(\mathbf{u}^{-1}(a))$.*

It turns out that $\Psi([N]) = \frac{\Phi(4,N+1)-1}{2}$. In combination with Theorem 128, this easily implies Theorem 40. We record this last fact below.

Lemma 129. *For all $N \ge 2$,*

$$\Psi([N]) = \frac{\Phi(4,N+1)-1}{2} = \min_{\substack{a+b=N \\ a,b \ge 1}} \left\{\Phi(4,a) + \Phi(3,b)\right\}.$$

Proof. The first identity is Lemma 2.2 from [16]. The second follows from (1.9). $\square$

Let $N \geq 1$ and $\mathbf{u}$ be the configuration with all N disks on peg 0. A configuration $\mathbf{c}$ of N disks on pegs $\{2,3\}$ such that $d(\mathbf{u},\mathbf{c}) \leq \frac{\Phi(4,N+1)-1}{2}$ in $\mathcal{H}(4,N)$ is called *a midpoint configuration of N disks on 4 pegs*. The existence of such configurations for all N follows from the Frame-Stewart algorithm. Lemma 129 together with Theorem 128 show that in fact $d(\mathbf{u},\mathbf{c}) = \frac{\Phi(4,N+1)-1}{2}$ whenever $\mathbf{c}$ is a midpoint configuration, but we will not use this stronger statement.

We shall also need the following two lemmas.

Lemma 130 (Lemma 2.6, [16])**.** *Let $A \subset \mathbb{N}$ finite, and s a natural number such that $A - [\Delta_4 s]$ has at most s elements. Then*

$$\Psi(A) - \Psi(A - \{a\}) \leq 2^{s-1}$$

for all $a \in A$.

Lemma 131 (Lemma 2.8, [16])**.** *Let $A, B \subset \mathbb{N}$ be finite sets. Then*

$$\Psi(A) + \Psi(B) \geq \frac{\Phi(4,N+3)-5}{4},$$

where $N := |A \cup B|$.

Finally, we shall need the following recursive lower bound for Γ.

Lemma 132 (Corollary 1, [23])**.** *Let $p \geq 4$ and $N \geq 2$. Then for every $1 \leq \ell \leq N-1$,*

$$\Gamma(p,N) \geq 2\min\{\Gamma(p,N-\ell), \Gamma(p-1,\ell)\}.$$

5.2 The length of the shortest essential path

We start with the following lemma.

Lemma 133. *Let $N \geq 1$ and $\mathbf{u}, \mathbf{v} \in \mathcal{H}(4, N)$ such that $\mathbf{u}^{-1}(\{2,3\}) = \emptyset$ and $\mathbf{v}^{-1}(\{0,1\}) = \emptyset$. Then*

$$d(\mathbf{u}, \mathbf{v}) \geq 1 + \frac{\Phi(4, N+2) - 5}{4}.$$

Moreover, this inequality is tight.

Proof. Let $\gamma : [T] \to \mathcal{H}(4, N)$ be a shortest path between $\mathbf{u}$ and $\mathbf{v}$. Let $t_1 \in [T-1]$ be the first time when the disk $N-1$ moves to one of the pegs 2 and 3. Then we may assume without lack of generality that $\gamma_{t_1}(N-1) = 0$ and $\gamma_{t_1+1}(N-1) = 2$.

Set

$$\begin{aligned} A &= \{z \in [N-1] : \gamma_{t_1}(z) = 1\} \\ B &= \{z \in [N-1] : \gamma_{t_1}(z) = 3\}. \end{aligned}$$

Note that $A \dot\cup B = [N-1]$, as no other disk besides $N-1$ is on pegs 0 or 2 at time t_1.

By Theorem 128,

$$d(\gamma(t_1), \gamma(0)) \geq \Psi(B),$$

as pegs 2 and 3 are empty in $\gamma(0) = \mathbf{u}$.

Similarly, by Theorem 128 and the fact that all disks are placed on pegs 2 and 3 in $\mathbf{v}$,

$$d(\gamma(t_1+1), \gamma(T-1)) \geq \Psi(A).$$

Consequently by Lemma 131, and the fact that the disk $N-1$ moves once at time t_1,

$$\begin{aligned} \ell(\gamma) &\geq d(\gamma(0), \gamma(t_1)) + 1 + d(\gamma(t_1+1), \gamma(T-1)) \\ &\geq \Psi(B) + 1 + \Psi(A) \\ &\geq 1 + \frac{\Phi(4, N+2) - 5}{4}. \end{aligned}$$

We now show that the inequality is tight. Let $a, b \in \mathbb{N}$ arbitrary such that $a + b = N$ and $b \geq 1$. Consider a configuration $\mathbf{u}_{a,b}$ with the disk $N-1$ on peg 0, disks $N-b, N-b+1, \ldots, N-2$ on peg 1, and disks $0, \ldots, a-1$ arranged on pegs 0 and 1 in such a way that they form a midpoint configuration of a disks on 4 pegs.

Then we can move the disks $0, \ldots, a-1$ to peg 3 using at most $\frac{\Phi(4,a+1)-1}{2}$ moves.

Afterwards, we can move the disk $N-1$ to peg 2.

Finally, we can move disks $N-b, \ldots, N-2$ to peg 2 using $2^{b-1} - 1$ moves.

Let $\mathbf{v}_{a,b}$ be the resulting configuration. It has disks $N-b, \ldots, N-1$ on peg 2, and disks $0, \ldots, a-1$ on peg 3. Also

$$d(\mathbf{u}_{a,b}, \mathbf{v}_{a,b}) \leq \frac{\Phi(4, a+1) - 1}{2} + 2^{b-1} \leq \frac{\Phi(4, a+1) + 2^b - 1}{2} = \frac{\Phi(4, a+1) + \Phi(3, b)}{2}.$$

We now minimize over all choices of a and b. This gives configurations $\mathbf{u}$ and $\mathbf{v}$ such that

$$\begin{aligned} d(\mathbf{u},\mathbf{v}) &\leq \min_{\substack{a+b=N\\ b\geq 1}} \frac{\Phi(4,a+1)+\Phi(3,b)}{2} \\ &= \frac{\Phi(4,N+2)-1}{4}, \quad \text{by Lemma 129,} \\ &= 1+\frac{\Phi(4,N+2)-5}{4}. \end{aligned}$$

□

We would now like to extend this result to configurations which may share a peg, i.e. there is a peg which is occupied in both the starting and ending configuration. Surprisingly, this requires some more effort.

Lemma 134. *Let $N \geq 1$ and $\mathbf{u},\mathbf{v} \in \mathcal{H}(4,N)$ such that $\mathbf{u}^{-1}(\{2,3\}) = \emptyset$ and $\mathbf{v}^{-1}(\{0,3\}) = \emptyset$. If $\gamma : [T] \to \mathcal{H}(4,N)$ is any essential path between $\mathbf{u}$ and $\mathbf{v}$ then $\ell(\gamma) \geq \Psi(\mathbf{u}^{-1}(1))$.*

Proof. We prove the lemma by induction on N.

If $N = 1$ then $\Psi([1]) = 1$ and the claim trivially holds.

So assume $N \geq 2$. Let $E := \mathbf{u}^{-1}(1)$.

If $N-1 \notin E$ then we can restrict the path γ to the moves of the first $N-1$ disks. This gives a (possibly shorter) path γ' between $u|_{[N-1]}$ and $v|_{[N-1]}$, which is essential in $\mathcal{H}(4,N-1)$. By the induction hypothesis applied to $u|_{[N-1]}, v|_{[N-1]}$ and γ', we obtain

$$\ell(\gamma) \geq \ell(\gamma') \geq \Psi(E),$$

proving the claim in this case.

Hence we may assume that $N-1 \in E$. Let t_1 be the first time when the disk $N-1$ moves. Then $\gamma_{t_1}(N-1) = 1$. Set $a := \gamma_{t_1+1}(N-1)$.

Case 1. $a \neq 2$.

Then $a \in \{0,3\}$. Let π be the involution on $\{0,1,2,3\}$ which exchanges elements 1 and a. We modify γ into a new path γ' by letting $\gamma'|_{[0,t_1+1]} = \gamma|_{[0,t_1+1]}$ and setting for all $t > t_1+1$,

$$\begin{aligned} \gamma'_t(D) &= \pi \circ \gamma_t(D),\ D \in [N-1], \\ \gamma'_t(N-1) &= a. \end{aligned}$$

At time t_1+1, peg 1 is empty and peg a only contains the disk $N-1$. Hence all moves represented by γ' are valid moves. However, γ' may contain repeated states, so we may need to delete some in order to make it into a proper path. Note that in $\gamma'(T-1)$ pegs 1 and $3-a$ are empty, as in $\gamma(T-1)$ pegs a and $3-a$ were empty. Consequently by Theorem 128,

$$\ell(\gamma) \geq \ell(\gamma') \geq d(\mathbf{u}, \gamma'(T-1)) \geq \Psi(\mathbf{u}^{-1}(1)).$$

Case 2. $a = 2$.

By Theorem 128 and the fact that the pegs 1 and 2 are empty in $\gamma(t_1)|_{[N-1]}$, we have

$$d(\gamma(0), \gamma(t_1)) \geq d(\gamma(0)|_{[N-1]}, \gamma(t_1)|_{[N-1]}) \geq \Psi(E - \{N-1\}).$$

Also, by Lemma 133 and the fact that pegs 1 and 2 are empty in $\gamma(t_1+1)|_{[N-1]}$, while pegs 0 and 3 are empty in $\gamma(T-1)|_{[N-1]}$, we have

$$d(\gamma(t_1+1), \gamma(T-1)) \geq d(\gamma(t_1+1)|_{[N-1]}, \gamma(T-1)|_{[N-1]}) \geq 1 + \frac{\Phi(4, N+1) - 5}{4}.$$

Hence adding the move of the disk $N-1$ gives

$$\ell(\gamma) \geq \Psi(E - \{N-1\}) + 1 + \frac{\Phi(4, N+1) - 1}{4}.$$

Write $N = \Delta_4 m + t, 0 \leq t \leq m$. By Lemma 130 applied to E and $s := m$, we get

$$\Psi(E) - \Psi(E - \{N-1\}) \leq 2^{m-1}.$$

Also $\Phi(4, N+1) = 1 + (m+t)2^m$ and so $\frac{\Phi(4,N+1)-1}{4} = (m+t)2^{m-2}$.

If $N = 2$ then $m = t = 1$ and so $m + t \geq 2$.

If $N \geq 3$ then $\nabla_4 N = m \geq 2$ and again $m + t \geq 2$.

Thus in any case $m + t \geq 2$ and $(m+t)2^{m-2} \geq 2^{m-1}$. Hence

$$\ell(\gamma) \geq \Psi(E) - 2^{m-1} + 1 + 2^{m-1} > \Psi(E) = \Psi(\mathbf{u}^{-1}(1)).$$

□

Lemma 135. *Let $N \geq 1$ and $\mathbf{u}, \mathbf{v} \in \mathcal{H}(4, N)$ such that $\mathbf{u}^{-1}(\{2,3\}) = \mathbf{v}^{-1}(\{2,3\}) = \emptyset$. If $\gamma : [T] \to \mathcal{H}(4, N)$ is any essential path between $\mathbf{u}$ and $\mathbf{v}$ then $\ell(\gamma) \geq \Psi(\mathbf{u}^{-1}(1))$.*

The proof is nearly identical to that of Lemma 134, and so we omit it.

Lemma 136. *If $N \geq 1$ then $\frac{\Phi(4,N+1)-1}{2} \geq \frac{\Phi(4,N+2)-1}{4}$.*

Proof. We show the equivalent statement $2\Phi(4, N+1) - 2 \geq \Phi(4, N+2) - 1$.

Write $N + 1 = \Delta_4 m + t, 0 \leq t \leq m$. Then

$$\Phi(4, N+1) = 1 + (m+t-1)2^m,$$
$$\Phi(4, N+2) = 1 + (m+t)2^m.$$

So the desired inequality takes the form

$$(m+t-1)2^{m+1} \geq (m+t)2^m,$$

that is, $2(m+t-1) \geq m+t$, which is equivalent to $m+t \geq 2$. As $N+1 \geq 2$, we always have $m+t \geq 2$, proving the inequality. □

We are now ready to prove the counterpart to Lemma 133.

Lemma 137. *Let $N \geq 1$ and $\mathbf{u}, \mathbf{v} \in \mathcal{H}(4, N)$ such that $\mathbf{u}^{-1}(\{2,3\}) = \emptyset$ and $\mathbf{v}^{-1}(\{0,3\}) = \emptyset$. If $\gamma : [T] \to \mathcal{H}(4, N)$ is any essential path between $\mathbf{u}$ and $\mathbf{v}$ then*

$$\ell(\gamma) \geq 1 + \frac{\Phi(4, N+2) - 5}{4}.$$

Proof. If $\mathbf{v}^{-1}(1) = \emptyset$ then all disks are on peg 2 in $\mathbf{v}$. But pegs 2 and 3 are empty in $\mathbf{u}$, so by Theorem 128 and Lemma 129,

$$\ell(\gamma) \geq d(\mathbf{v}, \mathbf{u}) \geq \Psi(\mathbf{v}^{-1}(2)) = \Psi([N]) = \frac{\Phi(4, N+1) - 1}{2}.$$

By Lemma 136, this is at least $1 + \frac{\Phi(4,N+2)-5}{4}$, proving the claim in this case.

So we may assume that $\mathbf{v}^{-1}(1) \neq \emptyset$. Let D be the largest disk on peg 1 in $\mathbf{v}$.

Case 1. $D = N - 1$.

Let t_1 be the last time when D is not on peg 1. Then $\gamma_{t_1+1}(D) = 1$. Set $a := \gamma_{t_1}(D)$. We define b and c as follows.

a	b	c
0	2	3
2	3	0
3	2	0

Then all disks in $[N-1]$ are on pegs b and c at time t_1. Set

$$B := \{z \in [N-1] : \gamma_{t_1}(z) = b\}$$
$$C := \{z \in [N-1] : \gamma_{t_1}(z) = c\}$$

By Theorem 128, $d(\gamma(t_1), \mathbf{u}) \geq \Psi(B)$, as $b \in \{2,3\}$ and pegs 2 and 3 are empty in $\mathbf{u}$. Also, $d(\gamma(t_1+1), \mathbf{v}) \geq \Psi(C)$, as $c \in \{0,3\}$ and pegs 0 and 3 are empty in $\mathbf{v}$.

Consequently by Lemma 131 and the fact that $B \dot\cup C = [N-1]$,

$$\ell(\gamma) \geq d(\gamma(t_1), \mathbf{u}) + 1 + d(\gamma(t_1+1), \mathbf{v}) \geq 1 + \frac{\Phi(4, N+2) - 5}{4},$$

as desired.

Case 2. $D < N - 1$.

Then $\mathbf{v}(N-1) = 2$. Let t_1 be the last time when the disk $N-1$ moves from pegs $\{0,1\}$ to pegs $\{2,3\}$. Let t_2 be the last time when D is not on peg 1.

Let $a := \gamma_{t_1+1}(N-1)$ and set

$$b := \begin{cases} 2, & \text{if } a = 3, \\ 3, & \text{if } a = 2. \end{cases}$$

Define

$$c := \begin{cases} 0, & \text{if } \gamma_{t_1}(N-1) = 1, \\ 1, & \text{if } \gamma_{t_1}(N-1) = 0. \end{cases}$$

Further set

$$B := \{z \in [N-1] : \gamma_{t_1}(z) = b\}$$
$$C := \{z \in [N-1] : \gamma_{t_1}(z) = c\}$$

Note that $B \dot{\cup} C = [N-1]$. We now consider two subcases.

Case 2.1. $t_2 > t_1$.

Then

$$d(\gamma(t_1), \mathbf{u}) \geq d(\gamma(t_1)|_{[N-1]}, \mathbf{u}|_{[N-1]}) \geq \Psi(B),$$

as $b \in \{2, 3\}$ and pegs 2 and 3 are empty in $\mathbf{u}$.

We will now show that $\ell(\gamma|_{[t_1+1,T-1]}) \geq \Psi(C)$.

If $c = 0$, then by Theorem 128,

$$\ell(\gamma|_{[t_1+1,T-1]}) \geq d(\gamma(t_1+1), \mathbf{v}) \geq d(\gamma(t_1+1)|_{[N-1]}, \mathbf{v}|_{[N-1]}) \geq \Psi(C),$$

as $c \in \{0, 3\}$ and pegs 0 and 3 are empty in $\mathbf{v}$.

If $c = 1$, then we claim that $\gamma|_{[t_1+1,T-1]}$ is an essential path when restricted to the moves of the first $N-1$ disks. Indeed, the disks on peg 1 at time t_1+1 will all have to move, to make room for the disk D. The disks on peg b at time t_1+1 will have to move, as either $b = 3$ and $\mathbf{v}^{-1}(3) = \emptyset$, or $b = 2$, and $N-1$ is not yet on peg 2. Hence by Lemma 134, if $b = 3$, and Lemma 135, if $b = 2$,

$$\ell(\gamma|_{[t_1+1,T-1]}) \geq \Psi(C).$$

So

$$\ell(\gamma) \geq \Psi(B) + 1 + \Psi(C) \geq 1 + \frac{\Phi(4, N+2) - 5}{4},$$

as desired.

Case 2.2. $t_2 < t_1$.

Then $c = 1$, otherwise $\gamma_{t_1}(N-1) = 1$ and D is not on peg 1 at time t_1. We claim that $\gamma|_{[0,t_1]}$ is an essential path when restricted to the moves of the first $N-1$ disks. Indeed, the disks on peg 1 all moved at least once in the time interval $[0, t_1]$, as D is already in final position at time t_1, and the disks on peg b all moved, as $b \in \{2, 3\}$ and $\mathbf{u}^{-1}(\{2, 3\}) = \emptyset$. Hence by Lemma 134,

$$\ell(\gamma|_{[0,t_1]}) \geq \Psi(C).$$

We will now show that

$$\ell(\gamma|_{[t_1+1,T-1]}) \geq \Psi(B).$$

If $b = 3$ then the disk $N-1$ moves from peg 0 to peg 2 at time t_1. By Theorem 128 and the fact that pegs 0 and 3 are empty in $\mathbf{v}$,

$$\ell(\gamma|_{[t_1+1,T-1]}) \geq d(\gamma(t_1+1), \gamma(T-1)) \geq \Psi(B).$$

If $b = 2$, let $t_3 > t_1$ be the last time when $\gamma_{t_3}(N-1) \neq 2$. Then at time t_3, peg $b = 2$ and some other peg do not contain any disks smaller than $N-1$. So by Theorem 128,

$$\ell(\gamma|_{[t_1+1,T-1]}) \geq d(\gamma(t_1+1)|_{[N-1]}, \gamma(t_3)|_{[N-1]}) \geq \Psi(B).$$

Thus in any case,

$$\ell(\gamma) \geq \Psi(C) + 1 + \Psi(B) \geq 1 + \frac{\Phi(4, N+2) - 5}{4},$$

as claimed. □

Proof of Theorem 44. If $N \leq 3$, we can place disk i on peg i, for all $i \in [N]$. Then we can move the disks in turn to peg 3, from the largest to the smallest one. Consequently, $\Gamma(4, N) = N$ for $N \leq 3$. As $\Phi(4, 3) = 5$, it follows that the theorem holds for $N \leq 3$.

So assume $N \geq 4$. We will first show the following inequality:

$$\Gamma(4, N) \geq 3 + \frac{\Phi(4, N) - 5}{4}. \tag{5.6}$$

Let $\gamma : [T] \to \mathcal{H}(4, N)$ be a shortest essential path. Let t_1 be any time when the disk $N-1$ moves. We may assume without lack of generality that $\gamma_{t_1}(N-1) = 0$ and $\gamma_{t_1+1}(N-1) = 1$. We shall further assume that $\gamma_{t_1+1}(N-2) = 2$.

We now choose $t_2 \in [T]$ such that the disk $N-2$ moves at time t_2, and the difference $|t_2 - t_1|$ is minimal. Clearly t_2 exists, although there may be two distinct choices, if the disk $N-2$ moves before and after time t_1. If there are two possibilities for t_2, we choose one arbitrarily. Then by definition of t_2, the disk $N-2$ does not move in the time interval $[\min\{t_1, t_2+1\}, \max\{t_1+1, t_2\}]$.

Note that we can always replace γ with γ^*, t_1 with $t_1' := T - t_1 - 2$ and t_2 with $t_2' := T - t_2 - 2$. Then γ^* is still essential, $N-1$ moves at time t_1', $N-2$ moves at time t_2', and the difference $|t_1' - t_2'| = |t_1 - t_2|$ is still minimal.

First suppose peg 3 is empty at time t_1+1. By replacing γ with γ^* if necessary, we may assume that $t_2 > t_1$. At time t_1+1, all disks in $[N-1]$ are on peg 2, while at time t_2, pegs 2 and $\gamma_{t_2+1}(N-2)$ do not contain any disks from $[N-2]$. Consequently, by restricting to the first $N-2$ disks and applying Theorem 128, we get $\ell(\gamma|_{[t_1+1,t_2]}) \geq \frac{\Phi(4,N-1)-1}{2}$. Adding the 2 moves of the disks $N-1$ and $N-2$ and using Lemma 136, we get

$$\ell(\gamma) \geq 2 + \frac{\Phi(4, N-1) - 1}{2} \geq 3 + \frac{\Phi(4, N) - 5}{4}.$$

Therefore we may assume that peg 3 is not empty at time t_1+1. Let D be the largest disk on peg 3 at time t_1+1 and t_3 any time when the disk D moves. Note that $D \leq N-3$.

Case 1. *$t_2 < t_1$ but $t_3 > t_1$.*

Let

$$A := \{z \in [N-2] : \gamma_{t_1+1}(z) = 2\}$$
$$B := \{z \in [D] : \gamma_{t_1+1}(z) = 3\}$$

Note that $A \dot{\cup} B = [N-2] - \{D\}$, in particular $|A \cup B| = N - 3$.

Let us look at the path $\gamma|_{[t_2+1,t_1]}$. If we go backwards from time t_1 to time t_2+1, all disks on peg 2, except $N-2$ (in other words, the disks in A), will have to move to make room for the move of the disk $N-2$ at time t_2+1. Hence by restricting to the first $N-2$ disks and then applying Theorem 128, we get $\ell(\gamma|_{[t_2+1,t_1]}) \geq \Psi(A)$. Similarly, by restricting to the disks in $[D]$, we get $\ell(\gamma|_{[t_1+1,t_3]}) \geq \Psi(B)$. Adding the 3 moves of the disks $N-1, N-2$ and D and using Lemma 131 with $|A \cup B| = N-3$ we get

$$\ell(\gamma) \geq 3 + \Psi(A) + \Psi(B) \geq 3 + \frac{\Phi(4, N) - 5}{4}.$$

Case 2. *$t_2 > t_1$ but $t_3 < t_1$.*

This case follows from the previous one by replacing γ with γ^*.

Case 3. *$t_2, t_3 > t_1$ or $t_2, t_3 < t_1$.*

By replacing γ with γ^* if necessary, we may suppose that $t_2, t_3 > t_1$. As the disk $N-2$ does not move in the time interval $[t_1, t_2]$, we have $\gamma_{t_2}(N-2) = 2$.

We shall consider two further subcases.

Case 3.1. *$t_3 < t_2$.*

Then $\gamma|_{[t_1+1,t_2]}$ is an essential path when restricted to the moves of the first $N-2$ disks. Indeed, all disks on peg 3 must move, because D moves, and all disks on peg 2 move, to make room for the move of the disk $N-2$ at time t_2. Hence by Lemmas 133 and 137 applied to $\mathbf{u} := \gamma(t_1+1)|_{[N-2]}$ and $\mathbf{v} := \gamma(t_2)|_{[N-2]}$, we get that

$$\ell(\gamma|_{[t_1+1,t_2]}) \geq 1 + \frac{\Phi(4, N) - 5}{4}.$$

Adding the further 2 moves of the disks $N-1$ and $N-2$ gives the result.

Case 3.2. *$t_3 > t_2$.*

If $\gamma_{t_2+1}(N-2) = 3$ then the disk D moves at least once in the time interval $[t_1+1, t_2]$ and we may apply the previous subcase.

Therefore we may assume that $\gamma_{t_2+1}(N-2) \in \{0, 1\}$ and the disk D does not move in the time interval $[t_1+1, t_2]$. As pegs 0 and 1 play a symmetric role in what follows, we may further assume that $\gamma_{t_2+1}(N-2) = 0$. Set

$$A := \{z \in [N-2] : \gamma_{t_2+1}(z) = 1\}$$
$$B := \{z \in [D] : \gamma_{t_2+1}(z) = 3\}$$

As D does not move in the time interval $[t_1+1, t_2]$, D is the largest disk on peg 3 at time t_2+1. Consequently $A\dot{\cup}B = [N-2] - \{D\}$, hence $|A \cup B| = N-3$.

By Theorem 128 and the fact that pegs 0 and 1 are empty in $\gamma(t_1+1)|_{[N-2]}$, we have $\ell(\gamma|_{[t_1+1,t_2]}) \geq \Psi(A)$. Similarly, by restricting to the moves of the disks in $[D]$, we see that $\ell(\gamma_{[t_2+1,t_3]}) \geq \Psi(B)$. Therefore by adding the 3 moves of the disks $N-1, N-2$ and D and using Lemma 131 with $|A \cup B| = N-3$ we get

$$\ell(\gamma) \geq 3 + \Psi(A) + \Psi(B) \geq 3 + \frac{\Phi(4,N)-5}{4}.$$

This completes the proof of (5.6).

We will now show that the bound can be achieved. Let $a, b \geq 0$ such that $a+b = N-3$. Consider a configuration $\mathbf{u}_{a,b}$ with the disk $N-1$ on peg 2, the disk $N-2$ on peg 1 and the disk $N-3$ on peg 0. We put disks $N-3-b, N-2-b, \ldots, N-4$ on peg 0, and distribute the remaining a disks on pegs 0 and 1 in such a way that they form a midpoint configuration on 4 pegs.

Then we can first move the disk $N-1$ to peg 3, followed by the disks $0, 1, \ldots, a-1$ to the same peg in at most $\frac{\Phi(4,a+1)-1}{2}$ moves.

Afterwards we move the disk $N-2$ to peg 2 (the peg is now free, and there are no more disks on top of the disk $N-2$). We further move the disks $N-3-b, N-2-b, \ldots, N-4$ to peg 2 in $2^b - 1$ moves.

Finally, we move the disk $N-3$ to peg 1.

Let $\mathbf{v}_{a,b}$ be the resulting configuration. We have just constructed an essential path $\gamma_{a,b}$ between $\mathbf{u}_{a,b}$ and $\mathbf{v}_{a,b}$ with

$$\ell(\gamma_{a,b}) \leq 2 + \frac{\Phi(4,a+1)-1}{2} + 2^b = 2 + \frac{\Phi(4,a+1)+\Phi(3,b+1)}{2}.$$

Minimizing over all choices of a and b yields an essential path γ of length at most

$$\begin{aligned} 2 + \min_{\substack{a+b=N-3 \\ a,b\geq 0}} \frac{\Phi(4,a+1)+\Phi(3,b+1)}{2} &= 2 + \min_{\substack{a+b=N-1 \\ a,b\geq 1}} \frac{\Phi(4,a)+\Phi(3,b)}{2} \\ &= 2 + \frac{\Phi(4,N)-1}{4}, \quad \text{by Lemma 129,} \\ &= 3 + \frac{\Phi(4,N)-5}{4}. \end{aligned}$$

□

A similar argument as above shows that $\Gamma(p,N) \leq p-1+\frac{\Phi(p,N)-(2(p-2)+1)}{4}$, for all $p \geq 3$ and $N \geq p-1$.

5.3 Proof of Theorem 42

Let us recall the statement of Theorem 42. Given $p \geq 4$ and $N \geq 1$, we write

$$N - 1 = \Delta_p m + \Delta_{p-1} t + r, \quad t \leq m,\ 0 \leq r < \Delta_{p-2}(t+1). \tag{5.7}$$

Theorem 42 then states that $H(p, N) \geq (m + t)2^{m-2(p-2)}$.

Let us first show that this decomposition of $N - 1$ exists.

Lemma 138. *For every $p \geq 3$, any natural number N has a unique decomposition $N = \Delta_p m_p + \Delta_{p-1} m_{p-1} + \ldots + \Delta_3 m_3$ with $m_p \geq m_{p-1} \geq \ldots \geq m_3 \geq 0$.*

Proof. We prove the claim by induction, first after $p \geq 3$, and then after $N \geq 0$.

If $p = 3$, then as $\Delta_3 m_3 = m_3$ for any $m_3 \geq 0$ we must have $m_3 = N$. This shows that the desired decomposition of N exists and is unique.

If $N = 0$, then the only valid choice is $m_p = m_{p-1} = \ldots = m_3 = 0$.

So assume $p \geq 4$, $N \geq 1$ and the claim holds for smaller values of p or N.

Note that for any $m_p \geq \ldots \geq m_3 \geq 0$ we have

$$\Delta_p m_p + \Delta_{p-1} m_{p-1} + \ldots + \Delta_3 m_3 \leq \sum_{i=3}^{p} \Delta_i m_p = \Delta_p(m_p + 1) - 1.$$

Thus in any decomposition of N of this form we have $\Delta_p m_p \leq N < \Delta_p(m_p + 1)$ and hence $m_p = \nabla_p N$. With this choice of m_p, by induction there exists a unique decomposition

$$N - \Delta_p m = \Delta_{p-1} m_{p-1} + \ldots + \Delta_3 m_3$$

with $m_{p-1} \geq \ldots m_3 \geq 0$.

On the other hand, $N - \Delta_p m_p < \Delta_p(m_p+1) - \Delta_p m_p = \Delta_{p-1}(m_p+1)$. Hence $m_{p-1} \leq m_p$. This shows that the obtained decomposition of N satisfies $m_p \geq m_{p-1} \geq \ldots \geq m_3 \geq 0$, and moreover, is unique. □

Now given $p \geq 4$ and $N \geq 1$, let $N - 1 = \Delta_p m_p + \Delta_{p-1} m_{p-1} + \ldots + \Delta_3 m_3$ be the decomposition guaranteed by Lemma 138. Taking $m = m_p, t = m_{p-1}$ and $r = \Delta_{p-2} m_{p-2} + \ldots + \Delta_3 m_3$ gives (5.7).

Proof of Theorem 42. We prove the stronger statement

$$\Gamma(p, N) \geq (m + t)2^{m-2(p-2)}$$

by induction, first after p, and then after N. The theorem then follows from the fact that $H(p, N) \geq \Gamma(p, N)$.

If $p = 4$, the claim reduces to the inequality

$$\Gamma(4, N) \geq (m + t)2^{m-4}.$$

But by Theorem 44,

$$\Gamma(4,N)=\begin{cases} N, & \text{if } N\le 2,\\ 2+(m+t)2^{m-2}, & \text{otherwise.}\end{cases}$$

As $0=\Delta_4 0, 1=\Delta_4 1$ and $2=\Delta_4 1+1$, the claim holds in this case. So assume $p\ge 5$.

For $m\le p-2$ the claim reduces to the inequality

$$\Gamma(p,N)\ge \frac{m+t}{2^{2(p-2)-m}}.$$

But $\frac{m+t}{2^{2(p-2)-m}}\le \frac{2(p-2)}{2^{p-2}}\le 1$, as $2^x\ge 2x$ holds for all $x\ge 1$. Hence in this case the claim is trivially true. So suppose $m\ge p-1$. Then we have two cases.

Case 1. $t+1\le m-1$.

Then $\Gamma(p,\Delta_p m+\Delta_{p-1}t+r+1)\ge \Gamma(p,\Delta_p m+\Delta_{p-1}t+1)$. But

$$\Delta_p m+\Delta_{p-1}t+1=\Delta_p(m-1)+\Delta_{p-1}(m-1)+\Delta_{p-2}m+\Delta_{p-1}(t+1)-\Delta_{p-2}(t+1)+1.$$

Furthermore $\Delta_{p-2}m\ge \Delta_{p-2}(m-1)+1$, with equality if $p=5$. Also by our assumption $t+1\le m-1$ we have $\Delta_{p-2}(m-1)-\Delta_{p-2}(t+1)\ge 0$. Hence

$$\Delta_p m+\Delta_{p-1}t+1\ge \Delta_p(m-1)+\Delta_{p-1}(m-1)+\Delta_{p-1}(t+1)+2.$$

So by Lemma 132,

$$\begin{aligned}\Gamma(p,\Delta_p m+\Delta_{p-1}t+1)\ge 2\min\{&\Gamma(p,\Delta_p(m-1)+\Delta_{p-1}(t+1)+1),\\ &\Gamma(p-1,\Delta_{p-1}(m-1)+1)\}.\end{aligned}$$

By induction,

$$\Gamma(p,\Delta_p(m-1)+\Delta_{p-1}(t+1)+1)\ge (m-1+t+1)2^{m-1-2(p-2)},$$

and

$$\begin{aligned}\Gamma(p-1,\Delta_{p-1}(m-1)+1)&\ge (m-1)2^{m-1-2(p-1-2)}\\ &\ge 4(m-1)2^{m-1-2(p-2)}.\end{aligned}$$

As $4(m-1)=m+3m-4\ge m+t$, it follows that

$$\Gamma(p,\Delta_p m+\Delta_{p-1}t+r+1)\ge (m+t)2^{m-2(p-2)}.$$

Case 2. $m\ge t\ge m-1$.

Then $t \geq m-1 \geq p-2 > 0$. So

$$\Delta_p m + \Delta_{p-1} t + 1 \geq \Delta_p m + 1 + \Delta_{p-1}(t-1) + \Delta_{p-2}(t-1) + 1.$$

Hence by Lemma 132,

$$\Gamma(p, \Delta_p m + \Delta_{p-1} t + r + 1) \geq 2 \min\{\Gamma(p, \Delta_p m + 1), \\ \Gamma(p-1, \Delta_{p-1}(t-1) + \Delta_{p-2}(t-1) + 1)\}.$$

By induction,

$$\Gamma(p, \Delta_p m + 1) \geq m 2^{m-2(p-2)} = 2m 2^{m-1-2(p-2)} \geq (m+t) 2^{m-1-2(p-2)}.$$

Also,

$$\Gamma(p-1, \Delta_{p-1}(t-1) + \Delta_{p-2}(t-1) + 1) \geq 2(t-1) 2^{t-1-2(p-2)+2} \geq 4(t-1) 2^{m-1-2(p-2)}.$$

As $4t - 4 = t + 3t - 4 \geq t + 3m - 7 = m + t + 2m - 7$ and $m \geq p - 1 \geq 4 > \frac{7}{2}$, we get

$$\Gamma(p-1, \Delta_{p-1}(t-1) + \Delta_{p-2}(t-1) + 1) \geq (m+t) 2^{m-1-2(p-2)}.$$

Therefore

$$\Gamma(p, \Delta_p m + \Delta_{p-1} t + r + 1) \geq (m+t) 2^{m-2(p-2)}.$$

□

Appendices

Appendix A

Subresultant theory

We collect here several known facts about resultants and subresultants. Some of the theorems below are used in Chapter 2.

A.1 The resultant

Suppose R is an integral domain. Let $f, g \in R[x]$ be non-zero polynomials and suppose $f = a_p x^p + \ldots + a_0, g = b_q x^q + \ldots + b_0$ with $a_p, b_q \neq 0$. The *Sylvester matrix* of f and g is the $(p+q) \times (p+q)$ matrix

$$S_{f,g} := \begin{pmatrix} a_p & \ldots & a_0 & & \\ & \ddots & & \ddots & \\ & & a_p & \ldots & a_0 \\ b_q & \ldots & b_0 & & \\ & \ddots & & \ddots & \\ & & b_q & \ldots & b_0 \end{pmatrix},$$

where the first q lines are formed by shifting the first row to the right, and the last p lines are formed by shifting the $(q+1)$th row to the right. If $p = q = 0$ we define $S_{f,g} = (1)$. The *resultant* of f and g, denoted by $\mathrm{res}(f, g)$, is the determinant of $S_{f,g}$.[1] We also define $\mathrm{res}(0, h) = \mathrm{res}(h, 0) = 0$, for any polynomial h.[2]

We note some special cases.

If $p = q = 0$ then $\mathrm{res}(f, g) = 1$.

If $p = 0, q \neq 0$ then $\mathrm{res}(f, g) = a_p^q$, and similarly if $p \neq 0, q = 0$ then $\mathrm{res}(f, g) = b_q^p$.

We first have the following result, perhaps the most common application of resultants.

Theorem A.1. *Suppose R is unique factorization domain (UFD). Then* $\gcd(f, g)$ *is non-constant iff* $\mathrm{res}(f, g) = 0$.

[1] Some authors prefer to define the Sylvester matrix and the resultant $\mathrm{res}_{p,q}$ for any $p \geq \deg(f)$ and $q \geq \deg(g)$. If both inequalities are strict, then $\mathrm{res}_{p,q}(f, g) = 0$, while if just one inequality is strict, $\mathrm{res}_{p,q}$ and our definition of res differ by only a constant.

[2] We shall also use the following conventions: $\deg(0) = -\infty$ and $\gcd(h, 0) = gcd(0, h) = h$ for any polynomial h.

Proof. The proof given here is taken from [120]. As f and g are non-zero, we may clearly assume that $p, q > 0$, for otherwise the theorem follows from the previous remarks. Note that $R[x]$ is a unique factorization domain. We now have the following claim.

Claim 11. $\gcd(f, g)$ *is non-constant iff there exists polynomials u and v of degree less than p, respectively q, not both vanishing, such that $vf + ug = 0$.*

Proof. If u and v exist, then any prime factor of f must occur in the factorization of ug. However, $\deg(u) < \deg(f)$, so at least one of these prime factors divides g. Hence $\gcd(f, g)$ is non-constant.

Conversely, if $h := \gcd(f, g)$, then we can write $f = wh, g = vh$ and clearly $vf - wg = 0$. As h is non-constant, $-w$ and v have degree less than p, respectively q, and they are non-zero. □

Now write

$$\begin{aligned} u &= u_{p-1}x^{p-1} + \ldots + u_0, \\ v &= v_{q-1}x^{q-1} + \ldots + v_0, \end{aligned}$$

where the coefficients u_i, v_j are to be determined. The identity $vf + ug = 0$ reduces to the system of linear equations

$$\begin{aligned} v_{q-1}a_p + u_{p-1}b_q &= 0 \\ v_{q-2}a_p + v_{q-1}a_{p-1} + u_{p-2}b_q + u_{p-1}b_{q-1} &= 0 \\ &\vdots \\ v_0a_0 + u_0b_0 &= 0. \end{aligned}$$

The system has a non-trivial solution iff the determinant of the matrix

$$\begin{pmatrix} a_p & 0 & 0 & \ldots & 0 & b_q & 0 & 0 & \ldots & 0 \\ a_{p-1} & a_p & 0 & \ldots & 0 & b_{q-1} & b_q & 0 & \ldots & 0 \\ & \vdots & & & & & \vdots & & & \\ 0 & 0 & 0 & \ldots & a_0 & 0 & 0 & 0 & \ldots & b_0 \end{pmatrix}$$

is 0. This indeed holds in any integral domain, and can be seen as follows. We first reduce the matrix to upper-triangular form using only additions and multiplications. If the determinant vanishes, we may assume that the first r diagonal entries are non-zero, and the remaining $p+q-r$ are 0. We set the free $p+q-r$ variables to equal the product of the first r diagonal entries, which is non-zero, as we work in an integral domain. Then we can express the first r variables in terms of the last $p+q-r$ variables, and division will work by construction. So we obtain a non-trivial solution. Conversely, if the determinant is non-zero, the only solution is the trivial one.

However, the matrix constructed above is the same as $S^{tr}_{f,g}$, hence u and v exists iff $\operatorname{res}(f, g) = 0$. By Claim 11, this proves the theorem. □

We shall later see a different proof of this theorem. We now have the following.

Lemma A.2. *We always have* $\mathrm{res}(f, g) = (-1)^{pq}\mathrm{res}(g, f)$.

Proof. The matrix $S_{g,f}$ can be obtained from $S_{f,g}$ in the following way. We interchange the $(q+1)$th row with the previous q rows one by one until it becomes the first row. This changes the determinant by $(-1)^q$. We proceed similarly to make the $(q+2)$th row become the second row, and so on. The total change of the determinant is $(-1)^{pq}$, which proves the lemma. □

Lemma A.3. *Let* $h \in R[x]$ *such that* $k := \max\{\deg(f + hg), 0\}$ *verifies* $k \le p$*. Then we have* $\mathrm{res}(f, g) = (-1)^{(p-k)(q+2)} b_q^{p-k}\mathrm{res}(f + hg, g)$*, unless* $q = 0$ *and* $f + hg = 0$*.*

Proof. If $q > p$ then necessarily $h = 0$, otherwise $k > p$. Then $f + hg = f$ and there is nothing to prove.

So we may suppose that $q \le p$. Also if $f + hg = 0$ the claim follows by definition and Theorem A.1. So assume $f + hg \ne 0$ and $h = c_t x^t + \ldots + c_0$.

If $q = 0$ then $\mathrm{res}(f, g) = b_q^p$ and $\mathrm{res}(f + hg, g) = b_q^k$. Hence

$$\mathrm{res}(f, g) = (-1)^{2(p-k)} b_q^{p-k}\mathrm{res}(f + hg, g).$$

So we may assume that $q > 0$.

We modify $S_{f,g}$ by row operations. We add to the ith row the row $p + i - j$ multiplied by c_j, for any $1 \le i \le q, 0 \le j \le t$. This does not change the determinant. If $k = p$, we have now obtained the matrix $S_{f+hg,g}$ and there is nothing more to be shown. Otherwise, select in each of the first $p - k$ columns the entry b_q. If we delete the rows and columns containing these entries, what remains is exactly $S_{f+hg,g}$. Then by definition of the determinant, we are done. □

Theorem A.4. *Suppose* R *is a field and* K *is its algebraic closure.*

(i) If f *has roots* $\alpha_1, \ldots, \alpha_p$ *in* K *then*

$$\mathrm{res}(f, g) = a_p^q \prod_{i=1}^{p} g(\alpha_i). \tag{A.1}$$

(ii) If g *has roots* $\beta_1, \ldots, \beta_q$ *in* K *then*

$$\mathrm{res}(f, g) = (-1)^{pq} b_q^p \prod_{j=1}^{q} f(\beta_j). \tag{A.2}$$

(iii) Under the assumptions of (i) and (ii) we also have

$$\mathrm{res}(f, g) = a_p^q b_q^p \prod_{i=1}^{p}\prod_{j=1}^{q}(\alpha_i - \beta_j). \tag{A.3}$$

Proof. We know that

$$f = a_p(x-\alpha_1)\dots(x-\alpha_p),$$
$$g = b_q(x-\beta_1)\dots(x-\beta_q),$$

and hence expanding the expressions in (i), (ii) and (iii) we see that they are all equivalent, even when $p=0$ or $q=0$. We now prove the theorem by induction on $p+q\geq 0$.

If $p=0$ or $q=0$, the claim follows by definition.

So assume $p>0$ and $q>0$. By Lemma A.2, we may assume that $p\geq q$. Hence $f=ug+r$ with $\deg(r)<\deg(f)$.

If $r=0$, then $\operatorname{res}(f,g)=0$ by Theorem A.1. But the expression in (i) also vanishes, hence the claim holds in this case.

So we may assume that $r\neq 0$. Let $k:=\deg(r)$. By Lemma A.3 we have $\operatorname{res}(f,g) = (-1)^{(p-k)(q+2)}b_q^{p-k}\operatorname{res}(r,g)$. The induction hypothesis and (ii) imply

$$\operatorname{res}(r,g) = (-1)^{kq}b_q^k\prod_{j=1}^{q} r(\beta_j).$$

But $f(\beta_j)=u(\beta_j)g(\beta_j)+r(\beta_j)=r(\beta_j)$, hence

$$\operatorname{res}(f,g) = (-1)^{pq}b_q^p\prod_{j=1}^{q} r(\beta_j) = (-1)^{pq}b_q^p\prod_{j=1}^{q} f(\beta_j).$$

This proves the theorem. □

We now give another proof of Theorem A.1.

Proof of Theorem A.1. We may assume w.l.o.g. that f and g are non-constant.

Let S be the quotient field of R and K its algebraic closure. As the greatest common divisor of f and g can be computed in R up to a constant by the Euclidean algorithm, without using divisions, it is non-constant iff $gcd_S(f,g)$ is non-constant. But the later holds iff f and g have a common root in K, which by Theorem A.4, (iii), holds iff $\operatorname{res}(f,g)=0$. This proves the theorem. □

We will now consider the problem of several polynomials in more than one variable.

Let $f_1,\dots,f_m\in R[x_1,\dots,x_n]$. Let $y_3,\dots,y_m$ be new indeterminates and define $R' := R[x_2,\dots,x_n,y_3,\dots,y_m]$. Let F_1,F_2 be polynomials in $R'[x_1]$ defined as follows:

$$F_1 := f_1$$
$$F_2 := f_2 + y_3f_3 + \ldots + y_mf_m.$$

We define the resultant of the polynomials $f_1,\dots,f_m$ in terms of x_1, denoted by $\operatorname{res}_{x_1}(f_1,\dots,f_m)$, as the resultant of F_1 and F_2. Note that this is a polynomial in $x_2,\dots,x_n$ and $y_3,\dots,y_m$.

We first have a lemma.

Lemma A.5. *Suppose R is a UFD. Then* $\gcd(f_1,\dots,f_m)=\gcd(F_1,F_2)$.

Proof. By assumption $R[x_1,\ldots,x_n]$ and R' are also UFD. Now if $g := \gcd(f_1,\ldots,f_m)$ and $g' := \gcd(F_1,F_2)$ then $g|g'$, as $g|F_1$ and $g|F_2$. Also $g' \in R[x_1,\ldots,x_n]$, because $g'|f_1$ and R' is a UFD. Giving values $y_i = 0$ we see that $g'|f_2$. Also if we let $y_j = 1$ and $y_i = 0, i \neq j$, we see that $g'|f_2 + f_j$. Hence $g'|f_j, 2 \leq j \leq m$. Then $g'|g$ and the claim follows. □

Theorem A.6. *Assume R is a field and let K be its algebraic closure. Suppose that the leading coefficient of x_1 in f_1 does not vanish, for any choice of $x_2,\ldots,x_n$ in K. Let $(a_2,\ldots,a_n) \in K^{n-1}$. Then there exists an $a_1 \in K$ such that $(a_1,\ldots,a_n)$ is a common zero for $f_1,\ldots,f_m$ iff $\operatorname{res}_{x_1}(f_1,\ldots,f_m)(a_2,\ldots,a_n) = 0$.*

Proof. We replace x_i by $a_i, 2 \leq i \leq n$, in $f_1,\ldots,f_m$. Then the degree of F_1 stays the same, but the degree of F_2 may decrease.

If $F_2 = 0$ then a_1 can be taken to be any root of $f_1(x_1,a_2,\ldots,a_m)$ and the assumption trivially holds.

So assume $F_2 \neq 0$. By definition of the Sylvester matrix we know that

$$\operatorname{res}_{x_1}(f_1,\ldots,f_m)(a_2,\ldots,a_n) = c\operatorname{res}_{x_1}(f_1(a_2,\ldots,a_n),\ldots,f_m(a_2,\ldots,a_n)),$$

where $0 \neq c \in K$ depends only on the leading coefficient of x_1 in $f_1(a_2,\ldots,a_n)$. Thus we may suppose w.l.o.g. that the degree of F_2 also stays the same. By replacing R with K, we can also suppose w.l.o.g. that $n = 1$.

By Lemma A.5, $\gcd(f_1,\ldots,f_m) = \gcd(F_1,F_2)$, and hence a_1 exists iff $\gcd(F_1,F_2)$ is non-constant. But by Theorem A.1 this happens iff $\operatorname{res}(F_1,F_2) = \operatorname{res}_{x_1}(f_1,\ldots,f_m)$ is zero, hence the claim holds. □

The presentation of the resultant in this section was vaguely inspired by an unpublished set of notes of Svante Janson [68].

A.2 Polynomial remainder sequences

Let R be an integral domain and $f,g \in R[x]$.

We say that that f is similar to g ($f \sim g$) iff there exists non-zero $a,b \in R$ such that $af = bg$. Clearly $\sim$ is an equivalence relation. For $g \neq 0$, we also define the *pseudo-quotient* $u := \operatorname{pquo}(f,g)$ and the *pseudo-remainder* $r := \operatorname{prem}(f,g)$ by the relation $b_q^{k+1}f - ug = r$, where $\deg(r) < \deg(g)$, b_q is the leading coefficient of g and $k := \max\{\deg(f) - \deg(g), -1\}$.

Lemma A.7. *The pseudo-quotient and the pseudo-remainder exist and are uniquely defined.*

Proof. We prove by induction on $t \geq -1$ that for any two polynomials f and $g \neq 0$ with $t \geq k := \max\{\deg(f) - \deg(g), -1\}$, the relation $b_q^{t+1}f - ug = r$, where $\deg(r) < \deg(g)$ and b_q is the leading coefficient of g, uniquely defines u and r.

Suppose $f = a_px^p + \ldots + a_0$ and $g = b_qx^q + \ldots + b_0$.

If $t = -1$ or $p < q$ then u must be 0. So $r = b_q^{t+1}f$ and we are done.

Now assume $t \geq 0$ and $p \geq q$. Then $u = u_kx^k + \ldots + u_0$. We must have $u_k = b_q^t a_p$. Then set $f' := b_qf - a_px^kg$. We have $\deg(f') < \deg(f)$ and $b_q^{t+1}f - ug = b_q^t f' - u'g = r$, where $u' = u_{k-1}x^{k-1} + \ldots + u_0$. Clearly $t - 1 \geq \max\{\deg(f') - \deg(g), -1\}$ and so by the induction hypothesis, u' and r are uniquely defined. This means u is also uniquely defined, which proves the claim. □

A sequence of non-zero polynomials $f_1, f_2, \ldots, f_k$ with $f_i \sim \mathrm{prem}(f_{i-2}, f_{i-1}), 3 \leq i \leq k$, and $\mathrm{prem}(f_{k-1}, f_k) = 0$, is called a *polynomial remainder sequence*. Then by definition, there exists $\alpha_i, \beta_i \in R$ and $u_i \sim \mathrm{pquo}(f_{i-2}, f_{i-1})$ such that

$$\alpha_i f_{i-2} - u_i f_{i-1} = \beta_i f_i, \quad \deg(f_i) < \deg(f_{i-1}), \quad 3 \leq i \leq k.$$

Furthermore, if R is a UFD, $\gcd(f_1, f_2) \sim f_k$.

A.3 Subresultants

Let R be an integral domain and $f, g \in R[x]$ non-zero.

Set $p := \deg(f)$ and $q := \deg(g)$ and assume $f = a_p x^p + \ldots + a_0, g = b_q x^q + \ldots + b_0$.

The *subresultant sequence for f and g* is a list of polynomials

$$S_i(f, g) := \sum_{j=0}^{i} s_{ij}(f, g) x^j, \quad 0 \leq i \leq \min\{p, q\},$$

where $s_{ij}(f, g)$ is the determinant of the matrix $M_{ij}(f, g)$ built with rows $1, \ldots, q - i$ and $q + 1, \ldots, q + p - i$ of $S_{f,g}$, and columns $1, 2, \ldots, p + q - 2i - 1, p + q - i - j$ of $S_{f,g}$. When $p = q \neq 0$ we set $S_q(f, g) = g$ and define s_{qj} in the obvious way. For $p = q = 0$ we set $S_0(f, g) = 1$ and we also define $S_0(0, f) = S_0(f, 0) = 0$. This is supported by the following observation.

Lemma A.8. *We have $S_0(f, g) = s_{00}(f, g) = \mathrm{res}(f, g)$. If $p > q \geq 0$ then $S_q(f, g) = b_q^{p-q-1} g$, while if $q > p \geq 0$ we have $S_p(f, g) = a_p^{q-p-1} f$. In general we have $S_i(g, f) = (-1)^{(p-i)(q-i)} S_i(f, g)$ for any $0 \leq i \leq \min\{p, q\}$ with $i < \min\{p, q\}$ or $p \neq q$.*

Proof. The first two assertions follow by definition. We prove the third.

Let $0 \leq i \leq \min\{p, q\}$ such that either $i < \min\{p, q\}$ or $p \neq q$. Then for any $0 \leq j \leq i$, we can interchange the first $(q - i)$ rows with the last $(p - i)$ rows of the matrix $M_{ij}(f, g)$ to obtain the matrix $M_{ij}(g, f)$. As the change in the determinant is $(-1)^{(p-i)(q-i)}$, this proves the claim. □

Lemma A.8 should be compared with Lemma A.2. We now prove a corresponding result to Lemma A.3.

Lemma A.9. *Assume $q > 0$ and let $h \in R[x]$ such that $k := \deg(f + hg)$ verifies $k \leq p$.*

If $f + gh = 0$ then $S_i(f, g) = 0, 0 \leq i < q$.

If $k \geq 0$ and $t := \max\{\min\{p-k, p+q-2i-1\}, 0\}$ then we have $S_i(f, g) = (-1)^{t(q-i+2)} b_q^t S_i(f + hg, g), 0 \leq i \leq \min\{k, q\}$.

Proof. We may suppose $q \leq p$, otherwise there is nothing to prove. Suppose $h = c_r x^r + \ldots + c_0$.

If $p = i = q = k$, then as $q \neq 0$, by definition $S_q(f,g) = S_q(f+hg,g) = g$ and the claim holds.

If $p > i = q = k$, then also by definition $S_q(f+hg,g) = g$ and $p-k = p-q > p+q-2i-1 = p-q-1$. By Lemma A.8, $S_q(f,g) = b_q^{p-q-1} g$, as desired.

So we may assume that none of these cases holds and let $0 \leq j \leq i$.

We modify $M_{ij}(f,g)$ by row operations. We add to the ℓ_1th row the rows $p + \ell_1 - \ell_2 - i$ multiplied by c_{ℓ_2}, for any $1 \leq \ell_1 \leq q-i, 0 \leq \ell_2 \leq r$. This does not change the determinant. If $k = p$, we have now obtained the matrix $M_{ij}(f+hg,g)$ and there is nothing more to be shown. Also if $f + hg = 0$ and $i < q$, the matrix $M_{ij}(f+hg,g)$ has a 0 row, and consequently $s_{ij}(f,g) = 0$.

If $p - k \geq p + q - i - j$ then $k \leq i + j - q \leq 2i - q$, or $2i \geq k + q$. Then $k = q = i$. But we have assumed we are not in this situation, hence $p - k < p + q - i - j$. Therefore we may select in each of the first $\min\{p-k, p+q-2i-1\}$ columns of $M_{ij}(f,g)$ the entry containing b_q. If we delete the rows and columns containing these entries, what remains is exactly $M_{ij}(f+hg,g)$. Then by definition of the determinant, we are done. $\square$

Note that by taking $i = 0$ in Lemma A.9 one obtains Lemma A.3.

The following theorem is the main application of subresultants.

Theorem A.10. *Suppose R is a UFD. If f and g are non-zero, and $k \geq 0$ is minimal such that $s_{kk}(f,g) \neq 0$ then there exists non-zero $u, v \in R$ such that $u \gcd(f,g) = v S_k(f,g)$.*

Proof. By Lemma A.8, k is well-defined. We now prove the theorem by induction on $p+q$.

By definition and Lemma A.8 the claim holds if $p = 0$ or $q = 0$. Hence we may assume that $p, q > 0$. We may further assume that $p \geq q$, as otherwise by Lemma A.8 we may interchange f with g.

Let $r := \operatorname{prem}(f,g)$. By using the fact that $S_i(cf,g) \sim S_i(f,g)$ and $\gcd(cf,g) \sim \gcd(f,g)$ for any non-zero $c \in R$, we may assume (by multiplying f by a constant if necessary) that $f = ug + r$. Set $t := \deg(r)$ and note that $t + q < p + q$.

If $r = 0$ then $g|f$ and by Lemma A.9, $k = q$. Then by definition the theorem holds.

Consequently we may assume that $r \neq 0$. If $k \leq t$, by the induction hypothesis and Lemma A.9, $S_k(f,g) \sim S_k(r,g) \sim \gcd(r,g) \sim \gcd(f,g)$, and hence the claim holds in this case. But if $k > t$, then by assumption and Lemma A.9, $0 = S_t(f,g) \sim S_t(r,g) \sim r$, and so $r = 0$, a contradiction. $\square$

Note that Theorem A.10 implies Theorem A.1, as $\gcd(f,g)$ is non-constant iff $k > 0$, or what is the same, $S_0(f,g) = \operatorname{res}(f,g) = 0$.

More information about subresultants can be found in [8].

Appendix B

A measure-theoretic lemma

Let $A \subset \mathbb{R}$ be a Borel set. Recall that the Lebesgue measure of A is defined as

$$\mu(A) = \inf \left\{ \sum_{k=1}^{\infty} \ell(I_k) : \{I_k\}_{k\in\mathbb{N}} \text{ is any sequence of disjoint intervals with } A \subset \bigcup_k I_k \right\},$$

where $\ell(I_k)$ is the length of the interval I_k.

Lemma B.11. *Let $h : J_1 \to J_2$ be a homeomorphism between two intervals $J_1, J_2 \subseteq \mathbb{R}$. If $A \subseteq J_1$ is Borel, then so is $h(A)$. Furthermore if A has positive Lebesgue measure and the first-order derivative of h^{-1} exists and is bounded, then $h(A)$ has positive Lebesgue measure.*

Proof. As A is Borel, it can be obtained using a countable number of reunion, intersection or relative complement operations from the collection of closed-open subsets of $\mathbb{R}$. As $A = A \cap J_1$, J_1 is an interval and h is a homeomorphism, it follows that $h(A)$ can be constructed in the same manner. Thus $h(A)$ is Borel. In particular, it is Lebesgue measurable.

Assume that the first-order derivative of h^{-1} exists and is bounded. Let $C > 0$ such that $|h^{-1}(x)| \leq C$ for all $x \in J_2$. For any $a, b \in J_2$ we have $|h^{-1}(a) - h^{-1}(b)| \leq C|a - b|$. In particular, if $I \subseteq J_2$ is any interval, then $h^{-1}(I)$ must be an interval of length at most $C\ell(I)$.

We will show that $\mu(A) \leq C\mu(h(A))$. From this it follows that if A has positive measure, then so does $h(A)$.

Let $\varepsilon > 0$ be arbitrary. Then there exists a sequence of disjoint intervals $\{I_k\}_{k\in\mathbb{N}}$ with $h(A) \subset \bigcup_k I_k$ and $\sum_{k=1}^{\infty} \ell(I_k) < \mu(h(A)) + \varepsilon$. We may assume w.l.o.g. that $I_k \subset J_2$ for all k. Then $A \subset \bigcup_k h^{-1}(I_k)$ and by our previous observation $\mu(h^{-1}(I_k)) \leq C\ell(I_k)$. Consequently $\mu(A) \leq \sum_{k=1}^{\infty} C\ell(I_k) < C(\mu(h(A)) + \varepsilon)$. Letting ε tend to zero gives the result. $\square$

Zusammenfassung

Diese Dissertation besteht aus fünf Kapiteln.

Das erste Kapitel stellt die vier Probleme vor, welche Gegenstand der Dissertation sind. Jedem der vier Probleme ist eines der folgenden Kapitel gewidmet.

Das zweite Kapitel befasst sich mit dem ersten Problem, der Existenz von partiellen Isomorphismen (bijektive Abbildungen, die nur eine endlich bestimmte Menge von algebraischen Relationen erhalten) zwischen Teilmengen von $\mathbb{F}_p$ und Teilmengen von $\mathbb{C}$. Wir zeigen, dass für jede genügend kleine Teilmenge von $\mathbb{F}_p$ eine Teilmenge von $\mathbb{C}$ existiert, die sich algebraisch ähnlich verhält. Wir gehen auf einige Anwendungen dieses Ergebnisses ein, insbesondere zeigen wir, dass für kleine Teilmengen von $\mathbb{F}_p$ der Satz von Szemerédi–Trotter mit optimaler Potenz $\frac{4}{3}$ gilt. Außerdem geben wir einige Teilantworten auf eine alte Frage von A. Rényi. Der Inhalt des vorstehenden Kapitels ist in [55] erschienen.

Das dritte Kapitel befasst sich mit dem zweiten Problem, Turán-Dichten von Hypergraphen bestimmen. Dieses Problem führte zu einem klassiches Gebiet der Graphentheorie, in welchem es noch zahlreiche ungelöste Probleme gibt. Unser Hauptergebnis sagt aus, dass eine abstrakte algebraische Struktur über der Menge der Turán Dichten existiert. Hieraus folgen auf einfache Weise einige bekannte Resultate. Die vorstehenden Ergebnisse sind in [56] erschienen.

Das vierte Kapitel befasst sich mit dem dritten Problem, der Graceful-Tree-Vermutung, eine fünfzig Jahre alte Vermutung mit wichtigen Anwendungen im Gebiet der Graphen-Zerlegungen. Unser Hauptergebnis ist eine approximative Version der Vermutung für Bäume mit beschräktem Grad. Die Ergebnisse im vorstehenden Kapitel sind in Zusammenarbeit mit Anna Adamaszek, Michał Adamaszek, Peter Allen und Jan Hladký entstanden und in [1] erschienen.

Das fünfte Kapitel befasst sich mit dem vierten Problem, einer Verallgemeinerung des Türme-von-Hanoi-Spiels von É. Lucas mit 3 Stäben auf die Situation mit $p \geq 3$ Stäben. Konkret ist die Frage nach der Anzahl von Zügen, die für ein erfolgreiches Spiel notwendig sind, eine Frage, die seit einigen Jahrzehnten offen ist. Erst kürzlich ist eine Lösung für den Fall $p = 4$ von T. Bousch erschienen []. Unser Hauptergebnis ist eine verbesserte untere Schranke für die minimale Anzahl von notwendigen Zügen für den Fall $p \geq 5$, welches in [57] erschienen ist.

Eidesstattliche Erklärung

Gemäß § 7 (4) der Promotionsordnung des Fachbereichs Mathematik und Informatik der Freien Universität Berlin versichere ich hiermit, dass ich alle Hilfsmittel und Hilfen angegeben und auf dieser Grundlage die Arbeit selbständig verfasst habe. Des Weiteren versichere ich, dass ich diese Arbeit nicht schon einmal zu einem früheren Promotionsverfahren eingereicht habe.

Berlin, 2015

Codruţ Grosu

Curriculum Vitae

Codruţ Grosu

2002–2006	C. N. Tudor Vianu, Bucharest, Romania (high school)
	Mathematics and Computer Science
2006–2010	Politehnica University of Bucharest, Romania
	Bachelor of Science
	Computer Science and Engineering
2010–2012	Freie Universität, Berlin, Germany
	Master of Science
	Member of *Berlin Mathematical School*
2012–2015	Freie Universität, Berlin, Germany
	PhD
	Member of *Berlin Mathematical School*, *Methods for Discrete Structures*

Bibliography

[1] Anna Adamaszek, Michał Adamaszek, Codruţ Grosu, Peter Allen, and Jan Hladký, *Almost all trees are almost graceful*, Preprint, 2015.

[2] R.E.L. Aldred and Brendan McKay, *Graceful and harmonious labellings of trees*, Bull. Inst. Combin. Appl. **23** (1998), 69–72.

[3] Noga Alon and Joel Spencer, *The probabilistic method*, third ed., Wiley, Hoboken, NJ, 2008.

[4] Simon Aumann, Katharina Götz, Andreas Hinz, and Ciril Petr, *The number of moves of the largest disc in shortest paths on Hanoi graphs*, Electron. J. Combin. **21** (2014), #P4.38.

[5] Rahil Baber and John Talbot, *Hypergraphs Do Jump*, Combin. Probab. Comput. **20** (2011), no. 2, 161–171.

[6] ———, *New Turán densities for* 3*-graphs*, Electronic J. Combin. **19** (2012), 21pp.

[7] ———, *A solution to the* 2/3 *conjecture*, SIAM J. Discrete Math. **28** (2014), no. 2, 756–766.

[8] Saugata Basu, Richard Pollack, and Marie-Françoise Roy, *Algorithms in real algebraic geometry (algorithms and computation in mathematics)*, Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.

[9] Jean-Claude Bermond, *Graceful graphs, radio antennae and French windmills*, Graph Theory and Combinatorics, Pitman, London, 1979, pp. 18–37.

[10] Yuri Bilu, Vsevolod F. Lev, and Imre Z. Rusza, *Rectification Principles in Additive Number Theory*, Discrete Comput. Geom. **19** (1998), no. 3, 343–353.

[11] Béla Bollobás, Imre Leader, and Claudia Malvenuto, *Daisies and Other Turán Problems*, Combin. Probab. Comput. **20** (2011), no. 5, 743–747.

[12] Julia Böttcher, Jan Hladký, Diana Piguet, and Anusch Taraz, *An approximate version of the tree packing conjecture*, arXiv:1404.0697, to appear in Israel J Math.

[13] Jean Bourgain, *On the Erdős-Volkmann and Katz-Tao ring conjectures*, Geom. Funct. Anal. **13** (2003), no. 2, 334–365.

[14] Jean Bourgain and Moubariz Z. Garaev, *On a variant of sum-product estimates and explicit exponential sum bounds in prime fields*, Math. Proc. Cambridge Philos. Soc. **146** (2009), no. 1, 1–21.

[15] Jean Bourgain, Nets Hawk Katz, and Terence Tao, *A sum-product estimate in finite fields and applications*, Geom. Funct. Anal. **14** (2004), no. 1, 27–57.

[16] Thierry Bousch, *La quatrième tour de Hanoï*, Bull. Belg. Math. Soc. Simon Stevin **21** (2014), 895–912.

[17] W. G. Brown and Miklós Simonovits, *Digraph extremal problems, hypergraph extremal problems, and the densities of graph structures*, Discrete Math. **48** (1984), 147–162.

[18] W. S. Brown and Joseph Frederick Traub, *On Euclid's Algorithm and the Theory of Subresultants*, J. ACM **18** (1971), no. 4, 505–514.

[19] Richard Carr and Cormac O'Sullivan, *On the Linear Independence of Roots*, Int. J. Number Theory **5** (2009), no. 1, 161–171.

[20] Mei-Chu Chang, *Factorization in generalized arithmetic progressions and applications to the Erdős-Szemerédi sum-product problems*, Geom. Funct. Anal. **13** (2003), no. 4, 720–736.

[21] ______, *A sum-product estimate in algebraic division algebras*, Israel J. Math. **150** (2005), no. 1, 369–380.

[22] W. C. Chen, H. I. Lü, and Y. N. Yeh, *Operations of interlaced trees and graceful trees*, Southeast Asian Bull. Math. **21** (1997), 337–348.

[23] Xiao Chen and Jian Shen, *On the Frame-Stewart conjecture about the Towers of Hanoi*, SIAM J. Comput. **33** (2004), no. 3, 584–589.

[24] Gi-Sang Cheon and Ian M. Wanless, *Some results towards the Dittert conjecture on permanents*, Linear Algebra Appl. **436** (2012), 791–801.

[25] Gi-Sang Cheon and Haeng-Won Yoon, *A note on the Dittert conjecture for permanents*, Int. Math. Forum **1** (2006), no. 39, 1943–1949.

[26] Fan Chung and Ronald Graham, *Erdős on Graphs: His Legacy of Unsolved Problems*, A. K. Peters, Wellesley, Massachusetts, 1998.

[27] George E. Collins, *Subresultants and reduced polynomial remainder sequences*, J. ACM **14** (1967), no. 1, 128–142.

[28] Don Coppersmith and James Davenport, *Polynomials whose powers are sparse*, Acta Arith. **LVIII** (1991), no. 1, 79–87.

[29] Michael Drmota and Anna Lladó, *Almost every tree with m edges decomposes $K_{2m,2m}$*, Combin. Probab. Comput. **23** (2014), no. 1, 50–65.

[30] G. A. Edgar and Chris Miller, *Borel subrings of the reals*, Proc. Amer. Math. Soc. **131** (2003), 1121–1129.

[31] Michelle Edwards and Lea Howard, *A Survey Graceful Trees*, Atlantic Electronic Journal of Math **1** (2006), 5–30.

[32] György Elekes, *On the number of sums and products*, Acta Arith. **LXXXI** (1997), no. 4, 365–367.

[33] Eric Emtander, *Betti numbers of hypergraphs*, Communications in Algebra **37** (2009), no. 5, 1545–1571.

[34] Paul Erdős, *On extremal problems of graphs and generalized graphs*, Israel J. Math. **2** (1964), 183–190.

[35] ______, *On some extremal problems on r-graphs*, Discrete Math. **1** (1971), no. 1, 1–6.

[36] ______, *On some of my conjectures in number theory and combinatorics*, Proceedings of the fourteenth Southeastern conference on combinatorics, graph theory and computing (Boca Raton), vol. 39, 1983, pp. 3–19.

[37] Paul Erdős and Miklós Simonovits, *A limit theorem in graph theory*, Stud. Sci. Math. Hungar. **1** (1966), 51–57.

[38] Paul Erdős and Arthur Harold Stone, *On the structure of linear graphs*, Bulletin of the American Mathematical Society **52** (1946), 1087–1091.

[39] Paul Erdős and Bodo Volkmann, *Additive Gruppen mit vorgegebener Hausdorffscher Dimension*, J. Reine Angew. Math. **221** (1966), 203–208.

[40] Paul Erdős and Endre Szemerédi, *On sums and products of integers*, Studies in Pure Math., Birkhäuser, Basel, 1983, pp. 213–218.

[41] Victor Falgas-Ravry and Emil Vaughan, *Applications of the semi-definite method to the Turán density problem for 3-graphs*, Combin. Probab. Comput. **22** (2013), no. 1, 21–54.

[42] Wenjie Fang, *A computational approach to the graceful tree conjecture*, Preprint (http://arxiv.org/abs/1003.3045).

[43] Kevin Ford, *Sums and Products from a Finite Set of Real Numbers*, Ramanujan J. **2** (1998), no. 1-2, 59–66.

[44] James Sutherland Frame, *Solution to Advanced Problem 3918*, Amer. Math. Monthly **48** (1941), no. 3, 216–217.

[45] Peter Frankl, Yuejian Peng, Vojtěch Rödl, and John Talbot, *A note on the jumping constant conjecture of Erdős*, J. Combin. Theory Ser. B **87** (2007), no. 2, 204–216.

[46] Peter Frankl and Vojtěch Rödl, *Hypergraphs do not jump*, Combinatorica **4** (1984), no. 2–3, 149–159.

[47] Gregory Freiman, *Foundations of a Structural Theory of Set Addition*, Kazan. Gos. Ped. Inst., Kazan, 1966, English translation: Translation of Mathematical Monographs, vol. 37, American Mathematical Society, Providence, RI, 1973.

[48] Joseph A. Gallian, *A dynamic survey of graph labeling*, Electronic J. Combin. (2014), 384p.

[49] Moubariz Z. Garaev, *An explicit sum-product estimate in* $\mathbb{F}_p$, Intern. Math. Res. Notices **2007** (2007), no. 11, 1–11.

[50] Roman Glebov, Daniel Král', and Jan Volec, *A problem of Erdős and Sós on* 3*-graphs*, Israel J. Math, to appear., 2013.

[51] Solomon Wolf Golomb, *How to number a graph*, Graph Theory and Computing (R. C. Read, ed.), Academic Press, 1972, pp. 23–37.

[52] William Timothy Gowers, *Hypergraph regularity and the multidimensional Szemerédi theorem*, Ann. of Math. **166** (2007), no. 3, 897–946.

[53] Ronald L. Graham and Neil J. A. Sloane, *On additive bases and harmonious graphs*, SIAM J. Algebr. Discr. Methods **1** (1980), no. 4, 382–404.

[54] Ben Green and Imre Z. Rusza, *Sets with small subsets and rectification*, Bull. London Math. Soc. **38** (2006), no. 1, 43–52.

[55] Codruţ Grosu, $\mathbb{F}_p$ *is locally like* $\mathbb{C}$, J. London Math. Soc. **89** (2014), no. 3, 724–744.

[56] ______, *On the algebraic and topological structure of the set of Turán densities*, Preprint, Available online at `http://arxiv.org/abs/1403.4653`, 2014.

[57] ______, *A new lower bound for the Towers of Hanoi problem*, Preprint, Available online at `http://arxiv.org/abs/1508.04272`, 2015.

[58] Andrzej Grzesik, *On the maximum number of five-cycles in a triangle-free graph*, J. Combin. Theory Ser. B **102** (2012), no. 5, 1061–1066.

[59] Leonid Gurvits, *Van der Waerden/Schrijver-Valiant like conjectures and stable (aka hyperbolic) homogeneous polynomials: One theorem for all*, Electron. J. Combin. **15** (2008), 26pp.

[60] Bruce Hajek, *A Conjectured Generalized Permanent Inequality and a Multiaccess Problem*, Open Problems in Communication and Computation (T. Cover and B. Gopinath, eds.), New York Springer-Verlag, 1987, pp. 127–129.

[61] Hamed Hatami, Jan Hladký, Daniel Král', Serguei Norine, and Alexander Razborov, *Non-three-colorable common graphs exist*, Combin. Probab. Comput. **21** (2012), no. 5, 734–742.

[62] ______, *On the Number of Pentagons in Triangle-Free Graphs*, J. Combin. Theory Ser. A **120** (2013), no. 3, 722–732.

[63] Harald Andrés Helfgott and Misha Rudnev, *An explicit incidence theorem in* $\mathbb{F}_p$, Mathematika **57** (2011), no. 1, 135–145.

[64] Robert Hermann, *Linear Systems Theory and Introductory Algebraic Geometry*, Math Sci Press, Brookline, Mass., 1974.

[65] Pavel Hrnčiar and Alfonz Haviar, *All trees of diameter five are graceful*, Discrete Math. **233** (2001), 133–150.

[66] Suk Geun Hwang, *A Note on a Conjecture on Permanents*, Linear Algebra Appl. **76** (1986), 31–44.

[67] ______, *On a conjecture of E. Dittert*, Linear Algebra Appl. **95** (1987), 161–169.

[68] Svante Janson, *Resultant and discriminant of polynomials*, Available online at `http://www2.math.uu.se/~svante/papers/sjN5.pdf`.

[69] Svante Janson, Tomasz Łuczak, and Andrzej Rucinski, *Random graphs*, Wiley-Interscience Series in Discrete Mathematics and Optimization, Wiley-Interscience, New York, 2000.

[70] Timothy G. F. Jones, *Further improvements to incidence and Beck-type bounds over prime finite fields*, Preprint, Available online at `http://arxiv.org/abs/1206.4517`, 2012.

[71] Gyula O. H. Katona, T. Nemetz, and Miklós Simonovits, *On a graph problem of Turán (In Hungarian)*, Mat. Fiz. Lapok **15** (1964), 228–238.

[72] Nets Hawk Katz and Chun-Yen Shen, *A slight improvement to Garaev's sum product estimate*, Proc. Amer. Math. Soc. **136** (2008), no. 7, 2499–2504.

[73] Peter Keevash, *Hypergraph Turán Problems*, Surveys in Combinatorics (R. Chapman, ed.), Cambridge Univ. Press, 2011, pp. 83–140.

[74] Peter Keevash and Benny Sudakov, *On a hypergraph Turán problem of Frankl*, Combinatorica **25** (2005), no. 6, 673–706.

[75] Sandi Klavžar, Uroš Milutinović, and Ciril Petr, *On the Frame-Stewart algorithm for the multi-peg Tower of Hanoi problem*, Discrete Appl. Math. **120** (2002), no. 1–3, 141–157.

[76] Sergei Vladimirovich Konyagin and Misha Rudnev, *On New Sum-Product Type Estimates*, SIAM J. Discrete Math. **27** (2013), no. 2, 973–990.

[77] János Körner and Katalina Marton, *Random Access Communication and Graph Entropy*, IEEE Trans. Inform. Theory **34** (1988), no. 2, 312–314.

[78] Daniel Král', Chun-Hung Liu, Jean-Sébastien Sereni, Peter Whalen, and Zelealem Yilma, *A new bound for the* 2/3 *conjecture*, Combin. Probab. Comput. **22** (2013), no. 3, 384–393.

[79] Teresa Krick, Luis Miguel Pardo, and Martín Sombra, *Sharp estimates for the arithmetic Nullstellensatz*, Duke Math. J. **109** (2001), no. 3, 521–598.

[80] Liangpan Li, *Slightly improved sum-product estimates in fields of prime order*, Acta Arith. **147** (2011), no. 2, 153–160.

[81] Édouard Lucas, *Récréations mathématiques, vol. iii*, Gauthier-Villars, Paris, 1893, Reprinted several times by Albert Blanchard, Paris.

[82] W. F. Lunnon, *The Reve's Puzzle*, The Computer Journal **29** (1986), 478.

[83] Colin McDiarmid, *On the method of bounded differences*, Surveys in Combinatorics, 1989 (J. Siemons, ed.), London Math. Soc. Lecture Note Ser., vol. 141, Cambridge Univ. Press, 1989, pp. 148–188.

[84] Henryk Minc, *Theory of permanents 1978-1981*, Linear and Multilinear Algebra **12** (1983), 227–263.

[85] J. W. Moon, *On the maximum degree in a random tree*, Michigan Math. J. **15** (1968), no. 4, 429–432.

[86] Rajeev Motwani and Prabhakar Raghavan, *Randomized algorithms*, Cambridge University Press, New York, 2006.

[87] Brendan Nagle, Vojtěch Rödl, and Mathias Schacht, *The counting lemma for regular k-uniform hypergraphs*, Random Struct. Algor. **28** (2006), no. 2, 113–179.

[88] Melvyn Bernard Nathanson, *On sums and products of integers*, Proc. Amer. Math. Soc. **125** (1997), no. 1, 9–16.

[89] Paul Erdős, *On the number of terms of the square of a polynomial*, Nieuw Arch. Wiskunde **23** (1949), no. 2, 63–65.

[90] A. M. Pastel and H. Raynaud, *Numerotation gracieuse des oliviers*, Colloq. Grenoble, Publications Université de Grenoble, 1978, pp. 218–223.

[91] Yuejian Peng, *Non-jumping numbers for* 4*-uniform hypergraphs*, Graphs Combin. **23** (2007), no. 1, 97–110.

[92] ———, *Using Lagrangians of hypergraphs to find non-jumping numbers II*, Discrete Math. **307** (2007), no. 14, 1754–1766.

[93] ———, *Using Lagrangians of hypergraphs to find non-jumping numbers I*, Ann. Comb. **12** (2008), 307–324.

[94] ———, *On Jumping Densities of Hypergraphs*, Graphs and Combinatorics **25** (2009), no. 5, 759–766.

[95] Yuejian Peng and Cheng Zhao, *Generating non-jumping numbers recursively*, Discrete Appl. Math. **156** (2008), no. 10, 1856–1864.

[96] Oleg Pikhurko, *On Possible Turán Densities*, Israel J. Math. **201** (2014), 415–454.

[97] ———, *The maximal length of a gap between r-graph Turán densities*, Preprint, Available online at `http://arxiv.org/abs/1504.00769`, 2015.

[98] Michael Rand, *On the Frame-Stewart Algorithm for the Tower of Hanoi*, Tech. report, Boston College, 2009, Available online at `https://www2.bc.edu/~grigsbyj/Rand_Final.pdf`.

[99] Alexander Razborov, *Flag Algebras*, Journal of Symbolic Logic **72** (2007), no. 4, 1239–1282.

[100] ———, *On 3-hypergraphs with forbidden 4-vertex configurations*, SIAM J. Disc. Math. **24** (2010), no. 3, 946–963.

[101] Alfréd Rényi, *On the minimal number of terms of the square of a polynomial*, Hungarica Acta Math. **1** (1947), 30–34, Reprinted in *Selected papers of Alfréd Rényi*, vol. 1, 42–47, edited by Paul Turán, Akadémiai Kiadó, Budapest 1976.

[102] Gerhard Ringel, *Problem 25*, Theory of Graphs and its Applications (Proc. Sympos. Smolenice 1963), Nakl. CSAV, Praha, 1964, p. 162.

[103] Vojtěch Rödl and Mathias Schacht, *Generalizations of the Removal Lemma*, Combinatorica **29** (2009), no. 4, 467–501.

[104] Vojtěch Rödl and Jozef Skokan, *Regularity lemma for k-uniform hypergraphs*, Random Struct. Algor. **25** (2004), no. 1, 1–42.

[105] ———, *Applications of the regularity lemma for uniform hypergraphs*, Random Struct. Algor. **28** (2006), no. 2, 180–194.

[106] Alexander Rosa, *On certain valuations of the vertices of a graph*, Theory of Graphs (Internat. Symposium, Rome, July 1966), Gordon and Breach, N. Y. and Dunod Paris, 1967, pp. 349–355.

[107] Misha Rudnev, *An improved sum-product inequality in fields of prime order*, Intern. Math. Res. Notices **2012** (2012), no. 16, 3693–3705.

[108] Imre Ruzsa and Endre Szemerédi, *Triple systems with no six points carrying three triangles*, Combinatorics (Keszthely, 1976), Coll. Math. Soc. J. Bolyai 18, Volume II, 1978, pp. 939–945.

[109] Andrzej Schinzel, *On the number of terms of a power of a polynomial*, Acta Arith. **XLIX** (1987), no. 1, 55–70.

[110] Andrzej Schinzel and Umberto Zannier, *On the number of terms of a power of a polynomial*, Atti Accad. Naz. Lincei Cl. Sci. Fis. Mat. Natur. Rend. Lincei Mat. Appl. **20** (2009), no. 1, 95–98.

[111] Alexander Sidorenko, *Extremal combinatorial problems in spaces with continuous measure (in Russian)*, Issledovanie Operatchiy i ASU **34** (1989), 34–40.

[112] ———, *Asymptotic Solution of the Turán Problem for Some Hypergraphs*, Graphs and Combinatorics **8** (1992), no. 2, 199–201.

[113] ———, *What we know and what we do not know about Turán numbers*, Graphs and Combinatorics **11** (1995), no. 2, 179–199.

[114] Miklós Simonovits, *A method for solving extremal problems in graph theory, stability problems*, Theory of Graphs (Proc. Colloq., Tihany, 1966), Bolyai Soc. Math. Stud., Academic Press, 1968, pp. 279–319.

[115] Richard Sinkhorn, *A problem related to the van der Waerden permanent theorem*, Linear and Multilinear Algebra **16** (1984), 167–173.

[116] József Solymosi, *On the number of sums and products*, Bull. London Math. Soc. **37** (2005), no. 4, 491–494.

[117] ———, *Bounding multiplicative energy by the sumset*, Adv. Math. **222** (2009), no. 2, 402–408.

[118] József Solymosi and Terence Tao, *An Incidence Theorem in Higher Dimensions*, Discrete Comput. Geom. **48** (2012), no. 2, 255–280.

[119] Hugo Steinhaus, *Sur le distances des points dans les ensembles de mesure positive*, Fund. Math. **1** (1920), no. 1, 93–104.

[120] Jan Stevens, *Bézout's Theorem and Inflection Points*, Available online at `http://www.math.chalmers.se/~stevens/bezout.pdf`.

[121] Bonnie M. Stewart, *Advanced Problem 3918*, Amer. Math. Monthly **46** (1939), no. 6, 363.

[122] ———, *Solution to Advanced Problem 3918*, Amer. Math. Monthly **48** (1941), no. 3, 217–219.

[123] Paul K. Stockmeyer, *Variations on the Four-Post Tower of Hanoi Puzzle*, Congr. Numer. **102** (1994), 3–12.

[124] W.-C. Stephen Suen, *A correlation inequality and a poisson limit theorem for nonoverlapping balanced subgraphs of a random graph*, Random Structures Algorithms **1** (1990), no. 2, 231–242.

[125] Mario Szegedy, *In how many steps the k peg version of the Towers of Hanoi game can be solved?*, STACS 99 (Trier), Lecture Notes in Comput. Sci., vol. 1563, Springer-Verlag, Berlin, 1999, pp. 356–361.

[126] Terence Tao, *Polynomial bounds via nonstandard analysis*, 5 July 2011, `http://terrytao.wordpress.com/2011/07/05/polynomial-bounds-via-nonstandard-analysis/`.

[127] ———, *Rectification and the Lefschetz principle*, 14 March 2013, `http://terrytao.wordpress.com/2013/03/14/rectification-and-the-lefschetz-principle/`.

[128] Terence Tao and Van H. Vu, *Additive combinatorics*, Cambridge University Press, New York, USA, 2009.

[129] Csaba D. Tóth, *The Szemerédi-Trotter Theorem in the Complex Plane*, Combinatorica **35** (2015), no. 1, 95–126.

[130] W. Verdenius, *On the number of terms of the square and the cube of polynomials*, Indag. Math. **11** (1949), 459–465.

[131] Van H. Vu, Melanie Matchett Wood, and Philip Matchett Wood, *Mapping incidences*, J. London Math. Soc. **84** (2011), no. 2, 433–445.

[132] T.-M. Wang, C.-C. Yang, L.-H. Hsu, and E. Cheng, *Infinitely many equivalent versions of the graceful tree conjecture*, Appl. Anal. Discrete Math. **9** (2015), to appear.

[133] Richard Wilson, *Decompositions of complete graphs into subgraphs isomorphic to a given graph*, Proceedings of the Fifth British Combinatorial Conference (Univ. Aberdeen, Aberdeen, 1975), Congr. Num. XV, 1976, pp. 647–659.

[134] Joshua Zahl, *A Szemerédi-Trotter type theorem in* $\mathbb{R}^4$, Discrete Comput. Geom., to appear, 2012.

[135] Bing Zhou, *A Turán-type problem on* 3*-graphs*, Ars Combin. **31** (1991), 177–181.