Universitat Politècnica de Catalunya

Programa de Doctorat de Matemàtica Aplicada

Probabilistic and Extremal Studies in Additive Combinatorics

Doctoral thesis by
Maximilian Wötzel

Thesis advisors
Juanjo Rué and Oriol Serra

Department of Mathematics
Barcelona, September 2021

Abstract

The results in this thesis concern extremal and probabilistic topics in number theoretic settings. We prove sufficient conditions on when certain types of integer solutions to linear systems of equations in binomial random sets are distributed normally, results on the typical approximate structure of pairs of integer subsets with a given sumset cardinality, as well as upper bounds on how large a family of integer sets defining pairwise distinct sumsets can be. In order to prove the typical structural result on pairs of integer sets, we also establish a new multipartite version of the method of hypergraph containers, generalizing earlier work by Morris, Saxton and Samotij.

Keywords. Additive combinatorics, probabilistic combinatorics, extremal combinatorics, Sidon sets, inverse sumset theory, independent sets in hypergraphs

Contents

Acknowledgments

I want to first and foremost thank my parents for providing me with constant support, both by encouraging me to actually pursue PhD studies, but also by offering constant financial support and providing me with much needed shelter in Germany during the hot summer months.

Next I want to thank my advisors, Oriol Serra and Juanjo Rué, who were always there to guide me through all aspects of the PhD, scientific and otherwise. I am still surprised how smoothly they managed to maintain contact during the sudden disruption during the latter period, while having to juggle new teaching challenges at the same time. I am also grateful to Simon Griffiths and Rob Morris for allowing me to visit their institutions.

Lastly, I am also thankful to my co-authors (if not previously mentioned), Marcelo Campos, Javier Cilleruelo and Matthew Coulson, as well as to everyone I had the pleasure of getting to know during my stay in Barcelona, namely Christoph Spiegel, Pilar Cano, Vasiliki Velona, Gonzalo Fiz Pontiveros, Clement Requilé, Lluís Vena, Marc Noy, Guillem Perarnau, Ignasi Sau and Dimitrios Thilikos.

Introduction

The general goal in additive combinatorics – historically also called combinatorial or additive number theory – is to study the additive structure of sets in certain ambient groups. Extremal combinatorics studies how large a collection of finite objects can be before exhibiting certain structural requirements. Probabilistic combinatorics analyzes random combinatorial structures, identifying in particular the structure of typical combinatorial objects. Among the most celebrated outputs is the study of random graphs initiated by Erdős and Rényi [32]. A particularly striking example of how these three areas interweave is the development by Erdős of the probabilistic method in number theory and in combinatorics, which exhibits the existence of many extremal structures in additive settings using probabilistic means.

The topics in this thesis all lie in the intersection of these three areas, and concern the following problems.

Integer solutions to systems of linear equations. The study of how large a subset of the integers can be before containing solutions to a given system of linear equations is a classical topic in additive number theory. Landmark results by Roth [84] and Szemerédi [103] solved the specific case of arithmetic progressions, showing that any subsets of the integers of positive upper density will contain arithmetic progressions of arbitrary lengths. Other well-studied topics that can be stated in the language of linear equations are those of sum-free and Sidon sets. Recent years have seen a lot of developments [86, 59] regarding *threshold results* for certain integer solutions to an arbitrary given system of linear equations, answering the question when one expects the binomial random subset of an initial segment of integers to contain solutions almost surely. The next logical question is the following. Suppose we are in the range that there will exist integer solutions in the binomial random set, how are these solutions distributed? In Chapter 1, we will take steps towards answering this question by providing sufficient conditions for when a wide variety of solutions follow a normal distribution. We will also discuss how in certain cases, these sufficient conditions are also necessary.

Independent sets in hypergraphs. The method of hypergraph containers is a very general tool introduced independently by Balogh, Morris and Samotij [9] and Saxton and Thomason [92] that can be used to obtain results on the number and structure of independent sets in hypergraphs, provided that the hyperedges in this hypergraph are sufficiently evenly distributed. While this topic may seem out of place in the context of this thesis, the connection appears because many additive problems can be encoded as studying independent sets in hypergraphs. For instance, the problem of integer sets free of arithmetic progressions mentioned before can be understood as the study of independent sets in the hypergraph that has the integers as its vertex set, while the edges are defined by the arithmetic progressions. In Chapter 2 we are going to establish an extension of a recent asymmetric version of the container method due to Morris, Samotij and Saxton [77] to multipartite hypergraphs. This extension is one of the main ingredients in the proof of a structure theorem proved in Chapter 3.

Sets with bounded sumset. What can be said about the structure of two finite sets in an abelian group if their Minkowski sum is not much larger than the sets themselves? The extremal case when the Minkowski sum is as small as possible was solved by Kneser [64] and Kemperman [62]. In particular, this can happen if and only if the Minkowski sum is periodic with respect to a proper subgroup. A fundamental result by Freĭman [40, 41] concerns the more general problem of obtaining structural results if one only asks for the *growth* of a set to be linear. He obtained a structural characterization in terms of a bounded number of arithmetic progressions. In Chapter 3 we are going to establish two types of results. First, we are going to establish so-called *robust* versions of classical theorems of Kneser and Freĭman. Robust here refers to the fact that instead of asking for structural information on the constituent sets with the knowledge that their sumset is small, we only require that this holds for a large subset. The second part of the chapter concerns the typical structure of pairs of sets with small Minkowski sum, that is, what if we only want to give a structural statement for *almost all* pairs of sets with a sumset of a given size? We give an approximate structure theorem that holds for almost the complete range of possible sumset sizes.

Sidon set systems. Sidon sets are subsets of abelian groups such that all pairwise sums of their elements are distinct, or equivalently, subsets such that their sumset is as large as possible. There is a large body of results on these sets, covering among others the question of how large a Sidon set of an initial integer segment can be, or when one expects a binomial random subset of such a segment to be a Sidon set. In Chapter 4 we generalize the notion of Sidon sets to set systems and establish the corresponding bounds for the two questions above. We also prove a so-called *relative density* result conditional on a conjecture on the specific structure of maximal Sidon systems.

General notation We use $\mathbb{N} = \{1, 2, 3, \dots\}$ to denote the set of positive integers. If $n \in \mathbb{N}$ is a positive integer, $[n]$ will denote the set $\{1, \dots, n\}$, and for any real number $0 \leq p \leq 1$ and any finite set X, we write X_p for the binomial random set obtained by including independently each element in X with probability p. Here, p may be a function depending on for instance the cardinality of X. For any not necessarily finite set X, denote by 2^X the power set of X, and for any non-negative integer k by $\binom{X}{k} = \{F \in 2^X : |F| = k\}$ the set of k-element subsets of X. Finally, for positive integers m and r, we write $\mathbb{Z}^{r \times m}$ for the set of integer valued matrices with r rows and m columns.

Chapter 1

Integer solutions to systems of linear equations

The main contribution of this chapter are sufficient conditions in order for the number of nontrivial integer solutions to a system of linear equations in binomial random sets to follow a normal distribution. All original work presented in this chapter is based on [87] and was done jointly with Juanjo Rué.

Suppose $A \in \mathbb{Z}^{r \times m}$ is an integer matrix, where $m > r$ are positive integers, and for some integer vector $\mathbf{b} \in \mathbb{Z}^r$, denote by $S(A, \mathbf{b}) = \{\mathbf{x} \in \mathbb{Z}^m : A\mathbf{x} = \mathbf{b}\}$ the set of integer solutions to the system of linear equations $A\mathbf{x} = \mathbf{b}$.

The main question that is investigated in this chapter is of the following nature: For a specific subset $T \subset \mathbb{Z}$, what can be said about the number of solutions to the system $A\mathbf{x} = \mathbf{b}$ with all of its coordinates lying in T, that is $|S(A, \mathbf{b}) \cap T^m|$? Classical examples for T are for instance the first n positive integers and one then studies the asymptotic behavior of $|S(A, \mathbf{b}) \cap [n]^m|$ as n tends to infinity. The prototypical example in this setting is to consider the system defined by $A_{k\text{-AP}} = (a_{i,j}) \in \mathbb{Z}^{(k-2)\times k}$ with $a_{i,i} = a_{i,i+2} = 1$, $a_{i,i+1} = -2$, $i \in [k-2]$ and 0 everywhere else, and $\mathbf{b}_{k\text{-AP}} = \mathbf{0} \in \mathbb{Z}^{k-2}$, which encodes k-term arithmetic progressions (k-APs). Note that when defining the problem this way, one additional constraint needs to be imposed, since it clearly holds that $(x, x, \ldots, x) \in S(A_{k\text{-AP}}, \mathbf{0}) \cap [n]^k$ for every $x \in [n]$. It turns out that for many systems of linear equations including k-APs, it is natural to only consider solutions where all the variables are pairwise distinct, which we will from now on refer to as *proper solutions*.

Definition 1.1. Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$, and $\boldsymbol{b} \in \mathbb{Z}^r$. Then the set $S_0(A, \boldsymbol{b})$ of *proper* solutions is the subset of $S(A, \boldsymbol{b})$ with all coordinates being pairwise distinct, that is

$$S_0(A, \boldsymbol{b}) = \{(x_1, \ldots, x_m) \in S(A, \boldsymbol{b}) : x_i \neq x_j \text{ for all } 1 \leq i < j \leq m\}.$$

Using Fourier Analytical techniques, Roth in [84] proved that every set of positive relative density in $[n]$ must contain a 3-term arithmetic progression, that is, he established the following.

Theorem 1.2 ([84]). *Let $\delta > 0$. Then there exists an integer n_0 such that for every integer $n \geq n_0$ and every subset $U \subset [n]$ satisfying $|U| \geq \delta n$, it holds that*

$$S_0(A_{3\text{-}AP}, \mathbf{0}) \cap U^3 \neq \emptyset.$$

This result is interesting in a qualitative sense, but from a quantitative standpoint one would like to know whether it is actually possible to only have very few 3-APs. It turns out that this is not the case. By applying Theorem 1.2 several times to smaller sub-intervals, Varnavides in [107] managed to prove what is usually referred to as a *supersaturation result*, which is a topic that will be important in Chapters 2 and 3.

Theorem 1.3 ([107]). *Let $\delta > 0$. Then there exist $\epsilon > 0$ and an integer n_0 both only depending on δ, such that for every integer $n \geq n_0$ and every subset $U \subset [n]$ satisfying $|U| \geq \delta n$, it holds that*

$$|S_0(A_{3\text{-}AP}, \mathbf{0}) \cap U^3| \geq \epsilon n^2.$$

Note that it is not hard to check that the number of 3-APs (and in fact, k-APs when k is fixed) in $[n]$ is $O(n^2)$, since fixing the smallest element and the common difference completely determines the progression. Hence Varnavides' theorem states that if a set contains a positive proportion of the elements of $[n]$, then it will also contain a positive proportion of the 3-term arithmetic progressions of it.

In 1975, Szemerédi in [103] managed to generalize Theorem 1.2 to the case of arbitrary (fixed) k, as the first application of his celebrated regularity lemma. Since the arguments that established Theorem [107] work just the same in this case, we state this quantitative version of the result.

Theorem 1.4 ([103]). *Let $\delta > 0$ and let $k \geq 3$ be an integer. Then there exist $\epsilon > 0$ and an integer n_0 both only depending on δ and k, such that for every integer $n \geq n_0$ and every subset $U \subset [n]$ satisfying $|U| \geq \delta n$, it holds that*

$$|S_0(A_{k\text{-}AP}, \mathbf{0}) \cap U^k| \geq \epsilon n^2.$$

Note that in the other direction, constructions by Behrend [12] and Rankin [83] show that there are indeed subsets of $[n]$ of cardinality almost linear in n that are free of k-APs, and hence the statements in Theorems 1.2 and 1.4 are sharp in a qualitative sense. For a general matrix $A \in \mathbb{Z}^{r \times m}$, the system of linear equations $A\mathbf{x} = \mathbf{0}$ is called *density regular* if a theorem akin to Theorem 1.2 holds, that is, if for any fixed $\delta > 0$ there exists an $n_0 \in \mathbb{N}$ only depending on A and δ, such that for any $n \geq n_0$, every set $U \subset [n]$ satisfying $|U| \geq \delta n$ contains at least one proper solution $\mathbf{x} \in S_0(A, \mathbf{0}) \cap [n]^m$. In [38] Frankl, Graham and Rödl observed that a system is density regular if and only if the all 1s vector $\mathbf{1} \in \mathbb{Z}^m$ is a solution. They also managed to prove the following generalization of Theorems 1.3 and 1.4.

Theorem 1.5 ([38]). *Let $\delta > 0$ and for integers $m > r$ let $A \in \mathbb{Z}^{r \times m}$ be a full rank integer matrix such that $A \cdot \mathbf{1} = \mathbf{0}$, and for every pair of integers $1 \leq i < j \leq m$ there exists a solution $\mathbf{x} = (x_1, \ldots, x_m) \in S(A, \mathbf{0})$ satisfying $x_i \neq x_j$. Then there exist $\epsilon > 0$ and $n_0 \in \mathbb{N}$ only depending on A and δ such that for any $n \geq n_0$, every set $U \subset [n]$ of cardinality $|U| \geq \delta n$ satisfies*

$$|S_0(A, \mathbf{0}) \cap U^m| \geq \epsilon n^{m-r}.$$

Note that it is not hard to see that if $A \in \mathbb{Z}^{r \times m}$ is an integer matrix of rank $\ell \leq r$, then $|S(A, \mathbf{0}) \cap [n]^m| \leq n^{m-\ell}$, and hence Theorem 1.5 states that U will contain at least an ϵ proportion of all solutions to the system $A\mathbf{x} = \mathbf{0}$.

Producing large subsets $U \subset [n]$ that are free of solutions for an arbitrary system of linear equations $A\mathbf{x} = \mathbf{0}$ seems quite difficult, and in fact, they can be of widely different sizes. For instance, while the aforementioned constructions of Behrend and Rankin provide k-AP free sets $U \subset [n]$ that are larger than $n^{1-\epsilon}$ for any $\epsilon > 0$ (given that n is large enough), the matrix $A_S = \begin{pmatrix} 1 & 1 & -1 & -1 \end{pmatrix}$ where $A_S\mathbf{x} = \mathbf{0}$ defines the *Sidon equation* only allows for solution free sets of size $O(\sqrt{n})$ (cf. Chapter 4). Note that A_S satisfies the requirements of Theorem 1.5, since $1 + 1 - 1 - 1 = 0$ while at the same time for instance $(1, 5, 2, 4)$ is a proper solution. In both cases, these maximal constructions are heavily structured, and so one would expect the random binomial set $[n]_p$ to contain solutions before it reaches this maximum size. This is indeed the case, and as an easy example, let us show when we expect to see k-APs in a binomial random set $[n]_p$. We prove the following *threshold* result.

Proposition 1.6. *Let $n \geq k \geq 3$ be integers, and $0 \leq p = p(n) \leq 1$. Then*

$$\lim_{n\to\infty} \mathrm{P}(S_0(A_{k\text{-}AP}, \mathbf{0}) \cap [n]_p^k \neq \emptyset) = \begin{cases} 0, & \text{if } p = o(n^{-2/k}), \\ 1, & \text{if } p = \omega(n^{-2/k}). \end{cases}$$

In particular, if $p = \omega(n^{-2/k})$, then $|S_0(A_{k\text{-}AP}, \mathbf{0}) \cap [n]_p^k| \sim |S_0(A_{k\text{-}AP}, \mathbf{0}) \cap [n]^k|p^k$ almost surely.

Proof. Denote by $M = |S_0(A_{k\text{-AP}}, \mathbf{0}) \cap [n]^k|$ the number of k-APs in $[n]$. Let $P_1, \ldots, P_M$ denote the k-APs in $[n]$, and for $j \in [M]$ let I_j denote the indicator variable for the event $P_j \subset [n]_p$. Defining $X = |S_0(A_{k\text{-AP}}, \mathbf{0}) \cap [n]_p^k|$, we see that $X = \sum_{j\in[M]} I_j$, and hence

$$\mathbb{E}(X) = \sum_{j\in[M]} \mathbb{E}(I_j) = Mp^k = \Theta(n^2p^k).$$

If $p = o(n^{-2/k})$, we thus have $\mathbb{E}(X) = o(1)$, and hence using Markov's inequality we see that $\mathrm{P}(X \geq 1) \leq \mathbb{E}(X) = o(1)$ which is what we wanted to show.

For the other direction, first note that we can assume that $p = o(1)$, since otherwise $\mathbb{E}(|[n]_p|) = np = \Omega(n)$, and standard concentration bounds for random variables that are sums of independent and identically distributed indicator variables will tell us that this also holds almost surely, and hence an application of Theorem 1.4 gives the desired result. Let us now upper bound the variance $\mathbb{V}\mathrm{ar}(X)$. We will make use of the following fact. If P is some k-AP, then for any integer $0 \leq s \leq k$, there are $O(n^{\max(2-s,0)})$ k-APs Q such that $|P \cap Q| = s$. Indeed, if we fix a single element in an AP, there are $O(n)$ choices for the common difference and $O_k(1)$ choices for the position of the single fixed element. Similarly, if at least two elements are fixed, specifying any of the $O_k(1)$ possible positions of them completely determines the progression. We thus see

$$\begin{aligned} \mathbb{V}\mathrm{ar}(X) &= \sum_{1\leq i\leq j\leq M} \mathbb{E}(I_iI_j) - \mathbb{E}(I_i)\,\mathbb{E}(I_j) \\ &= M(p^k - p^{2k}) + 2\sum_{s\in[k-1]} \sum_{\substack{1\leq i<j\leq M \\ |P_i\cap P_j|=s}} p^{2k-s} - p^{2k} \\ &= M(p^k - p^{2k}) + O(n(p^{2k-1} - p^{2k})) + O(1) \\ &\sim \mathbb{E}(X). \end{aligned}$$

Here, for the last line we have also used the fact that $p = \omega(1/n)$. So after an application of a version of Chernoff's bound (see Lemma 4.19), we see that with probability $1 - 2\exp(-\sqrt{\mathbb{E}(X)}/4)$, it holds that

$$X = \mathbb{E}(X) \pm \mathbb{E}(X)^{3/4} \sim Mp^k = \omega(1),$$

which is what we wanted to show. ■

Theorems similar to Proposition 1.6 have been proved for several specific (systems of) linear equations. For instance, Godbole, Janson, Locantore Jr. and Rapoport in [45] determined the threshold probability p for when $[n]_p$ contains solutions to the equation

$$a_1 + a_2 + \cdots + a_k = b_1 + b_2 + \cdots + b_k$$

such that $\{a_1, \ldots, a_k\} \neq \{b_1, \ldots, b_k\}$ as multisets. This example also shows that while the notion of proper solutions is appropriate for k-term arithmetic progressions, there are systems of linear equations that contain solutions with repeated variables that are still of interest. In [86], generalizing an earlier definition by Ruzsa [88], Rué, Spiegel and Zumalacárregui defined a notion of *non-trivial solutions* that captures all the aforementioned examples, and proved a threshold theorem for when $[n]_p$ will almost surely contain non-trivial solutions to a given system.

We will need this notion to state our results later, so we will now give it in full. Suppose for positive integers $m > r$ we are given an integer matrix $A \in \mathbb{Z}^{r\times m}$ and an integer vector $\mathbf{b} \in \mathbb{Z}^r$. Then for a solution $\boldsymbol{x} = (x_1, \dots, x_m) \in S(A, \boldsymbol{b})$, we define $\mathfrak{p}(\boldsymbol{x}) \subset 2^{[m]}$ to be the partition of $[m]$ such that for any $i, j \in [m]$ it holds that $x_i = x_j$ if and only if i and j are in the same partition class of $\mathfrak{p}(\boldsymbol{x})$. One can view $\mathfrak{p}(\boldsymbol{x})$ as an ordered $|\mathfrak{p}(\boldsymbol{x})|$-tuple $(C_1, \dots, C_{|\mathfrak{p}(\boldsymbol{x})|})$ such that $\min C_i < \min C_j$ whenever $i < j$. Doing this, we can now define the matrix $A_{\mathfrak{p}(\boldsymbol{x})}$ in the following way. Suppose the columns of A are denoted by $\boldsymbol{c}_1, \dots, \boldsymbol{c}_m \in \mathbb{Z}^r$, then

$$A_{\mathfrak{p}(\boldsymbol{x})} = \left(\sum_{i\in C_1} \boldsymbol{c}_i \;\middle|\; \sum_{i\in C_2} \boldsymbol{c}_i \;\middle|\; \cdots \;\middle|\; \sum_{i\in C_{|\mathfrak{p}(\boldsymbol{x})|}} \boldsymbol{c}_i \right),$$

so $A_{\mathfrak{p}(\boldsymbol{x})} \in \mathbb{Z}^{r\times|\mathfrak{p}(\boldsymbol{x})|}$ is the $r \times |\mathfrak{p}(\boldsymbol{x})|$ matrix obtained by contracting all columns in the same partition class, with columns ordered by the minimum index in each class. Finally, define the set

$$\mathfrak{P}(A) = \{\mathfrak{p} \subset 2^{[m]} : \mathfrak{p} \text{ is a partition of } [m] \text{ and } \operatorname{rk}(A_{\mathfrak{p}}) = \operatorname{rk}(A)\}.$$

We are finally ready to introduce the notion of a *non-trivial* solution.

Definition 1.7. Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r\times m}$, and $\boldsymbol{b} \in \mathbb{Z}^r$. Then the set $S_1(A, \boldsymbol{b})$ of *non-trivial* solutions is the subset of $S(A, \boldsymbol{b})$ with associated partitions coming from $\mathfrak{P}(A)$, that is

$$S_1(A, \boldsymbol{b}) = \{\boldsymbol{x} \in S(A, \boldsymbol{b}) : \mathfrak{p}(\boldsymbol{x}) \in \mathfrak{P}(A)\}.$$

On the surface, this definition might seem quite arbitrary, but the interested reader is invited to peruse the discussions in [86] and [100] which show that in some sense it is quite natural in that it encompasses the natural notions of non-triviality for specific systems of linear equations studied in the literature.

We have already seen in Theorem 1.5 that the matrix A will need to satisfy some specific conditions, which motivates the following two definitions that were also used in [86] and [100]. In order to state it, we need to introduce the following notation. If $Q \subset [m]$, let $A^Q \in \mathbb{Z}^{r\times|Q|}$ denote the $r \times |Q|$ matrix obtained by only keeping the columns of A indexed by Q. We notice that the rank of the empty matrix $A^\varnothing$ is zero. Note that while we have thus far identified solution tuples $\boldsymbol{x} = (x_1, \dots, x_m) \in \mathbb{Z}^m$ with the corresponding column vectors, we will abuse notation slightly by letting $\boldsymbol{x}^Q$ denote the vector obtained by only keeping the *rows* of $\boldsymbol{x}$ indexed by Q. This should not be confused with the notation $\mathbf{x}^k$, where k is a positive integer, which we will use to denote the vector

$$\mathbf{x}^k = (\underbrace{\mathbf{x}, \dots, \mathbf{x}}_{k \text{ times}}) = (x_1, \dots, x_m, x_1, \dots, x_m, \dots, x_1, \dots, x_m) \in \mathbb{Z}^{km}.$$

Since index sets will be denoted by capital letters and positive integers by lowercase ones, the meaning will always be clear in context.

Definition 1.8. An integer matrix $A \in \mathbb{Z}^{r\times m}$ is said to be

i) *positive* if there exists an integer solution to $A\boldsymbol{x} = \mathbf{0}$ having all positive entries:

$$S(A, \mathbf{0}) \cap \mathbb{N}^m \neq \varnothing,$$

and if for any pair of indices $i, j \in [m]$, $i \neq j$, there exists a solution $(x_1, \dots, x_m) \in S(A, \mathbf{0})$ satisfying $x_i \neq x_j$.

ii) *abundant* if $\mathrm{rk}(A) > 0$ and the removal of at most two columns does not change the rank of A:

$$\mathrm{rk}(A^Q) = \mathrm{rk}(A)$$

for any $Q \subset [m]$ satisfying $|Q| \geq m - 2$.

Sometimes the second requirement for positivity in Definition 1.8 is called *irredundancy*, and we have already encountered it as a requirement in the statement of Theorem 1.5. For our upcoming applications, we will never consider positivity and irredundancy separately, and hence we have combined those two properties into a single one for expediency's sake. Before proceeding to the threshold theorem of Rué, Spiegel and Zumalacárregui, we need to define one more parameter, which will measure the densest subsystem of A. The motivation behind this can be compared to the graph setting, where one wants to study the number of occurrences of some fixed graph G as a subgraph of a binomial random graph with n vertices and edge probability p. Here, one naively would expect the threshold for a random graph to contain G to be when p is around $n^{-v(G)/e(G)}$, which is when the expected number of copies of G flips from 0 to positive. But this is not the case: if G contains a subgraph H with $v(H)/e(H) > v(G)/e(G)$, then $n^{-v(H)/e(H)}$ will define the threshold instead. A similar behavior occurs for the distribution of subgraph counts, as shown by Ruciński in [85].

Definition 1.9. For positive integers $m > r$ and a positive integer matrix $A \in \mathbb{Z}^{r\times m}$, define the *density* $c(A)$ of A by

$$c(A) = \max_{\emptyset \neq Q \subset [m]} \frac{|Q|}{|Q| - r_Q},$$

where $r_Q = r_Q(A) = \mathrm{rk}(A) - \mathrm{rk}(A^{\overline{Q}})$, and $\overline{Q} = [m] \setminus Q$.

Note that this is indeed well-defined, since for a positive matrix we see that $\mathrm{rk}(A^{\overline{Q}}) \geq \mathrm{rk}(A) - |Q| + 1$ for every $\emptyset \neq Q \subset m$ (see, for instance, the proof of Lemma 1.17). We can now state the threshold result on appearance of non-trivial solutions in binomial random sets due to Rué, Spiegel and Zumalacárregui.

Theorem 1.10 ([86]). *Let $m > r$ be positive integers and $A \in \mathbb{Z}^{r\times m}$ a positive integer matrix. Then*

$$\mathrm{P}(S_1(A, \mathbf{0}) \cap [n]_p^m \neq \emptyset) = \begin{cases} 0, & \text{if } p = o(n^{-1/c(A)}), \\ 1, & \text{if } p = \omega(n^{-1/c(A)}). \end{cases}$$

Note that it is not hard to check that $c(A_{k-\mathrm{AP}}) = k/2$ and hence this generalizes Proposition 1.6. Later, in his dissertation [100], Spiegel was able to prove Theorem 1.10 also in the non-homogeneous case $A\mathbf{x} = \mathbf{b}$ with $\mathbf{b} \in \mathbb{Z}^r$, assuming that A is abundant in addition to being positive. Before talking about distributions of solutions in binomial random sets, let us make one final detour. Proposition 1.6 and Theorem 1.10 both essentially answer the question of what happens if in Theorem 1.5, instead of an arbitrary subset $U \subset [n]$, one considers a binomial random one. A different probabilistic consideration is to replace $[n]$ itself by $[n]_p$ and ask the following question, which we will refer to as a *relative density* problem. Consider a fixed system $A\mathbf{x} = \mathbf{b}$ and a $\delta > 0$. Does there exist a threshold probability p such that almost surely, for every subset $U \subset [n]_p$ of relative density δ there exists a non-trivial solution $\mathbf{x} \in S_1(A, \mathbf{b}) \cap U^m$? Note that clearly, if $[n]_p$ itself does not contain any non-trivial solutions to the system, then neither will any subset of it, and hence this threshold, if it exists, will always lie beyond that of Theorem 1.10. For the specific case of 3-APs, Kohayakawa Łuczak and Rödl were able to prove such a result in [66]. This was then generalized to arbitrary fixed k by Schacht [93] and independently by Conlon and Gowers [26]. A nice combinatorial proof was later discovered by

Saxton and Thomason [92] and independently by Balogh, Morris and Samotij [10] using the method of hypergraph containers, which will be explored in more detail in Chapter 2. Using the container method, the relative density problem for arbitrary homogeneous systems of linear equations $A\mathbf{x} = \mathbf{0}$ was then fully solved independently by Spiegel [99] and Hancock, Staden and Treglown [59]. As a final mention, note that these relative density problems can of course be considered in sparse subsets of $[n]$ other than $[n]_p$, and in fact the methods by Schacht and Conlon and Gowers allow for the *transference* of any density result similar to Theorem 1.4, Szemerédi's Theorem, to any sparse set $T \subset [n]$, as long as T satisfies some pseudo-randomness conditions. A particular striking example of this type of transference was used by Green and Tao in their celebrated proof of the existence of arithmetic progressions with arbitrary lengths in the primes [53]. We will return to relative density problems in Chapter 4, establishing such a result for a generalization of Sidon sets.

Let us now proceed to the main point of investigation of this chapter. Theorem 1.10 and its generalization to non-homogeneous systems by Spiegel [100] tells us that essentially, whenever the expected number of solutions to the densest subsystem of A is infinite, there will exist at least one solution. Moreover, a concentration argument similar to the one employed in Proposition 1.6 shows that almost surely, it will differ from its expectation only in lower order terms. But it is still an interesting problem to ask the more fine-grained question of what the exact limiting distribution of the solutions in this range is. Note that for arithmetic progressions, this question was studied before by Barhoumi-Andréani, Koch and Liu in [11] and solved in a very strong sense. That is, not only did they determine the limiting distribution for the number of k-APs in $[n]_p$ even in the case $\log n \gg k \gg 1$, that is, for unbounded but sub-logarithmic k, they also proved a bivariate central limit theorem for the joint distribution when considering the counts of two distinct progression lengths. Conversely, our results that will be stated shortly are applicable only to progressions of a fixed length. Another related result was obtained by Griffiths, Koch and Secco in [54], where they proved results that also imply central limit theorems for the number of proper solutions to certain systems of equations. While their results are essentially equivalent to ours in many natural and well-studied cases such as k-APs, Sidon sets, sum-free sets, it is possible to construct examples of systems in which their statements are not applicable. For a random variable X with finite first moment $\mathbb{E}(X)$ and non-zero finite variance $\mathbb{V}\mathrm{ar}(X)$, denote by

$$\tilde{X} = (X - \mathbb{E}(X))/\sqrt{\mathbb{V}\mathrm{ar}(X)}$$

its normalization. We write $X_n \xrightarrow{d} Y$ when a sequence of random variables $\{X_n\}_{n\geq 1}$ tends in distribution to Y. We are now ready to state our first main result regarding the distribution of proper solutions to linear systems.

Theorem 1.11. *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r\times m}$ a positive and abundant integer matrix, and $\boldsymbol{b} \in \mathbb{Z}^r$ such that $S(A, \boldsymbol{b}) \neq \emptyset$. Furthermore, let n be an integer, $0 \leq p := p(n) \leq 1$ and X_n the random variable $|S_0(A, \boldsymbol{b}) \cap [n]_p^m|$, which counts the number of proper solutions $\boldsymbol{x} \in [n]_p^m$ to $A\boldsymbol{x} = \boldsymbol{b}$. Then*

$$\tilde{X}_n \xrightarrow{d} \mathcal{N}(0, 1)$$

if $n(1-p) \to \infty$ and $np^{c(A)} \to \infty$.

Note that the statement only concerns sufficient conditions, but the topic of whether or not they are necessary will be discussed as well in the final section of this chapter. Our second main result concerns the distribution of non-trivial solutions to certain systems. Note that when we want to investigate this distribution, we actually need to look at $c(A_{\mathfrak{p}})$ for any partition type $\mathfrak{p}$ that is to be considered. A special case in which it suffices to only consider $c(A)$ is that of *strictly balanced* systems.

Definition 1.12. Let $m > r$ be positive integers and $A \in \mathbb{Z}^{r\times m}$ be positive. Then the matrix A is called *strictly balanced* if

$$c(A) = \frac{m}{m - \mathrm{rk}(A)} > \max_{\emptyset \subsetneq Q \subsetneq [m]} \frac{|Q|}{|Q| - r_Q(A)},$$

and furthermore for every $\mathfrak{p} \in \mathfrak{P}(A) \setminus \{\{1\},\{2\},\dots,\{m\}\}$ such that $A_{\mathfrak{p}}$ is positive, it holds that $c(A) > c(A_{\mathfrak{p}})$.

We can prove the following theorem about the distribution of nontrivial solutions of strictly balanced systems of linear equations.

Theorem 1.13. *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r\times m}$ a positive and abundant strictly balanced integer matrix such that $S(A, \mathbf{0}) \neq \emptyset$. Furthermore, let n be an integer, $0 \leq p := p(n) \leq 1$ and X_n the random variable $|S_1(A, \mathbf{0}) \cap [n]_p^m|$ that counts the number of non-trivial solutions $\boldsymbol{x} \in [n]_p^m$ to $A\boldsymbol{x} = \mathbf{0}$. Then*

$$\tilde{X}_n \xrightarrow{d} \mathcal{N}(0,1)$$

if $n(1-p) \to \infty$ and $np^{c(A)} \to \infty$.

Note that when comparing Theorems 1.11 and 1.13 to Theorem 1.10 and its generalization, ours are located in a more restrictive setting. This will also be true for our main technical result, Theorem 1.23, which the two theorems above are corollaries of. But this is not surprising, since the study of the distribution is more delicate than the threshold behavior, as can be seen by another result from [86], in which the authors explore the behavior at the threshold of nontrivial solutions to $A\boldsymbol{x} = \mathbf{0}$ for strictly balanced systems, that is, they work exactly in the same setting as in Theorem 1.13.

The remainder of this chapter is structured as follows. In Section 1.1 we will state some previously known algebraic properties concerning systems of linear equations, as well as prove new ones needed for our application. In particular, we will introduce a new solution type that is more fine-grained than either proper or non-trivial solutions, which we will need for the statement of our main meta theorem, Theorem 1.23. Section 1.2 is then fully devoted to this meta theorem. First we are going to state it and show how it implies Theorems 1.11 and 1.13, and then follow up by giving its proof in full. Finally, in Section 1.3, we will discuss the necessity of the sufficient conditions.

1.1 Algebraic properties of systems of linear equations

Let us start by investigating the relation between proper and non-trivial solutions. For any $A \in \mathbb{Z}^{r\times m}$ and $\boldsymbol{b} \in \mathbb{Z}^r$ it is clear that

$$S_0(A, \boldsymbol{b}) \subset S_1(A, \boldsymbol{b}) \subset S(A, \boldsymbol{b}).$$

Furthermore, suppose $s > 1$ and we fix a specific $(C_1, \dots, C_s) = \mathfrak{p} \in \mathfrak{P}(A)$ such that $\min C_i < \min C_j$ whenever $i < j$. Then if $\boldsymbol{x} = (x_1, \dots, x_m) \in S_1(A, \boldsymbol{b})$ satisfies $\mathfrak{p}(\boldsymbol{x}) = \mathfrak{p}$, we see that $\tilde{\boldsymbol{x}} = (x_{\min C_1}, \dots, x_{\min C_s}) \in \mathbb{Z}^s$ has s pairwise distinct coordinates and $A_{\mathfrak{p}}\tilde{\boldsymbol{x}} = \boldsymbol{b}$, so $\tilde{\boldsymbol{x}} \in S_0(A_{\mathfrak{p}}, \boldsymbol{b})$. Conversely, for any proper solution $\boldsymbol{x} = (x_1, \dots, x_s)$ in $S_0(A_{\mathfrak{p}}, \boldsymbol{b})$, we see that the vector $\tilde{\boldsymbol{x}} \in \mathbb{Z}^m$ with i-th coordinate $\tilde{x}_i$ defined as

$$\tilde{x}_i = \sum_{j=1}^{s} x_j \cdot \delta_{i \in C_j}$$

satisfies $\mathfrak{p}(\tilde{\boldsymbol{x}}) = \mathfrak{p}$ and $A\tilde{\boldsymbol{x}} = \boldsymbol{b}$ (here $\delta_{i\in C_j}$ denotes the indicator function which takes values 0 or 1 if $i \notin C_j$ and $i \in C_j$, respectively). So we get the following lemma, which is essentially Lemma 1.10 in [100].

Lemma 1.14. *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r\times m}$, and $\boldsymbol{b} \in \mathbb{Z}^r$. Then for any partition $\mathfrak{p} \in \mathfrak{P}(A)$ and any finite $T \subset \mathbb{Z}$, it holds that*

$$|\{\boldsymbol{x} \in S_1(A,\boldsymbol{b}) : \mathfrak{p}(\boldsymbol{x}) = \mathfrak{p}\} \cap T^m| = |S_0(A_{\mathfrak{p}},\boldsymbol{b}) \cap T^{|\mathfrak{p}|}|.$$

Lemma 1.14 will be very helpful because if A is positive, then the approximate size of $S_0(A,\boldsymbol{b})$ is not difficult to determine for any vector $\boldsymbol{b}$. Specifically, we have the following result which is Lemma 1.4 in [100].

Lemma 1.15. *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r\times m}$, and $\boldsymbol{b} \in \mathbb{Z}^r$ such that $S(A,\boldsymbol{b}) \neq \emptyset$. Then if $n \in \mathbb{N}$ it holds that*

$$|S_0(A,\boldsymbol{b}) \cap [n]^m| \leq |S(A,\boldsymbol{b}) \cap [n]^m| \leq n^{m-\mathrm{rk}(A)}.$$

Furthermore, if A is positive we also have

$$|S_0(A,\boldsymbol{b}) \cap [n]^m| = \Omega(n^{m-\mathrm{rk}(A)}).$$

In [86] the authors managed to find a threshold probability function for when non-trivial solutions to a homogeneous system of linear equations appear in $[n]_p$, but for our current purposes we need to make one more specification regarding types of solutions. The issue essentially lies in the fact that while A is positive, the same might not hold for $A_{\mathfrak{p}}$ for some $\mathfrak{p} \in \mathfrak{P}(A)$. Note that this can only happen in the case $\boldsymbol{b} \neq \boldsymbol{0}$, since clearly $S_0(A_{\mathfrak{p}},\boldsymbol{0}) \neq \emptyset$ implies that $A_{\mathfrak{p}}$ is positive. As a problematic example, consider for instance the matrix $A = \begin{pmatrix}1 & 1 & 1 & 1 & -1\end{pmatrix}$ which is positive since for instance $\boldsymbol{x} = (1,2,3,4,10) \in \mathbb{N}^5$ is a proper solution to the system $A\boldsymbol{x} = \boldsymbol{0}$. But taking $\boldsymbol{b} = 6$, we see that $\boldsymbol{y} = (1,2,3,3,3)$ satisfies $A\boldsymbol{y} = \boldsymbol{b}$. We have $\mathfrak{p}(\boldsymbol{y}) = (\{1\},\{2\},\{3,4,5\})$ and hence $A_{\mathfrak{p}(\boldsymbol{y})} = \begin{pmatrix}1 & 1 & 1\end{pmatrix}$ which is clearly not positive.

The problem arises as follows: if $A_{\mathfrak{p}}$ is not positive but we still have a significant number of solutions of this type – this can for instance be achieved by using the previous example as a gadget in a larger system – the number of solutions in the random binomial set will depend heavily on whether or not a specific bounded number of elements are included, and so getting any good results on their distribution is unlikely.

Having stated this motivation, we first introduce the concept of *positive partitions* as the subset $\mathfrak{P}_0(A) \subset \mathfrak{P}(A)$ such that

$$\mathfrak{P}_0(A) = \{\mathfrak{p} \in \mathfrak{P}(A) : A_{\mathfrak{p}} \text{ is positive.}\} \tag{1.1}$$

For any $\mathfrak{P} \subset \mathfrak{P}(A)$ we can now define the concept of *type-$\mathfrak{P}$* solutions.

Definition 1.16. Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r\times m}$, $\boldsymbol{b} \in \mathbb{Z}^r$, and $\mathfrak{P} \subset \mathfrak{P}(A)$. Then the set $S_{\mathfrak{P}}(A,\boldsymbol{b})$ of *type-$\mathfrak{P}$* solutions is the subset of $S(A,\boldsymbol{b})$ with associated partitions coming from $\mathfrak{P}$, that is

$$S_{\mathfrak{P}}(A,\boldsymbol{b}) = \{\boldsymbol{x} \in S(A,\boldsymbol{b}) : \mathfrak{p}(\boldsymbol{x}) \in \mathfrak{P}\}.$$

If $\mathfrak{P} = \{\mathfrak{p}\}$, we will write $S_{\mathfrak{p}}(A,\boldsymbol{b})$ instead of $S_{\{\mathfrak{p}\}}(A,\boldsymbol{b})$.

Non-trivial solutions are therefore the same as type-$\mathfrak{P}(A)$ solutions, while proper solutions are type-$\{\{1\},\dots,\{m\}\}$ solutions. The following notational convention will be useful. If $\boldsymbol{x} = (x_1,\dots,x_m) \in \mathbb{Z}^m$ is a tuple, we will write $\{\boldsymbol{x}\}$ as a shorthand for the set $\{x_1,\dots,x_m\}$. The next lemma is a straightforward consequence of Lemma 1.14 and the upper bound in Lemma 1.15.

Lemma 1.17. *Let $m > r$ and n be positive integers, $A \in \mathbb{Z}^{r\times m}$ positive, $\boldsymbol{b} \in \mathbb{Z}^r$, $\mathfrak{P} \subset \mathfrak{P}_0(A)$, and $Z \subset [n]$ a fixed set. Then*

$$\left|\{\boldsymbol{y} \in S_{\mathfrak{P}}(A,\boldsymbol{b}) \cap [n]^m : \{\boldsymbol{y}\} \cap Z \neq \emptyset\}\right| = O(n^{m-\mathrm{rk}(A)-1}).$$

Furthermore, if A is also abundant, then

$$\left|\{\boldsymbol{y} \in S_{\mathfrak{P}}(A, \boldsymbol{b}) \cap [n]^m : |\{\boldsymbol{y}\} \cap Z| \geq 2\}\right| = O(n^{m-\mathrm{rk}(A)-2}).$$

Proof. We will consider each partition $\mathfrak{p} \in \mathfrak{P}$ separately and consider $\boldsymbol{y} \in S_0(A_{\mathfrak{p}}, \boldsymbol{b}) \cap [n]^{|\mathfrak{p}|}$, which is equivalent by Lemma 1.14. If $\boldsymbol{y}$ contains an element of Z, then there exists some non-empty index set $Q \subset [|\mathfrak{p}|]$ and a vector $\boldsymbol{z} \in Z^{|Q|}$ such that $\boldsymbol{y}^{\bar{Q}} \in S_0(A_{\mathfrak{p}}^{\bar{Q}}, \boldsymbol{b} - A_{\mathfrak{p}}^{Q}\boldsymbol{z}) \cap [n]^{|\mathfrak{p}|-|Q|}$.* Note that here, $\bar{Q} = [|\mathfrak{p}|] \setminus Q$. By Lemma 1.15, we have

$$\left|S_0(A_{\mathfrak{p}}^{\bar{Q}}, \boldsymbol{b} - A_{\mathfrak{p}}^{Q}\boldsymbol{z}) \cap [n]^{|\mathfrak{p}|-|Q|}\right| \leq n^{|\mathfrak{p}|-|Q|-\mathrm{rk}(A_{\mathfrak{p}}^{\bar{Q}})}.$$

$A_{\mathfrak{p}}$ is positive, which implies that

$$\mathrm{rk}(A_{\mathfrak{p}}^{\bar{Q}}) \geq \mathrm{rk}(A_{\mathfrak{p}}) - |Q| + 1 = \mathrm{rk}(A) - |Q| + 1.$$

Indeed, first note that positivity implies that the removal of any single column will not change the rank of the resulting matrix. To see this, consider for a contradiction that there is a column whose removal would lower the rank of A, that is, it does not lie in the span of the remaining columns. Then there exists an invertible matrix $P \in \mathbb{Z}^{r \times r}$ such that PA contains a row with only a single non-zero entry in a column indexed by $i \in [m]$. But since we clearly have $A\mathbf{x} = \mathbf{0}$ if and only if $PA\mathbf{x} = P \cdot \mathbf{0} = \mathbf{0}$, this would imply that $x_i = 0$ for any solution $\mathbf{x} \in S(A, \mathbf{0})$, a contradiction to the positivity of A.

Now, since the rank of a matrix is the dimension of its column space, any subsequent column removal after the first one decreases the rank by at most 1. Hence

$$|\mathfrak{p}| - |Q| - \mathrm{rk}(A_{\mathfrak{p}}^{\bar{Q}}) \leq |\mathfrak{p}| - \mathrm{rk}(A) - 1 \leq m - \mathrm{rk}(A) - 1. \tag{1.2}$$

Since there are only $O(1)$ choices for $\mathfrak{p}$, Q and $\boldsymbol{z}$, this implies the result.

For the second part, note first that for any $|\mathfrak{p}| < m$, (1.2) actually implies the statement already, so it only remains to consider the case $|\mathfrak{p}| = m$, that is, $\boldsymbol{y} \in S_0(A, \boldsymbol{b}) \cap [n]^m$. If A is abundant, the removal of any two columns keeps the rank constant, while any subsequent removal will decrease it by at most 1 at a time. Hence for any $Q \subset [m]$ with $|Q| \geq 2$ we have

$$\mathrm{rk}(A^{\bar{Q}}) \geq \mathrm{rk}(A) - |Q| + 2,$$

and so

$$m - |Q| - \mathrm{rk}(A^{\bar{Q}}) \leq m - \mathrm{rk}(A) - 2,$$

which is what we wanted to show. ∎

1.1.1 Compounded matrices

While Lemma 1.17 showed that the upper bound of Lemma 1.15 is always helpful in the situation of counting solutions with some entries fixed beforehand, the requirement of positivity is sometimes too restrictive to do the same for the lower bound. Suppose now that $\boldsymbol{x} \in S_1(A, \boldsymbol{b})$ is some fixed non-trivial solution and we want to count the number of solutions $\boldsymbol{y} \in S_1(A, \boldsymbol{b})$ that intersect $\boldsymbol{x}$, that is, the size of the set $S_0(A_{\mathfrak{p}(\boldsymbol{y})}^{\bar{Q}}, \boldsymbol{b} - A_{\mathfrak{p}(\boldsymbol{y})}^{Q}\tilde{\boldsymbol{x}})$, where $Q \subset [|\mathfrak{p}(\boldsymbol{y})|]$ is some index set and $\tilde{\boldsymbol{x}} \in \{\boldsymbol{x}\}^{|Q|}$. Lemma 1.17 with $Z = \{\boldsymbol{x}\}$ immediately gives helpful upper bounds, but there are two issues when trying to apply the lower bound of Lemma 1.15 directly.

*Recall from our earlier definitions that in contrast to the matrix case, $\boldsymbol{y}^{\bar{Q}}$ here means that we only keep the rows indexed by Q.

The first is that when summing over several distinct Q, the same solution will be counted multiple times, but this could be alleviated by just counting a single Q that maximizes the cardinality of the corresponding set. The bigger issue is that it is not clear at all that the matrix $A^{\overline{Q}}_{\mathfrak{p}(\boldsymbol{y})}$ will be positive, and in fact this is not true in general even when $A_{\mathfrak{p}(\boldsymbol{y})}$ is itself positive. Consider for instance the matrix $A = \begin{pmatrix} 1 & 1 & -1 \end{pmatrix}$ (associated with the *Schur equation* $x + y = z$), which is both positive and abundant. Then for any Q with $|Q| = 1$, the equation $A^{\overline{Q}}(x, y) = \mathbf{0}$ implies either $x = y$ or $x = -y$, so in either case positivity will be violated.

Instead, we will consider the concept of the *compounded* matrix already used in [86] to study the distribution at the threshold. For matrices $A \in \mathbb{Z}^{r_A \times m_A}$, $B \in \mathbb{Z}^{r_B \times m_B}$ and a bijection $M\colon P \to [m_B]$ with $P = \{p_1 < \cdots < p_{|P|}\} \subset [m_A]$, define the $(r_A + r_B) \times (m_A + m_B - |P|)$ matrix $A \overset{M}{\times} B$ as

$$A \overset{M}{\times} B = \left(\begin{array}{c|c|c|c|c} A^{[m_A] \setminus P} & \boldsymbol{a}_{p_1} & \ldots & \boldsymbol{a}_{p_{|P|}} & \mathbf{0} \\ \hline \mathbf{0} & \boldsymbol{b}_{M(p_1)} & \ldots & \boldsymbol{b}_{M(p_{|P|})} & B^{[m_B] \setminus M(P)} \end{array}\right), \tag{1.3}$$

where $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_{m_A}$ denote the columns of A and $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_{m_B}$ the columns of B. Note that while we can apply this operator iteratively, the operation is in general not associative or even well-defined, that is

$$(A \overset{M}{\times} B) \overset{M'}{\times} C \neq A \overset{M}{\times} (B \overset{M'}{\times} C).$$

An exception to this is the case when the domain $\operatorname{dom}(M)$ is the empty set, which we will abbreviate by writing $A \dot{\times} B$. In general, whenever no parentheses are used the compounded matrix is implied to be constructed from left to right iteratively, that is $A \overset{M}{\times} B \overset{M'}{\times} C = (A \overset{M}{\times} B) \overset{M'}{\times} C$.

The reason to study this matrix is immediately apparent: Let $\rho_A = \rho_A(M)\colon [m_A] \to [m_A]$ be the bijective function that maps the column indices of the first $[m_A]$ columns of $A \overset{M}{\times} A$ to the corresponding ones of A, and define $\rho_B\colon [m_A + m_B - |P|] \setminus [m_A - |P|] \to [m_B]$ similarly. Recall that for any integer $t > 1$, $\boldsymbol{b}^t$ denotes the vector $(\boldsymbol{b}, \ldots, \boldsymbol{b}) \in \mathbb{Z}^{rt}$. Then $\boldsymbol{z} = (z_1, \ldots, z_{m_A + m_B - |P|})$ is a proper integer solution to the system of linear equations $(A \overset{M}{\times} B)\boldsymbol{z} = \boldsymbol{b}^2$ if and only if

$$\boldsymbol{x} = (x_1, \ldots, x_{m_A}) := (z_{\rho_A^{-1}(1)}, \ldots, z_{\rho_A^{-1}(m_A)}) \text{ and } \boldsymbol{y} = (y_1, \ldots, y_{m_B}) := (z_{\rho_B^{-1}(1)}, \ldots, z_{\rho_A^{-1}(m_B)})$$

are proper solutions to the systems $A\boldsymbol{x} = \boldsymbol{b}$ and $B\boldsymbol{y} = \boldsymbol{b}$, and for any $i \in [m_A]$ and $j \in [m_B]$ it holds that $x_i = y_j$ if and only if $i \in P$ and $j = M(i)$. In other words, elements of $S_0(A \overset{M}{\times} B, \boldsymbol{b}^2)$ correspond to pairs of proper solutions in $S_0(A, \boldsymbol{b})$ and $S_0(B, \boldsymbol{b})$ that agree exactly in the coordinates indicated by the function M.

We first state an easy but important property that any compounded matrix satisfies before looking at specific ones that will be important in the sequel.

Lemma 1.18. *Let $m_A > r_A$ and $m_B > r_B$ be positive integers, $A \in \mathbb{Z}^{r_A \times m_A}$, $B \in \mathbb{Z}^{r_B \times m_B}$, $P \subset [m_A]$, $M\colon P \to [m_B]$ a bijection and $Q \subset [m_A + m_B - |P|]$. If we define $Q_1 = Q \cap [m_A]$ and $Q_2 = Q \setminus Q_1$, then*

$$\operatorname{rk}\big((A \overset{M}{\times} B)^{\overline{Q}}\big) \geq \operatorname{rk}(A^{[m_A] \setminus Q_1}) + \operatorname{rk}(B^{\overline{M(P)} \setminus Q_2}).$$

In particular, $\operatorname{rk}(A \overset{M}{\times} B) \geq \operatorname{rk}(A) + \operatorname{rk}(B^{\overline{M(P)}})$.

Proof. It is clear that the rank of $(A \overset{M}{\times} B)^{[m_A] \setminus Q_1}$ is at least $rk(A^{[m_A] \setminus Q_1}) =: r_1$, and similarly the rank of $(A \overset{M}{\times} B)^{[m_A + m_b - |P|] \setminus ([m_A] \cup Q_2)}$ is at least $\operatorname{rk}(B^{\overline{M(P)} \setminus Q_2}) =: r_2$, so let $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_{r_1}$ be a collection of r_1 linearly independent vectors from the first $m_A - |Q_1|$ columns of $(A \overset{M}{\times} B)^{\bar{Q}}$ such that the corresponding column vectors of $A^{[m_A]\setminus Q_1}$ – that is, the vectors obtained by only considering the first r_A rows – are also linearly independent, and $\boldsymbol{c}_{r_1+1}, \dots, \boldsymbol{c}_{r_1+r_2}$ a collection of r_2 independent vectors from the $m_B - |P| - |Q_2|$ last columns. Then since the entries in the first r_A rows of $\boldsymbol{c}_{r_1+1}, \dots, \boldsymbol{c}_{r_1+r_2}$ are 0 and the column vectors obtained by only considering the first r_A rows of $\boldsymbol{c}_1, \dots, \boldsymbol{c}_{r_1}$ are linearly independent, any linear representation of $\mathbf{0}$

$$\lambda_1 \boldsymbol{c}_1 + \cdots + \lambda_{r_1+r_2} \boldsymbol{c}_{r_1+r_2}$$

must satisfy $\lambda_1 = \cdots = \lambda_{r_1} = 0$. But since the remaining r_2 columns were also linearly independent, we must have $\lambda_{r_1+1} = \cdots = \lambda_{r_1+r_2} = 0$ as well, and hence the $r_1 + r_2$ columns $\boldsymbol{c}_1, \dots, \boldsymbol{c}_{r_1+r_2}$ are linearly independent. Since the rank of a matrix is the dimension of its column space we are done. ∎

Next we will show that for certain compounded matrices this is indeed the correct rank.

Lemma 1.19. *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$, and $Q \subset [m]$. Then*

$$\operatorname{rk}(A \overset{\mathrm{id}_Q}{\times} A) = \operatorname{rk}(A) + \operatorname{rk}(A^{\bar{Q}}).$$

Proof. The lower bound follows from Lemma 1.18. For the upper bound, note that by performing elementary row and column operations we see that

$$\operatorname{rk}(A \overset{\mathrm{id}_Q}{\times} A) = \operatorname{rk}\left(\begin{array}{c|c|c} A^{Q} & A^{\bar{Q}} & \mathbf{0} \\ \hline \mathbf{0} & -A^{\bar{Q}} & A^{\bar{Q}} \end{array}\right).$$

Denote this right-hand matrix by B, and let $r_1 = \operatorname{rk}(A)$ and $r_2 = \operatorname{rk}(A^{\bar{Q}})$. Note that the matrix obtained by only keeping the first r rows of B will have rank r_1, while the matrix obtained by only keeping the last r rows will have rank r_2. Suppose now that we have $r_1 + r_2 + 1$ row vectors from B. If at least $r_1 + 1$ rows come from the upper half, then by the previous observation of its rank, there exists a nontrivial linear combination of the row vectors that equals $\mathbf{0} \in \mathbb{Z}^{1 \times (2m - |Q|)}$. This can then be extended to a nontrivial linear combination of all row vectors of B by setting the remaining coefficients to 0. If on the other hand more than $r_2 + 1$ rows come from the bottom half, we can achieve a nontrivial linear combination of only these rows that equals $\mathbf{0}$, which can again be extended. Since one of those two cases must happen, we have the required upper bound. ∎

The next result now shows how positivity of A can sometimes be passed on to a compounded matrix.

Lemma 1.20. *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$ be positive, and $Q \subset [m]$. If there exists an integer vector $\boldsymbol{y} \in S(A^{\bar{Q}}, \mathbf{0})$ with all coordinates nonzero, then $A \overset{\mathrm{id}_Q}{\times} A$ is positive.*

Proof. Since A was positive, there exists a solution $\boldsymbol{x} = (x_1, \dots, x_m) \in S(A, \mathbf{0}) \cap \mathbb{N}^m$. But then if we define $\boldsymbol{y} = (\boldsymbol{x}^{\bar{Q}}, \boldsymbol{x}^{Q}, \boldsymbol{x}^{\bar{Q}})$, we clearly have $\boldsymbol{y} \in S(A \overset{\mathrm{id}_Q}{\times} A, \mathbf{0}) \cap \mathbb{N}^m$. It remains to be proved that for any two distinct indices $i, j \in [2m - |Q|]$, there exists a solution $(x_1, \dots, x_{2m-|Q|}) = \boldsymbol{x}$ in $S(A \overset{\mathrm{id}_Q}{\times} A, \mathbf{0})$ such that $x_i \neq x_j$. Let $\rho\colon [2m - |Q|] \to [m]$ denote the surjective function that maps the column indices of $A \overset{\mathrm{id}_Q}{\times} A$ to the corresponding column indices of A. If $\rho(i) \neq \rho(j)$, that is, i and j refer to distinct variables of A, this again just follows directly from the fact that

A was positive itself. So we need to check the case $\rho(i) = \rho(j) \in \bar{Q}$. By our assumption on $A^{\bar{Q}}$, there exists a vector $\boldsymbol{y} \in S(A^{\bar{Q}}, \mathbf{0})$ with all its coordinates nonzero. Clearly, the same will hold for $2\boldsymbol{y}$, and hence

$$\boldsymbol{x} = (\boldsymbol{y}, \mathbf{0}, 2\boldsymbol{y}) \in \mathbb{Z}^{2m-|Q|}$$

will satisfy $(A \overset{\mathrm{id}_Q}{\times} A)\boldsymbol{x} = \mathbf{0}$ and the ith and jth entry of $\boldsymbol{x}$ are different. ■

Note that if $A \in \mathbb{Z}^{r\times m}$ is abundant, then any $Q \subset [m]$ of size $|Q| = 1$ will satisfy the requirements of Lemma 1.20. This follows essentially from the same observation that was used in the proof of Lemma 1.17. If after the removal of a single column, there existed a coordinate $i \in [m-1]$ such that x_i had to be 0 in any solution $\mathbf{x} \in S(A^{\bar{Q}}, \mathbf{0})$, then the ith column vector would not lie in the span of the remaining ones, and hence the matrix obtained by removing it would be of lower rank than A, a contradiction to the abundance property. The next result shows that even in the case that Q itself does not satisfy them, we can find a superset that does, which will help us to get universal lower bounds that are sufficient for our applications.

Lemma 1.21. *Let $m > r$ be positive integers, $n \in \mathbb{N}$, $A \in \mathbb{Z}^{r\times m}$ positive, and $\boldsymbol{b} \in \mathbb{Z}^r$ such that $S(A, \boldsymbol{b}) \neq \emptyset$. Then for all $Q \subset [m]$ there exists a $Q' \supset Q$ such that*

$$\left|S_0(A \overset{\mathrm{id}_{Q'}}{\times} A, \boldsymbol{b}^2) \cap [n]^{2m-|Q'|}\right| = \Omega\left(n^{2m-\mathrm{rk}(A)-\mathrm{rk}(A^{\bar{Q}})-|Q|}\right).$$

Proof. If Q satisfies the assumptions of Lemma 1.20, we see that $A \overset{\mathrm{id}_Q}{\times} A$ is positive and hence we can apply the lower bound of Lemma 1.15 directly with $Q' = Q$, noting that $A \overset{\mathrm{id}_Q}{\times} A$ has $2m - |Q|$ columns and rank $\mathrm{rk}(A) + \mathrm{rk}(A^{\bar{Q}})$ by Lemma 1.19. Otherwise, denote by $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_m \in \mathbb{Z}^r$ the columns of A, and let $Q_1 \subset \bar{Q}$ be the index set of the columns that will always have coefficient zero in a linear combination of the column vectors that equals zero. We claim that these columns are linearly independent and their span does not contain any column $\boldsymbol{c}_i$ with $i \in \bar{Q} \setminus Q_1$. Indeed, linear independence holds because any non-trivial linear combination $\sum_{i\in Q_1} \lambda_i \boldsymbol{c}_i = \mathbf{0}$ could be extended to a linear combination $\sum_{i\in\bar{Q}} \lambda_i \boldsymbol{c}_i = \mathbf{0}$ such that $\lambda_i \neq 0$ for at least one $i \in Q_1$, a contradiction to the definition of Q_1. Similarly, if $\sum_{i\in Q_1} \lambda_i \boldsymbol{c}_i = \boldsymbol{c}_j$ for some $j \in \bar{Q} \setminus Q_1$, we obviously see that $\sum_{i\in Q_1} \lambda_i \boldsymbol{c}_i - \boldsymbol{c}_j = \mathbf{0}$ is a linear combination with at least one $\lambda_i \neq 0$, again a contradiction.

We thus see that defining $Q' = Q \cup Q_1$, the matrix $A^{\bar{Q}'}$ has rank $\mathrm{rk}(A^{\bar{Q}'}) = \mathrm{rk}(A^{\bar{Q}}) - |Q_1|$ and its number of columns is $m - |Q'| = m - |Q| - |Q_1|$, and hence

$$m - |Q'| - \mathrm{rk}(A^{\bar{Q}'}) = m - |Q| - \mathrm{rk}(A^{\bar{Q}}).$$

By construction, $A^{\bar{Q}'}$ satisfies the assumption of Lemma 1.20, and hence for $n \in \mathbb{N}$, applying Lemma 1.15 implies

$$\left|S_0(A \overset{\mathrm{id}_{Q'}}{\times} A, \boldsymbol{b}^2) \cap [n]^{2m-|Q'|}\right| = \Omega\left(n^{2m-\mathrm{rk}(A)-\mathrm{rk}(A^{\bar{Q}})-|Q|}\right),$$

which is what we wanted to show. ■

Lemmas 1.19 and 1.20 actually apply to a more general iterated construction. Using the assumptions made in Lemma 1.18, write $Q = \{q_1 < \cdots < q_{|Q|}\}$, and define for any integer $j \geq 1$ the bijection $M_j\colon [j(m-|Q|)+|Q|] \setminus [j(m-|Q|)] \to [|Q|]$ by $M_j(i) = q_{i-j(m-|Q|)}$. Then all the results mentioned in the aforementioned lemmas also apply in a natural way to the matrix

$$A \overset{\mathrm{id}_Q}{\times} A \overset{M_1}{\times} \cdots \overset{M_t}{\times} A = \begin{pmatrix} A^{\bar{Q}} & & & A^{Q} & \\ & \ddots & & \vdots & \\ & & A^{\bar{Q}} & A^{Q} & \\ & & & A^{Q} & A^{\bar{Q}} \end{pmatrix} \tag{1.4}$$

for any $t \geq 1$. The idea being that instead of just having a pair of proper solutions to $A\boldsymbol{x} = \boldsymbol{b}$, we now have a collection of $t+1$ of them that all mutually intersect exactly in the variables indexed by Q. Specifically, we get the following result.

Lemma 1.22. *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$ be positive and abundant and $Q = \{i\} \subset [m]$ of size 1, and for any positive integer j let M_j denote the function $jm - j + 1 \mapsto i$. Then for any positive integer t it holds that the matrix $A \overset{\mathrm{id}_Q}{\times} A \overset{M_1}{\times} \cdots \overset{M_t}{\times} A$ is positive and of rank*

$$\mathrm{rk}\big(A \overset{\mathrm{id}_Q}{\times} A \overset{M_1}{\times} \cdots \overset{M_t}{\times} A\big) = \mathrm{rk}(A) + (t+1)\,\mathrm{rk}(A^{\bar{Q}}).$$

Proof. The proof is identical to that of Lemmas 1.19 and 1.20, noting that since A is abundant, for any $Q \subset [m]$ of size 1 it will hold that $S(A^{\bar{Q}}, \mathbf{0})$ contains a solution with all entries non-zero. ■

1.2 The distribution of type-$\mathfrak{P}$ solutions: the main meta theorem

We now state our main result – which we refer to as the main meta theorem – that will imply Theorems 1.11 and 1.13.

Theorem 1.23. *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$ a positive and abundant integer matrix, $\boldsymbol{b} \in \mathbb{Z}^r$ such that $S(A, \boldsymbol{b}) \neq \emptyset$, and $\mathfrak{P} \subset \mathfrak{P}(A)$ satisfying $(\{1\}, \ldots, \{m\}) \in \mathfrak{P}$ and*

$$\{\mathfrak{p} \in \mathfrak{P} : S_0(A_{\mathfrak{p}}, \boldsymbol{b}) \neq \emptyset\} \subset \mathfrak{P}_0(A).$$

Furthermore, let n be an integer, $0 \leq p := p(n) \leq 1$ and X_n the random variable $|S_{\mathfrak{P}}(A, \boldsymbol{b}) \cap [n]_p^m|$ that counts the number of type–$\mathfrak{P}$ solutions $\boldsymbol{x} \in [n]_p^m$ to $A\boldsymbol{x} = \boldsymbol{b}$. Then

$$\tilde{X}_n \xrightarrow{d} \mathcal{N}(0,1)$$

if $n(1-p) \to \infty$ and $np^{c(A_{\mathfrak{p}})} \to \infty$ for all $\mathfrak{p} \in \mathfrak{P}$ with $S_{\mathfrak{p}}(A, \boldsymbol{b}) \neq \emptyset$.

Let us first see that this indeed implies Theorems 1.11 and 1.13.

Proof of Theorem 1.11. We apply Theorem 1.23 with $\mathfrak{P} = \{(\{1\}, \ldots, \{m\})\}$, noting that the non-emptiness of $S_0(A, \boldsymbol{b})$ follows from the assumptions and Lemma 1.15. ■

Proof of Theorem 1.13. We see that for any $\mathfrak{p} \in \mathfrak{P}(A)$, whenever $S_0(A_{\mathfrak{p}}, \mathbf{0}) \neq \emptyset$, the matrix $A_{\mathfrak{p}}$ is positive by definition, that is $\mathfrak{p} \in \mathfrak{P}_0(A)$. Since A is strictly balanced, it holds that

$$c(A) = \frac{m}{m - \mathrm{rk}(A)} > \max_{\emptyset \subsetneq Q \subsetneq [m]} \frac{|Q|}{|Q| - r_Q(A)} > c(A_{\mathfrak{p}})$$

for any $\mathfrak{p} \in \mathfrak{P}(A)$ with $|\mathfrak{p}| < m$ and hence $np^{c(A_{\mathfrak{p}})} \to \infty$. ■

The proof of Theorem 1.23 involves the analysis of several sub-cases, which we will split into different subsections. In general, it will use the method of moments and follow the ideas that were used by Ruciński in [85] where he proved a similar result to Theorem 1.23 in order to determine the distribution of the number of occurrences of a fixed graph G as a subgraph in a binomial random graph. However, let us stress that our work highly differs from [85] on the more complex algebraic structure of the patterns that are taken into account and in part already presented in the preceding section. A particularly delicate difference between the graph and the system setting is that for systems, the elements of $[n]_p$ take on the role of both vertices and edges when compared with the binomial random graph model $G(n,p)$. This duality results in the additional consideration of partition types that needs to be done for the former case, but not for that of graphs. We proceed now to the proof of the main meta theorem.

1.2.1 The proof of Theorem 1.23

Before going into further case analysis, let us first discuss the general proof strategy. Let $\mu_k = \mathbb{E}((X_n - \mathbb{E}(X_n))^k)$ denote the k-th central moment of X_n associated to the system of linear equations $A\mathbf{x} = \boldsymbol{b}$. Our final goal will always be to show that, independently of the system studied,

$$\mu_k = \begin{cases} (1+o(1))\frac{k!}{(k/2)!}2^{-k/2}\mu_2^{k/2} & \text{if } k \text{ is even,} \\ o(\mu_2^{k/2}) & \text{if } k \text{ is odd.} \end{cases} \tag{1.5}$$

These estimates would show that the moments $\mathbb{E}(\tilde{X}_n^k)$ converge to the moments of a normal distribution, which is uniquely determined by its moments.

Given a solution $\boldsymbol{x} = (x_1, \dots, x_m)$ in $S(A, \boldsymbol{b})$, denote by $\mathbb{I}_{\boldsymbol{x}}$ the indicator random variable for the event $\{x_1, \dots, x_m\} \subset [n]_p$. Abusing notation somewhat, if $\chi = (\boldsymbol{x}_1, \dots, \boldsymbol{x}_k) \in S(A, \boldsymbol{b})^k$ is a k-tuple of solutions, we will write $\{\chi\} = \bigcup_{i=1}^k \{\boldsymbol{x}_i\}$. As a visual shorthand, bold latin letters will indicate solutions, while greek letters will denote tuples of solutions. Note first that by definition we have

$$\mu_k = \sum_{\chi \in \mathfrak{S}_k} \mathbb{E}\left(\prod_{\boldsymbol{x} \in \chi}\left(\mathbb{I}_{\boldsymbol{x}} - p^{|\{\boldsymbol{x}\}|}\right)\right), \tag{1.6}$$

where $\mathfrak{S}_k$ is the set containing all k-tuples $\chi = (\boldsymbol{x}_1, \dots, \boldsymbol{x}_k) \in (S_{\mathfrak{P}}(A, \boldsymbol{b}) \cap [n]^m)^k$ such that for every $i \in [k]$ there exists a $j \in [k] \setminus \{i\}$ such that $\{\boldsymbol{x}_i\} \cap \{\boldsymbol{x}_j\} \neq \emptyset$. Indeed, for any k-tuple χ that contains some $\boldsymbol{x} \in S_{\mathfrak{P}}(A, \boldsymbol{b}) \cap [n]^m$ that is disjoint from the remaining $k-1$ solutions, we have

$$\mathbb{E}\left(\prod_{\boldsymbol{y} \in \chi}\left(\mathbb{I}_{\boldsymbol{y}} - p^{|\{\boldsymbol{y}\}|}\right)\right) = \mathbb{E}\left(\mathbb{I}_{\boldsymbol{x}} - p^{|\{\boldsymbol{x}\}|}\right)\mathbb{E}\left(\prod_{\substack{\boldsymbol{y} \in \chi \\ \boldsymbol{y} \neq \boldsymbol{x}}}\left(\mathbb{I}_{\boldsymbol{y}} - p^{|\{\boldsymbol{y}\}|}\right)\right) = 0.$$

Equation (1.6) behaves slightly different depending on the behavior of p, so suppose $p \to a$ for some constant $a \in [0,1]$. We split the analysis in three cases, depending on wether this limit belongs to $(0,1)$, is equal to 1 and is equal to 0.

Case 1: $0 < a < 1$.

In this case we see that (1.6) implies $\mu_k = \sum_{\mathfrak{S}_k} \Theta(1)$, so we need to analyze the cardinality of $\mathfrak{S}_k$. We are going to prove (1.5) by induction on k. (the base cases $k = 1, 2$ are clearly always true). Note that in particular, this implies that for any $\ell < k$ it holds that $\mu_\ell = O(\mu_2^{\ell/2})$ and hence by our previous observation

$$|\mathfrak{S}_\ell| = O(|\mathfrak{S}_2|^{\ell/2}).$$

Now, let $\mathfrak{S}'_k$ denote the subset of $\mathfrak{S}_k$ containing all k-tuples of solutions $(\boldsymbol{x}_1, \dots, \boldsymbol{x}_k)$ with $\boldsymbol{x}_i \in (S_{\mathfrak{P}}(A, \boldsymbol{b}) \cap [n]^m)$ such that for every i the choice of $j \neq i$ with $\{\boldsymbol{x}_i\} \cap \{\boldsymbol{x}_j\} \neq \emptyset$ is unique, that is, the solutions can be grouped into $k/2$ pairwise disjoint pairs. Note that clearly, $\mathfrak{S}'_k = \emptyset$ for every odd k. If $\overline{\mathfrak{S}}'_k = \mathfrak{S}_k \setminus \mathfrak{S}'_k$, we will show that

$$|\overline{\mathfrak{S}}'_k| = o(|\mathfrak{S}_2|^{k/2}), \tag{1.7}$$

which would imply

$$\sum_{\chi \in \overline{\mathfrak{S}}'_k} \mathbb{E}\left(\prod_{x \in \chi}\left(\mathbb{I}_x - p^{|\{x\}|}\right)\right) = o(\mu_2^{k/2}).$$

To see that this is true, suppose $\chi := (\boldsymbol{x}_1, \dots, \boldsymbol{x}_k) \in \overline{\mathfrak{S}}'_k$. Then there must exist an index $i = i(\chi) \in [k]$ such that

$$\chi^{[k] \setminus \{i\}} = (\boldsymbol{x}_1, \dots, \boldsymbol{x}_{i-1}, \boldsymbol{x}_{i+1}, \dots, \boldsymbol{x}_k)$$

is contained in $\mathfrak{S}_{k-1}$. Taking the minimum choice of i for each χ, we have thus defined a map $\pi\colon \overline{\mathfrak{S}}'_k \to \mathfrak{S}_{k-1}$, and hence

$$|\overline{\mathfrak{S}}'_k| \leq \max_{v \in \mathfrak{S}_{k-1}} |\pi^{-1}(v)| \cdot |\mathfrak{S}_{k-1}|.$$

Since by induction $|\mathfrak{S}_{k-1}| = O(|\mathfrak{S}_2|^{(k-1)/2})$, it suffices to show that $\max_{v \in \mathfrak{S}_{k-1}} |\pi^{-1}(v)| = o(|\mathfrak{S}_2|^{1/2})$. We will first determine a lower bound on $|\mathfrak{S}_2|$. It is clear that $|\mathfrak{S}_2| \geq |S_0(A \overset{\mathrm{id}_Q}{\times} A, \boldsymbol{b}^2) \cap [n]^{2m-|Q|}|$ for any $Q \subset [m]$. Since A is abundant, any Q containing exactly one index will satisfy the conditions of Lemma 1.20, so fix an arbitrary one. It follows by this and Lemma 1.15 applied to $A \overset{\mathrm{id}_Q}{\times} A$ that

$$|\mathfrak{S}_2| \geq |S_0(A \overset{\mathrm{id}_Q}{\times} A, \boldsymbol{b}^2) \cap [n]^{2m-1}| = \Omega(n^{2m-2\,\mathrm{rk}(A)-1}). \tag{1.8}$$

We now turn to giving an upper bound on $\max_{v \in \mathfrak{S}_{k-1}} |\pi^{-1}(v)|$, so fix an $v \in \mathfrak{S}_{k-1}$. By definition of $\mathfrak{S}_k$, any solution that could extend v to a k-tuple in $\mathfrak{S}_k$ must intersect $\{v\}$, and hence by Lemma 1.17 we see that

$$|\pi^{-1}(v)| = O(n^{m-\mathrm{rk}(A)-1}) = o\left(|\mathfrak{S}_2|^{1/2}\right)$$

by the previously obtained lower bound (1.8). Since $\mathfrak{S}'_k = \emptyset$ for odd k, this also immediately proves the odd case of (1.5) in this regime of p.

We now turn to the case of even k, noting that (1.7) tells us that essentially, the tuples that are summed over in μ_k are $k/2$ pairwise disjoint pairs of those summed over in μ_2. We begin by showing that for all but a negligible amount, pairs $(\boldsymbol{x}, \boldsymbol{y}) \in \mathfrak{S}_2$ will satisfy $\boldsymbol{x}, \boldsymbol{y} \in S_0(A, \boldsymbol{b})$ and $|\{\boldsymbol{x}\} \cap \{\boldsymbol{y}\}| = 1$. Let us first see that we can restrict ourselves to pairs of proper solutions. This follows from a similar argument as was used in the proof of Lemma 1.17. Namely, if we fix an arbitrary $\boldsymbol{x} \in S_{\mathfrak{P}}(A, \boldsymbol{b})$, then by (1.2)

$$\left|\{y \in S_0(A_{\mathfrak{p}}, \boldsymbol{b}) \cap [n]^{|\mathfrak{p}|} : \mathfrak{p} \in \mathfrak{P},\, |\mathfrak{p}| < m\} = O(n^{m-\mathrm{rk}(A)-2}).$$

Since there are at most $n^{m-\mathrm{rk}(A)}$ choices for $\boldsymbol{x}$, this is negligible when compared to $|\mathfrak{S}_2|$. Applying the second part of Lemma 1.17 directly, the number of pairs of proper solutions that intersect in at least two elements is negligible as well. Note that for any pair $(\boldsymbol{x}, \boldsymbol{y})$ satisfying the structure described above, we have $\mathbb{E}((\mathbb{I}_{\boldsymbol{x}} - p^{\{\boldsymbol{x}\}})(\mathbb{I}_{\boldsymbol{y}} - p^{\{\boldsymbol{y}\}})) \sim (a^{2m-1}(1-a))$.

Finally, note that for any pair $\chi \in \mathfrak{S}_2$, the number of pairs $v \in \mathfrak{S}_2$ that share elements with χ is negligible: Clearly, any such v will consist of an $\boldsymbol{x} \in S_0(A, \boldsymbol{b})$ satisfying $\{\boldsymbol{x}\} \cap \{\chi\} \neq \emptyset$

and a $\boldsymbol{y} \in S_0(A,\boldsymbol{b})$ satisfying $\{\boldsymbol{y}\} \cap \{\boldsymbol{x}\} \neq \emptyset$. Applying Lemma 1.17 again, the number of such pairs is $O(n^{2m-2\,\mathrm{rk}(A)-2}) = o(|\mathfrak{S}_2|)$. We conclude

$$|\mathfrak{S}_k| \sim \binom{|\mathfrak{S}_2|/2}{k/2} k!$$

and hence, as p tends to $a \in (0,1)$,

$$\mu_k \sim \binom{|\mathfrak{S}_2|/2}{k/2} k!(a^{2m-1}(1-a))^{k/2} \sim \frac{k!}{(k/2)!} 2^{-k/2} \mu_2^{k/2}.$$

■

Case 2: $a = 1$**.**

Recall that p is chosen such that $n(1-p)$ tends to infinity. This property will be especially relevant in this case. Instead of working directly with the random variables defined before, we will instead look at the complements. So if $\boldsymbol{x}$ is a solution to $A\boldsymbol{x} = \boldsymbol{b}$, define $q = 1-p$ and $\bar{\mathbb{I}}_{\boldsymbol{x}} = 1 - \mathbb{I}_{\boldsymbol{x}}$. We see that $\nu_{\boldsymbol{x}} := \mathbb{E}(\bar{\mathbb{I}}_{\boldsymbol{x}}) = 1 - p^{|\{\boldsymbol{x}\}|} \sim |\{\boldsymbol{x}\}|q$. Let $\chi = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k) \in \mathfrak{S}_k$ be a k-tuple of solutions in $S_{\mathfrak{P}}(A,\boldsymbol{b}) \cap [n]^m$. Using these definitions we see that (1.6) can be written as

$$\mu_k = \sum_{\chi \in \mathfrak{S}_k} \mathbb{E}\left(\prod_{\boldsymbol{x} \in \chi} (\mathbb{I}_{\boldsymbol{x}} - p^{|\{\boldsymbol{x}\}|})\right) = (-1)^k \sum_{\chi \in \mathfrak{S}_k} \mathbb{E}\left(\prod_{\boldsymbol{x} \in \chi} (\bar{\mathbb{I}}_{\boldsymbol{x}} - \nu_{\boldsymbol{x}})\right). \tag{1.9}$$

We can associate a set system $\mathcal{H} = \mathcal{H}(\chi)$ to the k-tuple χ. The vertex set $V(\mathcal{H})$ is just $\{\chi\}$, while the edges are $E(\mathcal{H}) = \{\{\boldsymbol{x}\} : \boldsymbol{x} \in \chi\}$. Suppose the vertex cover number $\tau(\mathcal{H})$ of the hypergraph is s, that is, there exist s elements $a_1, \ldots, a_s \in \{\chi\}$ such that removing them destroys all solutions in χ, and no collection of less than s elements in $[n]$ achieves this. Then an straightforward computation shows that

$$\mathbb{E}\left(\prod_{\boldsymbol{x} \in \chi} \bar{\mathbb{I}}_{\boldsymbol{x}}\right) = \Theta(q^s).$$

Furthermore, it is clear that for any $t \in [k-1]$, every $(k-t)$-sub-collection of χ will require removal of at least $s-t$ elements to destroy all solutions in it, that is, for any $I \subset [k]$ with $|I| = k-t$ it holds that

$$\mathbb{E}\left(\prod_{i \in I} \bar{\mathbb{I}}_{\boldsymbol{x}_i} \prod_{j \notin I} \nu_{\boldsymbol{x}_j}\right) = \begin{cases} \Theta(q^s), & \text{if } \chi^I \in \mathfrak{S}_{k-t}, \\ 0, & \text{otherwise.} \end{cases}$$

Putting this together, we see that (1.9) becomes

$$\mu_k = \sum_{s=1}^{k} \sum_{\chi \in \mathfrak{S}_{k,s}} \Theta(q^s), \tag{1.10}$$

where $\mathfrak{S}_{k,s}$ is the subset of k-tuples $\chi \in \mathfrak{S}_k$ such that $\tau(\mathcal{H}(\chi)) = s$. We will now show that for any fixed s, only χ of a certain structure contribute significantly to (1.10). For this, define by $\mathfrak{S}'_{k,s}$ the set of s-*milky ways*, which are $\chi \in \mathfrak{S}_{k,s}$ such that $\mathcal{H}(\chi)$ is the union of s disjoint components, with all edges in a component intersecting in a unique vertex. Note that this set is only non-empty for $s \leq \lfloor k/2 \rfloor$. Furthermore, note that each component corresponds to a matrix as described in (1.4), call it B_i, hence by considering the compounded matrix $B_1 \dot{\times} \cdots \dot{\times} B_s$ we see that Lemma 1.22 together with Lemma 1.15 implies that for any k and $s \leq \lfloor k/2 \rfloor$

$$|\mathfrak{S}'_{k,s}| = \Omega(n^{k(m-\mathrm{rk}(A)-1)+s}). \tag{1.11}$$

The combination of these two lemmas also gives a matching upper bound: For any fixed way of dividing up the k solutions into s components, there are at most $O(n^{k(m-\operatorname{rk}(A)-1)+s})$ corresponding k-tuples. Since k and s are fixed, the number of these partitions is $O(1)$, and hence

$$|\mathfrak{S}'_{k,s}| = O(n^{k(m-\operatorname{rk}(A)-1)+s}). \tag{1.12}$$

We will now show by induction on k that for any k and s,

$$|\overline{\mathfrak{S}}'_{k,s}| := |\mathfrak{S}_{k,s} \setminus \mathfrak{S}'_{k,s}| = O(n^{k(m-\operatorname{rk}(A)-1)+s-1}).$$

This holds trivially for $k = 1$, and for $k = 2$, it follows immediately from Lemma 1.17, since A is abundant and $s = 1$ being the only vertex cover number leading to a non-empty set. So suppose $k \geq 3$ and note that the bound (1.12) together with the induction hypothesis in particular implies $|\mathfrak{S}_{\ell,s}| = O(n^{\ell(m-\operatorname{rk}(A)-1)+s})$ for any $\ell < k$.

We will split up $\overline{\mathfrak{S}}'_{k,s}$ even further, into disjoint sets $\overline{\mathfrak{S}}''_{k,s}$, $\overline{\mathfrak{S}}'''_{k,s}$, and $\overline{\mathfrak{S}}''''_{k,s}$, which are defined as follows:

a) The set $\overline{\mathfrak{S}}''_{k,s}$ will contain all $\chi = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k) \in \overline{\mathfrak{S}}'_{k,s}$ such that there are solutions $\boldsymbol{x}_i, \boldsymbol{x}_j \in \chi$ satisfying $\{\chi^{[k]\setminus\{i,j\}}\} \cap \{\chi^{\{i,j\}}\} = \emptyset$, and $|\{\boldsymbol{x}_i\} \cap \{\boldsymbol{x}_j\}| \geq 2$.

b) The set $\overline{\mathfrak{S}}'''_{k,s}$ contains all $\chi = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k) \in \overline{\mathfrak{S}}'_{k,s} \setminus \overline{\mathfrak{S}}''_{k,s}$, such that there exists a solution $\boldsymbol{x}_i \in \chi$ satisfying $|\{\boldsymbol{x}\} \cap \{\chi^{[k]\setminus\{i\}}\}| \geq 2$ and $\chi^{[k]\setminus\{i\}} \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}$.

c) Finally, $\overline{\mathfrak{S}}''''_{k,s} = \overline{\mathfrak{S}}'_{k,s} \setminus (\overline{\mathfrak{S}}''_{k,s} \cup \overline{\mathfrak{S}}'''_{k,s})$ are the remaining non-milky ways.

We proceed to analyze the cardinality of each set.

a) We start with case $\overline{\mathfrak{S}}''_{k,s}$. For every $\chi \in \overline{\mathfrak{S}}''_{k,s}$, there are indices i and j such that $\chi^{[k]\setminus\{i,j\}}$ is contained in $\mathfrak{S}_{k-2,s-1}$, and taking the lexicographic smallest pair we have defined a map π from $\overline{\mathfrak{S}}''_{k,s}$ to $\mathfrak{S}_{k-2,s-1}$, which implies

$$|\overline{\mathfrak{S}}''_{k,s}| \leq |\mathfrak{S}_{k-2,s-1}| \max_{v \in \mathfrak{S}_{k-2,s-1}} |\pi^{-1}(v)| = O\left(n^{(k-2)(m-\operatorname{rk}(A)-1)+s-1} \max_{v \in \mathfrak{S}_{k-1,s}} |\pi^{-1}(v)|\right).$$

Lemmas 1.15 and 1.17 imply that for any $v \in \mathfrak{S}_{k-2,s-1}$ we have

$$|\pi^{-1}(v)| = O(n^{2m-2\operatorname{rk}(A)-2}),$$

and hence

$$|\overline{\mathfrak{S}}''_{k,s}| = O(n^{k(m-\operatorname{rk}(A)-1)+s-1}). \tag{1.13}$$

b) We continue with the analysis of $\overline{\mathfrak{S}}'''_{k,s}$. In this case, taking the smallest possible index i, this again defines a map $\pi\colon \overline{\mathfrak{S}}'''_{k,s} \to \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}$ and we see by induction hypothesis that

$$\begin{aligned}|\overline{\mathfrak{S}}'''_{k,s}| &\leq |\mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}| \max_{v \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}} |\pi^{-1}(v)| \\ &= O\left(n^{(k-1)(m-\operatorname{rk}(A)-1)+s} \max_{v \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}} |\pi^{-1}(v)|\right).\end{aligned}$$

Again, for any $v \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}$, we see that $|\pi^{-1}(v)|$ is at most the number of solutions $\boldsymbol{y} \in S_{\mathfrak{P}}(A, \boldsymbol{b}) \cap [n]^m$ that contain at least 2 elements from $\{v\}$, which is $O(n^{m-\operatorname{rk}(A)-2})$ by Lemma 1.17, and hence

$$|\overline{\mathfrak{S}}'''_{k,s}| = O(n^{k(m-\operatorname{rk}(A)-1)+s-1}). \tag{1.14}$$

c) Finally, it remains to look at $\overline{\mathfrak{S}}''''_{k,s}$. If $\chi = (\boldsymbol{x}_1, \dots, \boldsymbol{x}_k)$ in $\overline{\mathfrak{S}}''''_{k,s}$, we claim that there must exist an index i such that $\chi^{[k]\setminus\{i\}} \in \mathfrak{S}_{k-1,s-1}$. In the sequel, we will only consider the components of $\mathcal{H}(\chi)$ that do not intersect in a unique vertex, of which there is at least one since $\chi \notin \mathfrak{S}'_{k,s}$. First note that it is clear that since $\chi \notin \overline{\mathfrak{S}}'_{k,s} \cup \overline{\mathfrak{S}}''_{k,s}$, there must exist an index i such that $\chi^{[k]\setminus\{i\}} \in \mathfrak{S}_{k-1,s-1} \cup \mathfrak{S}_{k-1,s}$ or in other words, the remaining solutions are still intersecting. Since $\chi \notin \overline{\mathfrak{S}}'''_{k,s}$, for all of these indices, it must hold that $|\{\boldsymbol{x}_i\} \cap \{\chi^{[k]\setminus\{i\}}\}| = 1$. Finally, for at least one index i of the previously considered, it must actually hold that there is a unique j_i such that

$$\{\boldsymbol{x}_i\} \cap \{\chi^{[k]\setminus\{i\}}\} = \{\boldsymbol{x}_i\} \cap \{\boldsymbol{x}_{j_i}\}. \tag{1.15}$$

Indeed, since the components we consider are not sunflowers, there must exist an $\boldsymbol{x}_\ell$ that intersects the rest of its component in at least two points, and so the negation of (1.15) would imply $\chi \in \overline{\mathfrak{S}}'''_{k,s}$, since $\boldsymbol{x}_\ell$ would be a valid choice. We see that for any valid i satisfying (1.15) it holds that $\chi^{[k]\setminus\{i\}} \in \mathfrak{S}_{k-1,s-1}$. Having established this, we repeat the arguments already used, taking the minimal valid i and defining the appropriate function $\pi\colon \overline{\mathfrak{S}}''''_{k,s} \to \mathfrak{S}_{k-1,s-1}$ and conclude

$$\begin{aligned}
|\overline{\mathfrak{S}}''''_{k,s}| &\leq |\mathfrak{S}_{k-1,s-1}| \max_{v \in \mathfrak{S}_{k-1,s-1}} |\pi^{-1}(v)| \\
&= O(n^{(k-1)(m-\mathrm{rk}(A)-1)+s-1}) \max_{v \in \mathfrak{S}_{k-1,s-1}} |\pi^{-1}(v)| \\
&= O(n^{k(m-\mathrm{rk}(A)-1)+s-1}).
\end{aligned}$$

Together with (1.11), we see that (1.13), (1.14) and (1.15) imply that only s-milky ways contribute meaningfully when $s \leq \lfloor k/2 \rfloor$. Sadly, this bound is not quite strong enough when $s > \lfloor k/2 \rfloor$, since s-milky ways do not exist here. For this, we will prove by induction on k that for any k and $s \geq \lfloor k/2 \rfloor$ it holds that

$$|\mathfrak{S}_{k,s}| = O(n^{k(m-\mathrm{rk}(A)-1)+\lfloor k/2 \rfloor}). \tag{1.16}$$

Again, the cases $k = 1$ and $k = 2$ follow from previous observations. Moreover, the statement is trivially true for any k and $s = \lfloor k/2 \rfloor$, so suppose $k \geq 3$ and $s > \lfloor k/2 \rfloor$. If $\chi = (\boldsymbol{x}_1, \dots, \boldsymbol{x}_k) \in \mathfrak{S}_{k,s}$, then there must exist a least index i such that $\chi^{[k]\setminus\{i\}} \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}$, which allows us again to define a projection map $\pi\colon \mathfrak{S}_{k,s} \to \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}$. Since $s \geq \lfloor k/2 \rfloor + 1$ we have $s - 1 \geq \lfloor (k-1)/2 \rfloor$, and hence by induction hypothesis we have

$$\begin{aligned}
|\mathfrak{S}_{k,s}| &\leq |\mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}| \max_{v \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}} |\pi^{-1}(v)| \\
&= O(n^{(k-1)(m-\mathrm{rk}(A)-1)+\lfloor (k-1)/2 \rfloor})| \max_{v \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}} |\pi^{-1}(v)| \\
&= O(n^{k(m-\mathrm{rk}(A)-1)+\lfloor (k-1)/2 \rfloor}),
\end{aligned}$$

which in particular implies (1.16). Since $q \to 0$, we thus have

$$|\mathfrak{S}_{k,s}| q^s = o\big(|\mathfrak{S}_{k,\lfloor k/2 \rfloor}| q^{\lfloor k/2 \rfloor}\big),$$

and so the terms with $s > \lfloor k/2 \rfloor$ in (1.10) can be discarded. Now, if $\chi \in \mathfrak{S}'_{k,s}$ is a milky way, we see that

$$\mathbb{E}\left(\prod_{\boldsymbol{x} \in \chi} \mathbb{I}_{\boldsymbol{x}} \right) \sim q^s,$$

while for any $I \in [k]$ with $|I| \leq k-1$ we have

$$\mathbb{E}\left(\prod_{i\in I}\bar{\mathbb{I}}_{x_i}\prod_{j\notin I}\nu_{x_j}\right) = O(q^{s+1}) = o(q^s),$$

since the removal of any single solution does not impact the cover number, while any subsequent removal can reduce it by at most one. Furthermore, by our previous observations we have already seen that for any k and $s \leq \lfloor k/2 \rfloor$

$$|\mathfrak{S}'_{k,s}|q^s = \Theta(n^{k(m-\mathrm{rk}(A)-1)}(nq)^s),$$

so since $nq \to \infty$ by assumption only the $s = \lfloor k/2 \rfloor$ term contributes significantly. Hence (1.10) can be rewritten as

$$\mu_k \sim \sum_{s=1}^{\lfloor k/2 \rfloor} |\mathfrak{S}'_{k,s}|q^s \sim |\mathfrak{S}'_{k,\lfloor k/2 \rfloor}|q^{\lfloor k/2 \rfloor}. \tag{1.17}$$

This proves (1.5) for odd k, since

$$\mu_k = O(n^{k(m-\mathrm{rk}(A)-1)}(nq)^{(k-1)/2}) = o(n^{k(m-\mathrm{rk}(A)-1)}(nq)^{k/2}),$$

while

$$\mu_2^{k/2} = \Omega(n^{k(m-\mathrm{rk}(A)-1)}(nq)^{k/2}).$$

For the even case, it is easy to see that one can repeat the argument from Case 1 to show that only those milky ways with all solutions being proper contribute meaningfully and that among those, the number of pairs in $\mathfrak{S}'_{2,1}$ that intersect another pair is negligible, so just like in that case we conclude

$$\mu_k \sim \binom{|\mathfrak{S}'_{2,1}|/2}{k/2} k! q^{k/2} \sim \frac{k!}{(k/2)!}2^{-k/2}\mu_2^{k/2}.$$

■

Case 3: $a = 0$.

Finally, in this case the crucial property of p that will be used is the fact that $np^{c(A_\mathfrak{p})}$ tends to infinity for every considered partition type $\mathfrak{p} \in \mathfrak{P}$ with non-empty solution set. Essentially, this means that the expected number of solutions $\mathbf{x} \in [n]_p^{|\mathfrak{p}|}$ to the densest sub-system of each system $A_\mathfrak{p}\mathbf{x} = \mathbf{b}$ is infinite.

Let $\chi = (x_1, \dots, x_k) \in \mathfrak{S}_k$. Then for any $I \subsetneq [k]$ we see that

$$\mathbb{E}\left(\prod_{i\in I}\mathbb{I}_{x_i}\prod_{j\in[k]\setminus I}p^{|\{x_j\}|}\right) = p^{|\cup_{i\in I}\{x_i\}|+\sum_{j\in[k]\setminus I}|\{x_j\}|} = o\left(p^{|\cup_{i=1}^k\{x_i\}|}\right),$$

since every x_j with $j \notin I$ intersects at least one other solution. Since we clearly have $\mathbb{E}(\mathbb{I}_{x_1}\cdots\mathbb{I}_{x_k}) = p^{|\cup_{i=1}^k\{x_i\}|}$, it thus follows that

$$\mathbb{E}\left((\mathbb{I}_{x_1} - p^{|\{x_1\}|})\cdots(\mathbb{I}_{x_k} - p^{|\{x_k\}|})\right) \sim p^{|\cup_{i=1}^k\{x_i\}|}. \tag{1.18}$$

Furthermore, since $np^{c(A_\mathfrak{p})} \to \infty$ for any $\mathfrak{p} \in \mathfrak{P}$, wee see that for any nonempty $Q \subset [|\mathfrak{p}|]$, it holds that

$$n^{|Q|-r_Q(A_\mathfrak{p})}p^{|Q|} = \omega(1). \tag{1.19}$$

Let us prove (1.5) via induction on k, the cases $k = 1$ and $k = 2$ clearly being true. Note that in particular, the induction hypothesis implies that for any $\ell < k$ it holds that $\mu_\ell = O(\mu_2^{\ell/2})$. Let us decompose μ_k as

$$\mu_k \sim \sum_{\chi \in \mathfrak{S}_k} \mathbb{E}\left(\prod_{\boldsymbol{x} \in \chi} \mathbb{I}_{\boldsymbol{x}}\right) = \sum_{\chi \in \mathfrak{S}'_k} \mathbb{E}\left(\prod_{\boldsymbol{x} \in \chi} \mathbb{I}_{\boldsymbol{x}}\right) + B_k.$$

Recall that $\mathfrak{S}'_k$ denoted the subset of $\mathfrak{S}_k$ such that every $\boldsymbol{x} \in \chi$ had a unique partner $\boldsymbol{y} \in \chi$ and was disjoint from all other components. Our first step will be to show that

$$B_k = o(\mu_2^{k/2}),$$

which would in particular imply (1.5) in the case of odd k, since here $\mathfrak{S}'_k = \emptyset$. To see this, note that $\chi = (\boldsymbol{x}_1, \dots, \boldsymbol{x}_k) \in \mathfrak{S}_k \setminus \mathfrak{S}'_k$ implies that there must exist an index $i \in [k]$ such that $\chi^{[k]\setminus\{i\}} \in \mathfrak{S}_{k-1}$. Let i_χ denote the least index $i \in [k]$ for which this is true, then we can define the map $\pi \colon \mathfrak{S}_k \setminus \mathfrak{S}'_k \to \mathfrak{S}_{k-1}$ by $\pi(\chi) = \chi^{[k]\setminus\{i_\chi\}}$. Furthermore, let $Q_\chi \subset [|\mathfrak{p}(\boldsymbol{x}_{i_\chi})|]$ denote the index set of all the components of $\boldsymbol{x}_\chi$ that are contained in $\{\pi(\chi)\}$. Then

$$\begin{aligned}
B_k &= \sum_{\chi \in \mathfrak{S}_k \setminus \mathfrak{S}'_k} \mathbb{E}\left(\prod_{\boldsymbol{x} \in \chi} \mathbb{I}_{\boldsymbol{x}}\right) \\
&= \sum_{\chi \in \mathfrak{S}_k \setminus \mathfrak{S}'_k} \mathbb{E}\left(\prod_{\boldsymbol{x} \in \pi(\chi)} \mathbb{I}_{\boldsymbol{x}}\right) p^{|\mathfrak{p}(\boldsymbol{x}_{i_\chi})| - |Q_\chi|} \\
&\leq \sum_{v \in \mathfrak{S}_{k-1}} \mathbb{E}\left(\prod_{\boldsymbol{y} \in v} \mathbb{I}_{\boldsymbol{y}}\right) \sum_{\mathfrak{p} \in \mathfrak{P}} \sum_{Q \subset [|\mathfrak{p}|]} \sum_{\boldsymbol{z} \in \{v\}^{|Q|}} p^{|\mathfrak{p}| - |Q|} \left| S_0(A_{\mathfrak{p}}^{\bar{Q}}, b - A_{\mathfrak{p}}^{Q} \boldsymbol{z}) \cap [n]^{|\mathfrak{p}| - |Q|} \right| \qquad (1.20) \\
&\leq \mu_{k-1} \sum_{\mathfrak{p} \in \mathfrak{P}} \sum_{Q \subset [|\mathfrak{p}|]} n^{|\mathfrak{p}| - |Q| - \operatorname{rk}(A_{\mathfrak{p}}^{\bar{Q}})} p^{|\mathfrak{p}| - |Q|} \\
&= \mu_{k-1} \sum_{\mathfrak{p} \in \mathfrak{P}} \sum_{Q \subset [|\mathfrak{p}|]} n^{|\mathfrak{p}| - \operatorname{rk}(A) - (|Q| - r_Q(A_{\mathfrak{p}}))} p^{|\mathfrak{p}| - |Q|}
\end{aligned}$$

Since $\mu_{k-1} = O(\mu_2^{(k-1)/2})$ it thus suffices to show that the remaining expression in (1.20) is $o(\sqrt{\mu_2})$. But by Lemma 1.21 there exists a $Q' \supset Q$ such that, using $p \to 0$ we see that

$$\begin{aligned}
\mu_2 &\geq p^{2|\mathfrak{p}| - |Q'|} \left| S_0(A_{\mathfrak{p}} \overset{\mathrm{id}_{Q'}}{\times} A_{\mathfrak{p}}, (\boldsymbol{b}, \boldsymbol{b})) \cap [n]^{2|\mathfrak{p}| - |Q'|} \right| \\
&\geq p^{2|\mathfrak{p}| - |Q|} \left| S_0(A_{\mathfrak{p}} \overset{\mathrm{id}_{Q'}}{\times} A_{\mathfrak{p}}, (\boldsymbol{b}, \boldsymbol{b})) \cap [n]^{2|\mathfrak{p}| - |Q'|} \right| \\
&= \Omega(n^{2|\mathfrak{p}| - 2\operatorname{rk}(A) - (|Q| - r_Q(A_{\mathfrak{p}}))} p^{2|\mathfrak{p}| - |Q|}),
\end{aligned}$$

and hence

$$\sqrt{\mu_2} = \Omega(n^{|\mathfrak{p}| - \operatorname{rk}(A) - (|Q| - r_Q(A_{\mathfrak{p}}))/2} p^{|\mathfrak{p}| - |Q|/2}) = \omega(n^{|\mathfrak{p}| - \operatorname{rk}(A) - (|Q| - r_Q(A_{\mathfrak{p}}))} p^{|\mathfrak{p}| - |Q|})$$

which follows from (1.19). As stated before, this finishes the proof in the case of odd k, so suppose k is even. If $\mathfrak{p}, \mathfrak{q} \in \mathfrak{P}$ and $M \colon P \to Q$ is a bijection between nonempty $P \subset [|\mathfrak{p}|]$ and $Q \subset [|\mathfrak{q}|]$, we will call the triple $(\mathfrak{p}, \mathfrak{q}, M)$ *leading* if

$$p^{|\mathfrak{p}| + |\mathfrak{q}| - |P|} \left| S_0(A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}}, (\boldsymbol{b}, \boldsymbol{b})) \cap [n]^{|\mathfrak{p}| + |\mathfrak{q}| - |P|} \right| = \Omega(\mu_2),$$

and hence

$$\mu_2 \sim \sum_{\substack{(\mathfrak{p}, \mathfrak{q}, M) \\ \text{leading}}} p^{|\mathfrak{p}| + |\mathfrak{q}| - |\operatorname{dom}(M)|} \left| S_0(A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}}, (\boldsymbol{b}, \boldsymbol{b})) \cap [n]^{|\mathfrak{p}| + |\mathfrak{q}| - |\operatorname{dom}(M)|} \right|.$$

Let us first make an observation that will be helpful later. Suppose $|\mathfrak{p}| \geq |\mathfrak{q}|$, and let $M\colon P \to Q$ be a bijection between some index sets $P \subset [|\mathfrak{p}|]$ and $Q \subset [|\mathfrak{q}|]$. Then by Lemma 1.21 there exists a $P' \supset P$ such that using $np \to \infty$ and $p \to 0$ we see

$$\begin{aligned}
\mu_2 &= \Omega\left(p^{2|\mathfrak{p}|-|P'|}\left|S_0\left(A_{\mathfrak{p}} \overset{\mathrm{id}_{P'}}{\times} A_{\mathfrak{p}}, (\boldsymbol{b},\boldsymbol{b})\right) \cap [n]^{2|\mathfrak{p}|-|P'|}\right|\right) \\
&= \Omega\left(p^{2|\mathfrak{p}|-|P|}\left|S_0\left(A_{\mathfrak{p}} \overset{\mathrm{id}_{P'}}{\times} A_{\mathfrak{p}}, (\boldsymbol{b},\boldsymbol{b})\right) \cap [n]^{2|\mathfrak{p}|-|P'|}\right|\right) \\
&= \Omega\left(n^{2|\mathfrak{p}|-2\operatorname{rk}(A)-(|P|-r_P(A_{\mathfrak{p}}))}p^{2|\mathfrak{p}|-|P|}\right) \qquad (1.21) \\
&= \Omega\left(n^{|\mathfrak{p}|+|\mathfrak{q}|-2\operatorname{rk}(A)-(|P|-r_P(A_{\mathfrak{p}}))}p^{|\mathfrak{p}|+|\mathfrak{q}|-|P|}\right) \\
&= \Omega\left(p^{|\mathfrak{p}|+|\mathfrak{q}|-|P|}\left|S_0\left(A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}}, (\boldsymbol{b},\boldsymbol{b})\right) \cap [n]^{2|\mathfrak{p}|-|P|}\right|\right).
\end{aligned}$$

Here the last line followed from the fact that $\operatorname{rk}(A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}}) \geq \operatorname{rk}(A_{\mathfrak{q}}) + \operatorname{rk}(A_{\mathfrak{p}}^{\overline{P}}) = 2\operatorname{rk}(A) - r_P(A_{\mathfrak{p}})$. Hence, $(\mathfrak{p}, \mathfrak{q}, M)$ can be a leading overlap only if

i) $|\mathfrak{p}| = |\mathfrak{q}|$,

ii) $r_P(A_{\mathfrak{p}}) = r_Q(A_{\mathfrak{q}})$,

iii) $\operatorname{rk}(A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}}) = 2\operatorname{rk}(A) - r_P(A_{\mathfrak{p}})$,

iv) For any matrix $C \in \left\{A_{\mathfrak{p}} \overset{\mathrm{id}_P}{\times} A_{\mathfrak{p}}, A_{\mathfrak{q}} \overset{\mathrm{id}_Q}{\times} A_{\mathfrak{q}}, A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}}\right\}$ it holds that

$$|S_0(C \cap [n]^{2|\mathfrak{p}|-|P|})| = \Theta(n^{2|\mathfrak{p}|-(|P|-r_P(A_{\mathfrak{p}}))}),$$

and

v) $(\mathfrak{p}, \mathfrak{p}, \mathrm{id}_P)$ and $(\mathfrak{q}, \mathfrak{q}, \mathrm{id}_Q)$ are leading overlaps,

since otherwise some of the Ωs in (1.21) would turn into ωs. Note that these things in particular imply that there cannot exist any $P' \supsetneq P$ satisfying

$$|S_0(A_{\mathfrak{p}} \overset{\mathrm{id}_{P'}}{\times} A_{\mathfrak{p}}, \boldsymbol{b}^2) \cap [n]^{2|\mathfrak{p}|-|P'|})| = \Omega(n^{2|\mathfrak{p}|-(|P|-r_P(A_{\mathfrak{p}}))}),$$

since $p \to 0$ would then imply that $(\mathfrak{p}, \mathfrak{p}, \mathrm{id}_P)$ is not a leading triple. By the proof of Lemma 1.21 this means that there does not exist any index $i \in [|\mathfrak{p}|]$ such that $x_i = 0$ for every solution $\boldsymbol{x} \in S(A_{\mathfrak{p}}, \boldsymbol{0})$. Before continuing, let us quickly try to understand the preceding statements. Essentially, one would naively hope that all leading triples are of the form $(\mathfrak{p}, \mathfrak{p}, \mathrm{id}_P)$ for some $P \subset [|\mathfrak{p}|]$, but this is too optimistic because it ignores the inherent symmetry of some systems of linear equations. For instance, it is clear that for $A = \begin{pmatrix} 1 & 1 & 1 & -1 \end{pmatrix}$, the partitions $\mathfrak{p} = (\{1,2\},\{3\},\{4\})$ and $\mathfrak{q} = (\{1,3\},\{2\},\{4\})$ are technically different, but essentially the same in the sense that there exists a bijection between $S_0(A_{\mathfrak{p}}, \boldsymbol{b})$ and $S_0(A_{\mathfrak{q}}, \boldsymbol{b})$.

We can now continue with the actual proof by defining $\mathfrak{S}_k'' \subset \mathfrak{S}_k'$ to be the set of k-tuples $\chi = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k)$ such that whenever $1 \leq i < j \leq k$ and $|\{\boldsymbol{x}_i\} \cap \{\boldsymbol{x}_j\}| = s > 0$, it holds that $(\mathfrak{p}(\boldsymbol{x}_i), \mathfrak{p}(\boldsymbol{x}_j), M_{i,j})$ is a leading triple. Here $M_{i,j}$ is the bijection defining the incidences between the distinct components of $\boldsymbol{x}_i$ and $\boldsymbol{x}_j$. We will show that

$$\sum_{\chi \in \mathfrak{S}_k' \setminus \mathfrak{S}_k''} \mathbb{E}\left(\prod_{\boldsymbol{x} \in \chi} \mathbb{I}_{\boldsymbol{x}}\right) = o(\mu_2^{k/2}).$$

To see this, we will again define a map $\pi\colon \mathfrak{S}'_k \setminus \mathfrak{S}''_k \to \mathfrak{S}'_{k-2}$ in the following way. Among all $1 \leq i < j \leq k$ such that $(\mathfrak{p}(\boldsymbol{x}_i), \mathfrak{p}(\boldsymbol{x}_j), M_{i,j})$ is not a leading triple, let $\{i_\chi, j_\chi\}$ be the set minimizing $\min\{i, j\}$. Then we can define $\pi(\chi) = \chi^{[k]\setminus\{i_\chi, j_\chi\}}$ and see that

$$\begin{aligned}\sum_{\chi \in \mathfrak{S}'_k \setminus \mathfrak{S}''_k} \mathbb{E}\left(\prod_{\boldsymbol{x}\in\chi} \mathbb{I}_{\boldsymbol{x}}\right) &= \sum_{\chi \in \mathfrak{S}'_k \setminus \mathfrak{S}''_k} \mathbb{E}\left(\prod_{\boldsymbol{x}\in\pi(\chi)} \mathbb{I}_{\boldsymbol{x}}\right) p^{|\mathfrak{p}(\boldsymbol{x}_{i_\chi})|+|\mathfrak{p}(\boldsymbol{x}_{j_\chi})|-|\operatorname{dom}(M_{i_\chi,j_\chi})|} \\ &\leq \mu_{k-2} \sum_{\substack{(\mathfrak{p},\mathfrak{q},M)\\ \text{not leading}}} p^{|\mathfrak{p}|+|\mathfrak{q}|-|\operatorname{dom}(M)|} \left|S_0(A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}}, (\boldsymbol{b}, \boldsymbol{b})) \cap [n]^{|\mathfrak{p}|+|\mathfrak{q}|-|\operatorname{dom}(M)|}\right| \\ &= o(\mu_2^{k/2}),\end{aligned}$$

the last line following from the induction hypothesis $\mu_{k-2} = O(\mu_2^{(k-2)/2})$ and the definition of leading triples, noting that there are only $O(1)$ choices for $\mathfrak{p}$, $\mathfrak{q}$ and M. Let us enumerate the set of leading triples by $(\mathfrak{p}_1, \mathfrak{q}_1, M_1), \ldots, (\mathfrak{p}_u, \mathfrak{q}_u, M_u)$, and for any choice of $0 \leq i_1, \ldots, i_u \leq k/2$ such that $\sum i_j = k/2$, define the matrix $A(i_1, \ldots, i_u)$ by

$$A(i_1,\ldots,i_u) = \underbrace{(A_{\mathfrak{p}_1} \overset{M_1}{\times} A_{\mathfrak{q}_1}) \dot{\times} \cdots \dot{\times} (A_{\mathfrak{p}_1} \overset{M_1}{\times} A_{\mathfrak{q}_1})}_{i_1 \text{ times}} \dot{\times} \cdots \dot{\times} \underbrace{(A_{\mathfrak{p}_u} \overset{M_u}{\times} A_{\mathfrak{q}_u}) \dot{\times} \cdots \dot{\times} (A_{\mathfrak{p}_u} \overset{M_u}{\times} A_{\mathfrak{q}_u})}_{i_u \text{ times}},$$

and write ℓ_j for the expression $|\mathfrak{p}_j| + |\mathfrak{q}_j| - |\operatorname{dom}(M_j)|$. We also need the following result.

Claim. *For any integers $0 \leq i_1, \ldots, i_u \leq k/2$ satisfying $\sum i_j = k/2$ it holds that*

$$|S_0(A(i_1,\ldots,i_u), \boldsymbol{b}^k) \cap [n]^{i_1\ell_1+\cdots+i_u\ell_u}| \sim \prod_{j\in[u]} |S_0(A_{\mathfrak{p}_j} \overset{M_j}{\times} A_{\mathfrak{q}_j}, \boldsymbol{b}^2) \cap [n]^{\ell_j}|^{i_j}.$$

Proof. Let $(\mathfrak{p}, \mathfrak{q}, M)$ be a leading triple with $P := \operatorname{dom}(M)$. We will first show that $A_{\mathfrak{p}}^{\overline{P}}$ does not contain any column c_i such that

$$\operatorname{rk}(A_{\mathfrak{p}}^{\overline{P}\setminus\{i\}}) = \operatorname{rk}(A_{\mathfrak{p}}^{\overline{P}}) - 1.$$

But this follows since for any such hypothetical column, every solution $\boldsymbol{x} \in S(A_{\mathfrak{p}}^{\overline{P}}, \boldsymbol{0})$ would satisfy $x_i = 0$, and we have seen before that this cannot happen for any leading triple. We are now able to prove the claim. Consider any $k/2$-tuple χ of solutions in $\bigtimes_{j=1}^{u} \big(S_0(A_{\mathfrak{p}_j} \overset{M_j}{\times} A_{\mathfrak{q}_j}, \boldsymbol{b}^2) \cap [n]^{\ell_j}\big)^{i_j}$, of which there are $\prod_{j\in[u]} |S_0(A_{\mathfrak{p}_j} \overset{M_j}{\times} A_{\mathfrak{q}_j}, \boldsymbol{b}^2) \cap [n]^{\ell_j}|^{i_j}$ many. When these solutions are all pairwise disjoint, they in fact define a solution in $S_0(A(i_1, \ldots, i_u), \boldsymbol{b}^k) \cap [n]^{\ell_1 i_1+\cdots+\ell_u i_u}$. Otherwise, there exist indices $s, t \in [u]$ and a bijection $M\colon P \to Q$ with $\emptyset \neq P \subset [\ell_s]$ and $\emptyset \neq Q \subset [\ell_t]$ such that some pair $\boldsymbol{x}, \boldsymbol{y} \in \chi$ defines a solution in $S_0\big((A_{\mathfrak{p}_s} \overset{M_s}{\times} A_{\mathfrak{q}_s}) \overset{M}{\times} (A_{\mathfrak{p}_t} \overset{M_t}{\times} A_{\mathfrak{q}_t}), \boldsymbol{b}^4\big) \cap [n]^{\ell_s+\ell_t-|P|}$. But if we define $Q_1 = Q \cap [|\mathfrak{p}_t|]$ and $Q_2 = Q \cap [\ell_t] \setminus [|\mathfrak{p}_t|]$, we see that

$$\begin{aligned}\operatorname{rk}\big((A_{\mathfrak{p}_s} \overset{M_s}{\times} A_{\mathfrak{q}_s}) \overset{M}{\times} (A_{\mathfrak{p}_t} \overset{M_t}{\times} A_{\mathfrak{q}_t})\big) &\geq \operatorname{rk}\big(A_{\mathfrak{p}_s} \overset{M_s}{\times} A_{\mathfrak{q}_s}\big) + \operatorname{rk}\big((A_{\mathfrak{p}_t} \overset{M_t}{\times} A_{\mathfrak{q}_t})^{\overline{Q}}\big) \\ &\geq \operatorname{rk}\big(A_{\mathfrak{p}_s} \overset{M_s}{\times} A_{\mathfrak{q}_s}\big) + \operatorname{rk}(A_{\mathfrak{p}_t}^{\overline{Q}_1}) + \operatorname{rk}(A_{\mathfrak{q}_t}^{\overline{\operatorname{im}(M_t)}\setminus Q_2}) \\ &\geq \operatorname{rk}\big(A_{\mathfrak{p}_s} \overset{M_s}{\times} A_{\mathfrak{q}_s}\big) + \operatorname{rk}(A_{\mathfrak{p}_t}) - |Q_1| + 1 + \operatorname{rk} A_{\mathfrak{q}_t}^{\overline{\operatorname{im}(M_t)}} - |Q_2| + 1 \\ &= \operatorname{rk}\big(A_{\mathfrak{p}_s} \overset{M_s}{\times} A_{\mathfrak{q}_s}\big) + 2\operatorname{rk}(A) - |Q| + 2 - r_{\operatorname{im}(M_t)}(A_{\mathfrak{q}_t}).\end{aligned}$$

Here, for the third inequality we have used the previously established fact that removing any one column from $A_{\mathfrak{q}_t}^{\overline{\operatorname{im}(M_t)}}$ will not reduce the rank. The same clearly holds for $A_{\mathfrak{p}_t}$ since this matrix is positive. We are almost done, since now we can use the upper bound from Lemma 1.15 and see that

$$\begin{aligned}
&\left|S_0\big((A_{\mathfrak{p}_s} \overset{M_s}{\times} A_{\mathfrak{q}_s}) \overset{M}{\times} (A_{\mathfrak{p}_t} \overset{M_t}{\times} A_{\mathfrak{q}_t}), \boldsymbol{b}^4\big) \cap [n]^{\ell_s+\ell_t-|P|}\right| \\
&\le n^{\ell_s+\ell_t-|Q|-\mathrm{rk}\left((A_{\mathfrak{p}_s} \overset{M_s}{\times} A_{\mathfrak{q}_s}) \overset{M}{\times} (A_{\mathfrak{p}_t} \overset{M_t}{\times} A_{\mathfrak{q}_t})\right)} \\
&\le n^{\ell_s+\ell_t-\mathrm{rk}(A_{\mathfrak{p}_s} \overset{M_s}{\times} A_{\mathfrak{q}_s})-2\,\mathrm{rk}(A)+r_{\mathrm{im}(M_t)}(A_{\mathfrak{q}_t})} \\
&= o\big(|S_0(A_{\mathfrak{p}_s} \overset{M_s}{\times} A_{\mathfrak{q}_s}, \boldsymbol{b}^2) \cap [n]^{\ell_s}| \cdot |S_0(A_{\mathfrak{p}_t} \overset{M_t}{\times} A_{\mathfrak{q}_t}, \boldsymbol{b}^2) \cap [n]^{\ell_t}|\big),
\end{aligned}$$

the last line following from the observations made from (1.21). Since there are only $O(1)$ many choices for s, t, and M, this proves the claim. ■

Using the claim, we can now conclude

$$\begin{aligned}
\mu_k &\sim \frac{k!}{2^{k/2}} \sum_{\substack{i_1,\dots,i_u \\ \sum i_j = k/2}} \frac{1}{i_1! \cdots i_u!} p^{i_1\ell_1+\cdots+i_u\ell_u} |S_0(A(i_1,\dots,i_u), \boldsymbol{b}^k) \cap [n]^{i_1\ell_1+\cdots+i_u\ell_u}| \\
&\sim \frac{k!}{(k/2)!2^{k/2}} \sum_{\substack{i_1,\dots,i_u \\ \sum i_j = k/2}} \binom{k/2}{i_1,\dots,i_u} \prod_{j\in[u]} (p^{\ell_j}|S_0(A_{\mathfrak{p}_j} \overset{M_j}{\times} A_{\mathfrak{q}_j}, \boldsymbol{b}^2) \cap [n]^{\ell_j}|)^{i_j} \\
&\sim \frac{k!}{(k/2)!2^{k/2}} \mu_2^{k/2}.
\end{aligned}$$

which is what we wanted to prove. ■

1.3 Concluding remarks

Theorem 1.23 establishes sufficient conditions for the choice of p in order to guarantee normal limiting distributions for the number of solutions to linear systems of equations of the form $Ax = \boldsymbol{b}$. Specifically, we showed that $n(1-p) \to \infty$ and $np^{c(A_{\mathfrak{p}})} \to \infty$ (for all partitions under consideration) sufficed, and we used them in their full strength in Cases 2 and 3, respectively of the proof of Theorem 1.23. Both of these conditions are analogous to those that appear in Ruciński's proof of normality for the number of copies of a given subgraph H in the binomial random model $G(n,p)$. Note that when comparing to the graph setting, the elements of $[n]$ in systems take on the function of both vertices and edges, and hence the analogue of our $n(1-p) \to \infty$ requirement is to ask that $n^2(1-p)$, the expected number of edges, is unbounded.

In fact, in [85] Ruciński showed that those conditions were necessary as well as sufficient. In our context we can say something similar regarding the condition that $n(1-p)$ is unbounded. Observe that the expression $n(1-p)$ can be interpreted as the expected number of elements that are *not* chosen in $[n]_p$, hence if $n(1-p) \not\to \infty$, then $[n]_p$ is typically all the interval $[n]$ with the exception of a bounded number of elements. Then it is easy to show that under this condition, X_n is strongly concentrated around its mean value, concluding that $\tilde{X}_n \xrightarrow{d} 0$, and hence, we do not have a normal limiting distribution for the number of solutions.

However, the argument used by Ruciński in order to study the second condition requires a delicate study of moments of order 4 and 6, which we are unable to adapt. Roughly speaking, in our algebraic setting the structure of solutions is more complicated compared with the graph setting, as solutions with repeated components can be valid ones (the analogy would be to consider also subgraphs of the fixed graph H that we want to count the corresponding number of subcopies). Hence, an open question arising from our work is to obtain an *only if* statement of Theorem 1.23. As an intermediate step, one could investigate the two main corollaries of our meta theorem, Theorems 1.11 and 1.13, and in fact, the following arguments suggest that the sufficient conditions on p stated in these theorems are also necessary.

Let us start with the setting of Theorem 1.13, that of non-trivial solutions in strictly balanced homogeneous systems of linear equations. Here, as discussed before, Rué, Spiegel and Zumalacárregui in [86] already established a threshold result showing that when $np^{c(A)} = o(1)$, asymptotically almost surely $S_1(A, \mathbf{0}) \cap [n]_p^m$ is empty. Using a concentration argument similar to the one above when discussing the necessity of $n(1-p) \to \infty$, we see that in this case $|S_1(A, \mathbf{0}) \cap [n]_p^m|$ will converge in distribution to the constant 0 distribution. In addition to this threshold results, the authors also studied the distribution at the threshold, that is, the case $np^{c(A)} \to a > 0$ where a is a constant. Their results show that here, the random variable $|S_1(A, \mathbf{0}) \cap [n]_p^m|$ converges in distribution to a Poisson. Putting all of this together, we see that this establishes the only if direction for Theorem 1.13.

Let us now turn to the setting of Theorem 1.11, that of proper solutions, where there are no repeated variables by definition. The argument that $np^{c(A)} = o(1)$ implies that $\tilde{X}_n \xrightarrow{d} 0$ is the same as before, so what remains is to understand the case when $np^{c(A)}$ tends to some positive constant. When A is strictly balanced, one can use the same arguments that were used by Rué, Spiegel and Zumalacárregui in the proof of the strictly balanced homogeneous case for the distribution of non-trivial solutions to see that one will also have a Poisson distribution in the setting of Theorem 1.11. When A is not strictly balanced, an analysis as was performed by Ruciński in the subgraph setting is needed, but since we are now in the situation that all variables are pairwise distinct, the same result should follow, which would establish the necessity of $np^{c(A)} \to \infty$.

To conclude, let us mention that once one has proved a normal limiting distribution, a natural next question is the study of local limit theorems as well as anticoncentration results and tail estimates in a general context. This has been a very active trend of research in the last years, see for instance [109, 14].

Chapter 2

Independent sets in hypergraphs and the method of hypergraph containers

The main contribution of this chapter is a new multipartite generalization of the asymmetric container lemma introduced recently by Morris, Samotij and Saxton [77]. All original work presented in this chapter is based on [21] and was done jointly with Marcelo Campos, Matthew Coulson and Oriol Serra.

Broadly speaking, this chapter will concern the study of the number and structure of independent sets in k-uniform hypergraphs, where k will usually be a fixed constant. More specifically, we will present a very powerful tool, the method of hypergraph containers, developed independently by Balogh, Morris and Samotij [9] and Saxton and Thomason [92] which uses ideas that go back to Kleitman and Winston [63] to prove strong results on such topics as long as the edges of the hypergraph are sufficiently evenly distributed.

Before going into any more detail, let us see how this is connected to the topics in this thesis. The key observation is that many structures can be encoded as edges of hypergraphs, and hence independent sets represent objects avoiding this structure. For instance, for some pair of positive integers $n \geq k$, one can consider the k-uniform hypergraph $\mathcal{H} = (V, E)$ with

$$V = [n] \quad \text{and} \quad E = \{\{x, x+d, \ldots, x+(k-1)d\} : x, d \in [n], x+(k-1)d \leq n\},$$

that is, the edges are exactly the k-term arithmetic progressions (k-APs) in the first n positive integers. Hence an independent set $I \subset V$ is a subset of the first n positive integers not containing any k-APs. Let us denote by $r_k(n)$ the smallest number r such that any subset $A \subset [n]$ of size $|A| = r$ will contain at least one k-AP. As discussed in Chapter 1, by proving Theorem 1.4, Szemerédi proved in [103] that for any fixed $\delta > 0$ there exists an integer n_0 such that for every $n > n_0$, any subset A of $[n]$ of relative density δ will contain a k-AP, and hence $r_k(n) = o(n)$. On the other hand Rankin in [83] generalizing an earlier result for the $k = 3$ case by Behrend [12] gave a construction of k-AP free sets that established the lower bound

$$r_k(n) \geq n e^{-O((\log n)^{1/(k-1)})}.$$

Clearly, if A is a set free of k-APs, then the same will be true for any subset of A, and hence the number of subsets of $[n]$ not containing any k-AP can be lower bounded by $2^{r_k(n)}$. Cameron and Erdős in [18] conjectured that this should be almost the truth, asking whether the number of subsets of $[n]$ free of k-APs is $2^{(1+o(1))r_k(n)}$. Using the method of hypergraph containers, Balogh, Liu and Sharifzadeh in [8] made major progress on answering this question by showing that for infinitely many values of n, there are $2^{O(r_k(n))}$ sets in the first n integers that are free of k-term arithmetic progressions.

Note that beforehand, it is not clear at all that this more abstract view is actually helpful, since in general the structure and number of independent sets can be very unpredictable, but this is where the strength of the concept of containers enters the picture. Roughly speaking, the idea behind the method of hypergraph containers is that as long as the edges of a hypergraph $\mathcal{H}$ with vertex set V are distributed evenly enough, there exists a family $\mathcal{C} \subset 2^V$ of subsets of V called *containers* such that every independent set of $\mathcal{H}$ is contained in one of the containers. There are two crucial aspects that this family satisfies. Firstly, it is not too large, and secondly, the induced sub-hypergraph $\mathcal{H}[C]$ on any container $C \in \mathcal{C}$ will contain very few hyperedges. In order to state what "evenly enough" means, we need the following definitions. Denote by $d_{\mathcal{H}}(A) = |\{L \in E(\mathcal{H}) : A \subset L\}|$ the degree of a subset $A \subset V(\mathcal{H})$ of the vertices of $\mathcal{H}$, that is, the number of edges that contain A. Then for any integer ℓ we define $\Delta_\ell(\mathcal{H})$ as the maximum degree of a set of ℓ vertices of $\mathcal{H}$, that is,

$$\Delta_\ell(\mathcal{H}) = \max\{d_{\mathcal{H}}(A) : A \subset V(\mathcal{H}), |A| = \ell\}.$$

Finally, let

$$\mathcal{I}(\mathcal{H}) = \{I \subset V(\mathcal{H}) : 2^I \cap E(\mathcal{H}) = \emptyset\}$$

denote the family of independent sets of $\mathcal{H}$. Then the *hypergraph container lemma* states the following.

Proposition 2.1 ([77]). *Let $k \in \mathbb{N}$ and set $\delta = 2^{-k(k+1)}$. Let $\mathcal{H}$ be a k-uniform hypergraph and suppose that*

$$\Delta_\ell(\mathcal{H}) \le \left(\frac{b}{|V(\mathcal{H})|}\right)^{\ell-1} \frac{|E(\mathcal{H})|}{r} \tag{2.1}$$

for some $b, r \in \mathbb{N}$ and every $\ell \in [k]$. Then there exists a family $\mathcal{C}$ of subsets of $V(\mathcal{H})$ called containers *and a function $f : 2^{V(\mathcal{H})} \to \mathcal{C}$ such that:*

(a) for every independent set $I \in \mathcal{I}(\mathcal{H})$, there exists a fingerprint *$S \subset I$ with $|S| \le (k-1)b$ and $I \subset f(S)$.*

(b) $|C| \le |V(\mathcal{H})| - \delta r$ for every $C \in \mathcal{C}$.

As indicated, this version of the container lemma differs slightly from the original one proved in [9] and was first formulated in [77]. A first glance at Proposition 2.1 might not immediately reveal its strength, since an absolute size difference of δr between the original cardinality of the vertex set and that of the container will be rather small. The key observation is that as long as (2.1) is satisfied, the container lemma can be reapplied to the induced hypergraphs $\mathcal{H}[C]$ on *each* container $C \in \mathcal{C}$. Since every independent set I of $\mathcal{H}$ was fully contained in some set C by property (a), we have $I \in \mathcal{I}(\mathcal{H}[C])$ and hence the containment property is passed down at each step. Furthermore, note that $\mathcal{C}$ will be small, since its size is determined by the fingerprints S, each of which is small itself, specifically we have $|\mathcal{C}| \le \sum_{s \le (k-1)b} \binom{V(\mathcal{H})}{s}$. Repeated applications of the lemma will clearly increase this upper bound, but since the size of the vertex set of the hypergraph the lemma is applied to also shrinks at each step, the final family will still be sufficiently small.

Following the presentation also given in the survey [10] by Balogh, Morris and Samotij, let us now consider a specific example, the hypergraph $\mathcal{H}_n = (V, E)$ that encodes triangles. Its vertices V are the edges of the complete graph K_n and the hyperedges E are all the triples that form a triangle. Let $d = 3|E|/|V|$ denote the average degree of $\mathcal{H}_n$. Then for some fixed $\epsilon > 0$, successively applying Proposition 2.1 with parameters $b = |V(\mathcal{H}_n)|/2\sqrt{d}$ and $r = \epsilon|V|/6$ for as long as (2.1) holds and the current hypergraph contains at least ϵn^3 edges, we will eventually arrive at the following container family for triangles.

Theorem 2.2 (Theorem 2.1 in [10]). *For each $\epsilon > 0$, there exists $C > 0$ such that the following holds. For each $n \in \mathbb{N}$, there exists a collection $\mathcal{G}$ of graphs on n vertices with*

$$|\mathcal{G}| \leq n^{Cn^{3/2}}$$

such that

(a) each $G \in \mathcal{G}$ contains fewer than ϵn^3 triangles,

(b) each triangle-free graph on n vertices is contained in some $G \in \mathcal{G}$.

In order to count the number of or give structural information on graphs without any triangles, we can now apply supersaturation and stability results to the containers obtained via Theorem 2.2. The following supersaturation theorem for triangles follows from Szemerédi's regularity lemma, but can also be proved with the the classical theorem of Mantel [74] by using an averaging argument over constant sized induced subgraphs of a graph G.

Lemma 2.3. *For every $\delta > 0$, there exists $\epsilon > 0$ such that if G is a graph on n vertices and at least $(\frac{1}{4} + \delta)n^2$ edges, then G contains at least ϵn^3 triangles.*

One can now apply this to the containers from Theorem 2.2 to give a quick proof about the number of triangle-free graphs on n vertices, first proved by Erdős, Kleitman and Rothschild in [34].

Theorem 2.4 ([34]). *The number of triangle-free graphs on n vertices is $n^{n^2/4+o(n^2)}$.*

Proof. The lower bound follows from the existence of $K_{\lfloor n/2 \rfloor, \lceil n/2 \rceil}$. We will use the container method for the upper bound. Let ϵ denote the output of the supersaturation lemma when applied with some $\delta = o(1)$, then apply Theorem 2.2 with this ϵ. Since every graph $G \in \mathcal{G}$ has fewer than ϵn^3 triangles, they all have at most $n^2/4 + \delta n^2$ edges. Since every triangle-free graph is contained in some $G \in \mathcal{G}$ by the containment property, we see that the number of triangle-free graphs with n vertices is at most

$$\sum_{G \in \mathcal{G}} 2^{|E(G)|} \leq |\mathcal{G}| 2^{n^2/4+\delta n^2} \leq 2^{n^2/4+Cn^{3/2}\log(n)\log(2)^{-1}} = 2^{n^2/4+o(n^2)}.$$

■

Note that we could have also applied this technique in a relative sense to get an upper bound for the number of triangle-free graphs with n vertices and m edges of the form

$$\exp(Cn^{3/2}) \binom{(1/4+\delta)n^2}{m}.$$

We now turn to showing how one can use the container method to say something about the typical structure of triangle-free graphs in the sparse setting. Note that for the dense one, also in [34], Erdös, Kleitman and Rothschild proved that almost all triangle-free graphs are bipartite. We require the following stability result for triangles.

Lemma 2.5. *For every $\delta > 0$, there exists $\epsilon > 0$ such that if G is a graph on n vertices such that*

$$|E(G)| \geq \left(\frac{1}{2} - \epsilon\right) \binom{n}{2},$$

then either G contains at least ϵn^3 triangles or one can remove at most δn^2 edges from G to make it bipartite.

A short proof due to Füredi of Lemma 2.5 can for instance be found in [44]. We can now prove the following structural result originally proved by Łuczak in [67].

Theorem 2.6 ([67])**.** *For every $\alpha > 0$, there exists $C > 0$ such that for any $m = \omega(n^{3/2} \log n)$ it holds that almost all triangle-free graphs on n vertices and m edges can be made bipartite by removing at most αm edges.*

Proof. We begin by applying the stability result, Lemma 2.5, with $\delta = \delta(\alpha) > 0$ sufficiently small and Theorem 2.2 with the $\epsilon > 0$ then given by that lemma. Since every graph G in the container family $\mathcal{G}$ has fewer than ϵn^3 triangles, we see that because of Lemma 2.5 for every $G \in \mathcal{G}$ one of two things must be true:

(a) G can be made bipartite by removing at most δn^2 edges, or

(b) G contains less than $\left(\frac{1}{2} - \epsilon\right)\binom{n}{2}$ edges.

Note that by looking at every subgraph of $K_{\lfloor n/2 \rfloor, \lceil n/2 \rceil}$ with m edges, there clearly are at least $\binom{n^2/4}{m}$ many graphs that satisfy the conclusion of the theorem, and hence it suffices to show that the number of violations is negligible when compared to this value. We begin by considering the triangle-free graphs that are contained in a $G \in \mathcal{G}$ of the second type, that is, those that have few edges. Using a standard bound for the binomial coefficient and summing up over $\mathcal{G}$, we see that there are at most

$$\begin{aligned} n^{O(n^{3/2})} \binom{\left(\frac{1}{2} - \epsilon\right)\binom{n}{2}}{m} &\leq \exp(O(n^{3/2} \log n))(1-\epsilon)^m \binom{n^2/4}{m} \\ &\leq \exp(O(n^{3/2} \log n) - \epsilon m) \binom{n^2/4}{m} \end{aligned}$$

such triangle-free graphs, which is negligible compared to $\binom{n^2/4}{m}$ because of the lower bound on m.

We can thus turn to bad triangle-free graphs that are contained in containers of the first type, that is, those that can be made bipartite by removing at most δn^2 of their edges. Let $G \in \mathcal{G}$ be an arbitrary container of this type, and denote by G' its largest bipartite subgraph. Then any triangle-free graph H with n vertices and m edges that is contained in G and violates the structural conclusion of the theorem must contain fewer than $(1 - \alpha)m$ edges of G'. Hence the number of possible choices for H can be upper bounded by

$$\binom{|E(G)| - |E(G')|}{\alpha m} \binom{|E(G)|}{(1-\alpha)m} \leq \binom{\delta n^2}{\alpha m} \binom{\binom{n}{2}}{(1-\alpha)m} \leq 2^{-m} \binom{n^2/4}{m},$$

as long as δ is sufficiently small. Again, summing this up over all $\exp(O(n^{3/2} \log n))$ containers in $\mathcal{G}$ will result in something that is negligible when compared to $\binom{n^2/4}{m}$, and we are done. ■

Remark. Theorem 2.6 actually holds when only requiring $m \geq Cn^{3/2}$ and this can also be proved via the method of hypergraph containers, but one needs to be a bit more careful in the analysis.

Having seen this introduction on how to use the method of hypergraph containers to study objects that avoid some fixed small substructure, the next section will discuss the problem of induced (and more generally multicolored) substructures. The main result will be a multipartite version of Proposition 2.1 that can handle these problems better than the original method. Finally, in Section 2.2 we discuss some possible applications of this new method.

2.1 Forbidden multicolored structures

As we have seen in the introductory section, the method of hypergraph containers in its original form using Proposition 2.1 is very applicable when studying objects that avoid substructures, and this is true both in the graph setting that was explored in more detail, as well as in the additive one. For instance, in addition to the result on the number of sets free of k-APs by Balogh, Liu and Sharifzadeh in [8], the method can also be used to give a simple proof of the fact that for every $\delta \in (0,1)$ and every $k \in \mathbb{N}$, there exists a constant C such that if $p \geq Cn^{-1/(k-1)}$ then for n large enough, almost every subset of the binomial random set $[n]_p$ of relative density δ will contain a k-term arithmetic progression. This was first proved by Balogh, Morris and Samotij in [9] as one of the original applications of the method.

The problems that are going to be discussed in this section feel more naturally at home in the graph (and hypergraph) setting, but as we will see later, they surprisingly have also found applications in the additive one. The main question is: Can we use the method of hypergraph containers to study graphs (and more generally hypergraphs) on n vertices that avoid some small fixed graph H as an induced subgraph? This can be understood as a problem about 2-colorings (for instance with colors red and blue) of the edges of the complete graph K_n on n vertices. Define a map f_n from the set of graphs on n vertices to the set of red-blue edge-colorings of K_n in the following way. The image $f_n(G)$ is defined as

$$f_n(G)(\{v,w\}) = \begin{cases} \text{red}, & \text{if } \{v,w\} \in E(G) \\ \text{blue}, & \text{otherwise.} \end{cases}$$

Then, up to relabeling, G not containing an induced copy of H is equivalent to $f_{|V(H)|}(H)$ not being a restriction of $f_n(G)$ to a subset $S \subset [n]$ of $|V(H)|$ elements. Having reformulated the problem in this way, it is now natural to investigate the generalization to edge-colorings using r colors for some fixed $r \geq 2$. It turns out that one can indeed study this using the method of hypergraph containers, using the following encoding. For some fixed positive integer s, consider an r-edge-colored K_s, and let $c\colon E(K_s) \to [r]$ denote the function that returns the color of each edge. Then we consider the $\binom{s}{2}$-uniform hypergraph with vertex set $E(K_n) \times [r]$ and edges

$$\bigcup_{i \in [r]} \{(\phi(u)\phi(v), i) : c(uv) = i\}$$

for every injection $\phi\colon [s] \to [n]$. Note that every r-colored K_n that avoids a colored copy of K_s with the specific color pattern given by c as a subgraph will represent an independent set in the hypergraph, but the converse will not be true in general. Still, using this construction, Saxton and Thomason in [92] were able to show that the number of graphs on n vertices that avoid some fixed graph H as an induced subgraph is at most

$$2^{(1-1/c(H))\binom{n}{2}+o(n^2)},$$

where $c(H)$ denotes the largest integer c such that for some pair (c_1, c_2) satisfying $c_1 + c_2 = c$, the vertex set of H cannot be partitioned into c_1 cliques and c_2 independent sets.

One limitation of this method is that it forgets the asymmetry between the colors and instead essentially considers this r-colored K_s quantitatively the same as a 1-colored K_s. For instance, an induced H is treated the same as a clique on $|V(H)|$ vertices. This quantitative loss prevents one from proving sharp threshold behaviors for this type of problem, which is regrettable, since the method of hypergraph containers was an excellent tool for establishing such thresholds when dealing with non-induced substructures. In order to remedy this shortcoming, recently Morris, Samotij and Saxton in [77] proved an asymmetric version of Proposition 2.1 and used it to obtain the following result about the typical structure of graphs avoiding cycles of length four as induced subgraphs.

Theorem 2.7 ([77]). *For every $\epsilon > 0$ there exists $\delta > 0$ such that if n is a sufficiently large integer and m is an integer satisfying*

$$n^{4/3}(\log n)^4 \leq m \leq \delta n^2,$$

then almost all graphs G on n vertices and m edges that do not contain an induced copy of C_4 admit a partition of their vertex set $V(G) = A \cup B$ such that $|E(G[A])| \geq (1-\epsilon)\binom{|A|}{2}$ and $|E(G[B])| \leq \epsilon m$.

As mentioned before, the asymmetric version of the container method was primarily developed to study problems of this kind, and it was not immediately clear that additive applications could be possible. Nevertheless, in [19] Campos managed to modify the asymmetric version of Proposition 2.1 slightly and used it to prove results on the number and typical structure of subsets A of the first n positive integers that have a bounded sumset. This result and generalizations of it will be more closely investigated in Chapter 3.

The asymmetric version of the container lemma can be understood as dealing with bipartite hypergraphs, and so it makes sense to consider an r-partite generalization of it. The remainder of this section will be concerned with establishing exactly such a multipartite version. In order to state the result, we first need to introduce some notation. Let r be a positive integer. For an r-vector $x = (x_1, \ldots, x_r)$ we call an r-partite hypergraph $\mathcal{H}$ with vertex set $V(\mathcal{H}) = V_1 \cup \cdots \cup V_r$ *x-bounded* if $|E \cap V_i| \leq x_i$ for every hyperedge $E \in E(\mathcal{H})$ and every $1 \leq i \leq r$. Denote by $\mathcal{I}$ the family of independent sets of $\mathcal{H}$, and for any $m \in \mathbb{N}$, define

$$\mathcal{I}_m(\mathcal{H}) := \{I : I \in \mathcal{I} \text{ and } |I \cap V_r| \geq |V_r| - m\}.$$

For a subset of vertices $L \subset V(\mathcal{H})$, the codegree is defined as $d_{\mathcal{H}}(L) = |\{E \in E(\mathcal{H}) : L \subset E\}|$. Also, given a vector $v = (v_1, v_2, \ldots, v_r) \in \mathbb{Z}^r$, denote

$$\Delta_v(\mathcal{H}) := \max\{d(L) : L \subset V(\mathcal{H}),\ |L \cap V_i| = v_i,\ 1 \leq i \leq r\}.$$

Finally, for any vector y, $|y|$ will denote its 1-norm $\sum |y_i|$.

Theorem 2.8. *For all non-negative integers r, r_0 and each $R > 0$ the following holds. Suppose that $\mathcal{H}$ is a non-empty r-partite $(1, \ldots, 1, r_0)$-bounded hypergraph with $V(\mathcal{H}) = V_1 \cup V_2 \cup \ldots \cup V_r$, $m \in \mathbb{N}$, $w = (|V_1|, |V_2|, \ldots, |V_{r-1}|, m)$ and b, q are integers with $b \leq \min_i w_i$ and $q \leq m$, satisfying*

$$\Delta_v(\mathcal{H}) \leq R \left(\prod_{i=1}^{r} w_i^{v_i}\right)^{-1} b^{|v|-1} e(\mathcal{H}) \left(\frac{m}{q}\right)^{\mathbb{1}[v_r > 0]} \tag{2.2}$$

for every vector $v = (v_1, v_2, \ldots, v_r) \in \left(\prod_{i=1}^{r-1}\{0,1\}\right) \times \{0, 1, \ldots, r_0\}$. Then there exists a family $\mathcal{S} \subset \prod_{i=1}^{r}\binom{V_i}{\leq b}$ and functions $f \colon \mathcal{S} \to \prod_{i=1}^{r} 2^{V_i}$ and $g \colon \mathcal{I}_m(\mathcal{H}) \to \mathcal{S}$, such that, letting $\delta = 2^{-(r_0+r-1)(2r_0+r)}R^{-1}$, the following three things are true.

(i) If $f(g(I)) = (A_1, A_2, \ldots, A_r)$ with $A_i \subset V_i$, then $I \cap V_i \subset A_i$ for all $1 \leq i \leq r$.

(ii) For every $(A_1, A_2, \ldots, A_r) \in f(\mathcal{S})$, either $|A_i| \leq (1-\delta)|V_i|$ for some $1 \leq i \leq r-1$, or $|A_r| \leq |V_r| - \delta q$.

(iii) If $g(I) = (S_1, S_2, \ldots, S_r)$ and $f(g(I)) = (A_1, A_2, \ldots, A_r)$, then $S_i \subset I \cap V_i$ for all $1 \leq i \leq r$. Furthermore, $|S_i| > 0$ only if $|A_j| \leq |V_j| - \delta w_j$ for some $j \geq i$.

Before proceeding to the proof of Theorem 2.8, let us quickly compare it to Proposition 2.1. Let us start by recalling the nomenclature. The elements of $\mathcal{S}$ are called *fingerprints*, while the set $f(\mathcal{S})$ is the family of *containers*. Now, consider property (i). Since we are dealing with r-partite hypergraphs, it makes sense that one needs to be more restrictive in the sense that not *every* independent set of $\mathcal{H}$ can be considered, since any subset A of the vertices that only intersects at most $r-1$ components of $\mathcal{H}$ will be independent. Next, let us look at property *(ii)*. We can only guarantee that one component of a container (tuple) shrinks with each application of Theorem 2.8, but since r is usually a fixed constant, this will not result in a large quantitative difference when compared to the original method. On the other hand, note that property (ii) considers one of the components to be special in the sense that it shrinks at a different rate than the remaining $r-1$ ones. This was the essential modification made by Campos in [19] to the asymmetric version by Morris, Samotij and Saxton in order to obtain his additive results. It is not clear whether there are applications in the graph setting that make use of this feature. Finally, let us discuss the requirement of $(1,\dots,1,r_0)$-boundedness. One could think that this might be detrimental when considering r-edge-colored complete graphs that do not contain a colored copy of K_s as a subgraph when $\binom{s}{2} > r$, since then clearly some edges of this K_s will be colored the same. But this is not actually an issue, since the integer s is usually considered to be fixed, and hence one can just apply Theorem 2.8 with $r = \binom{s}{2}$. That is, if the particular r-coloring of K_s contains for instance 2 red edges, one can just define $\mathcal{H}$ to have two identical components that both represent red edges. Let us now continue to the proof.

2.1.1 The proof of Theorem 2.8

The proof of Theorem 2.8 is an adaptation of the proof of the original asymmetric container lemma due to Morris, Samotij and Saxton [77] and Campos' modified proof in [19]. Since it is rather lengthy, it will be split into several subsections.

Setup

Let $r, r_0 \in \mathbb{N}$, $m \in \mathbb{N}$ and let R be a positive real. Let b be positive integer and suppose that $\mathcal{H}$ is a $(1,\dots,1,r_0)$-bounded r-partite hypergraph* with vertex set $V = V_1 \cup \cdots \cup V_r$ satisfying (2.2) for each vector $v \in \left(\prod_{i=1}^{r-1}\{0,1\}\right) \times \{0,1,\dots,r_0\}$, $b \le \min_i |V_i|$ and $b \le m$ as in the statement of Theorem 2.8, and let $w = (|V_1|, |V_2|, \dots, |V_{r-1}|, m)$. We claim that, without loss of generality we may assume that $m \le |V_r|$. Indeed, if $m > |V_r|$, then we may replace m with $|V_r|$ as $\mathcal{I}_m(\mathcal{H}) \subseteq \mathcal{I}_{m'}(\mathcal{H})$ for any $m' \ge m$, and the right-hand side of (2.2) is a non-increasing function in m. We shall be working only with hypergraphs with edge cardinalities coming from the set

$$\mathcal{U} := \left\{ x \in \left(\prod_{i=1}^{r-1}\{0,1\}\right) \times \{1,2,\dots,r_0\} : \ x_i \le x_{i+1} \text{ and } r_0 x_{r-1} \le x_r \text{ for } 1 \le i < r \right\}.$$

The maximum codegrees we must check for each edge size $x \in \mathcal{U}$ will come from the set

$$\mathcal{V}(x) := \left(\prod_{i=1}^{r}\{0,\dots,x_i\}\right) \setminus \{(0,\dots,0)\}.$$

We now define a collection of numbers that will be upper bounds on the maximum codegrees of the hypergraphs constructed by our algorithm. To be more precise, for each $x \in \mathcal{U}$ and all $v \in \mathcal{V}(x)$, we shall force the maximum v-codegree of the x-bounded hypergraph not to exceed the quantity Δ_v^x, defined as follows.

*We remark that from now on all hypergraphs are allowed to have multi-edges, and the edges are counted with multiplicity.

Definition 2.9. For every $x \in \mathcal{U}$ and every $v \in \mathcal{V}(x)$, we define the number Δ_v^x using the following recursion:

(1) If $x = (1, \ldots, 1, r_0)$, set $\Delta_v^x := \Delta_v(\mathcal{H})$ for all $v \in \mathcal{V}(x)$.

(2) Given $x \in \mathcal{U}$, let $i' = \min\{i : x_i > 0\}$ and $x - e_{i'} = x' \in \mathcal{U}$ where $e_1, \ldots, e_r$ are the standard basis vectors of $\mathbb{R}^r$. If $v \in \mathcal{V}(x)$ satisfies $v_{i'} > 0$, denote similarly $v - e_{i'} =: v' \in \mathcal{V}(x')$.[†] Then define
$$\Delta_{v'}^{x'} := \max\left\{2\Delta_v^x, \frac{b}{w_{i'}}\Delta_{v'}^x\right\}.$$

The above recursive definition will be convenient in some parts of the analysis. In other parts, we shall require the following explicit formula for Δ_v^x, which one easily derives from Definition 2.9 using a straightforward induction on $r_0 + r - 1 - |x|$.

Observation 2.10. *For all x and v as in Definition 2.9,*
$$\Delta_v^x = \max\left\{2^{|z|}\prod_{i=1}^{r-1}\left(\frac{b}{|V_i|}\right)^{1-v_i-z_i}\left(\frac{b}{m}\right)^{r_0-v_r-z_r}\Delta_{v+z}(\mathcal{H}) : z \in \left(\prod_{i=1}^{r-1}\{0, 1-x_i\}\right) \times [0, r_0 - x_r]\right\}.$$

For future reference, we note the following two simple corollaries of Observation 2.10 and our assumptions on the maximum degrees of $\mathcal{H}$, see (2.2). Suppose that $x \in \mathcal{U}$ such that $i \in [r]$ is the least index with $e_i \in \mathcal{V}(x)$. If $i < r$, then by definition of $\mathcal{U}$ it holds that $x_j = 0$ for all $1 \le j < i$, $x_j = 1$ for all $i \le j < r$ and $x_r = r_0$, so
$$\Delta_{e_i}^x \le 2^i R \prod_{j=1}^{i-1}\left(\frac{b}{|V_j|}\right)\frac{e(\mathcal{H})}{|V_i|}. \tag{2.3}$$

If $i = r$, then $x_j = 0$ for all $1 \le j < r$ and
$$\Delta_{e_i}^x \le 2^{r+r_0} R \prod_{j=1}^{r-1}\left(\frac{b}{|V_j|}\right)\left(\frac{b}{m}\right)^{r_0-x_r}\frac{e(\mathcal{H})}{q}. \tag{2.4}$$

We will build a sequence of hypergraphs with decreasing maximum edge size, starting with $\mathcal{H}$, and making sure that for each hypergraph $\mathcal{G}$ in the sequence we have an appropriate bound on its maximum codegrees. To this end we define the following set of pairs with large codegree.

Definition 2.11. Given $x \in \mathcal{U}$, $v \in \mathcal{V}(x)$, and an x-bounded hypergraph $\mathcal{G}$, we define
$$M_v^x(\mathcal{G}) = \left\{L \in \prod_{i=1}^{r}\binom{V_i}{v_i} : d_{\mathcal{G}}(L) \ge \Delta_v^x/2\right\}.$$

The algorithm

We shall now define precisely a single round of the algorithm we use to prove the container lemma. To this end, fix some $x \in \mathcal{U}$, set $i' := \min\{i : x_i > 0\}$ and
$$x' = x - e_{i'}. \tag{2.5}$$

[†] In this case i' depends on x, so v' also depends on x not only on v, but we omit it from the notation to avoid clutter.

Suppose that $\mathcal{G}$ is an x-bounded hypergraph with $V(\mathcal{G}) = V(\mathcal{H})$. A single round of the algorithm takes as input an arbitrary $I \in \mathcal{I}(\mathcal{G})$ and outputs an x'-bounded hypergraph $\mathcal{G}_*$ satisfying $V(\mathcal{G}_*) = V(\mathcal{G})$ and $I \in \mathcal{I}(\mathcal{G}_*)$ as well as a set $S \subseteq I \cap V_{i'}$ such that $|S| \leq b$. Crucially, the number of possible outputs of the algorithm (over all possible inputs $I \in \mathcal{I}(\mathcal{G})$) is at most $\binom{|V_{i'}|}{\leq b}$.

Assume that there is an implicit linear order $\preccurlyeq$ on $V(\mathcal{G})$. The *i'-maximum vertex* of a hypergraph $\mathcal{A}$ with $V(\mathcal{A}) = V(\mathcal{G})$ is the $\preccurlyeq$-smallest vertex among all $v \in V_{i'}$ of maximal degree.

The algorithm. Set $\mathcal{A}^{(0)} = \mathcal{G}$, $S = \emptyset$ and $\mathcal{G}_*^{(0)} = (V(\mathcal{G}), \emptyset)$. Do the following for each integer $j \geq 0$ in turn:

(S1) If $|S| = b$ or $\mathcal{A}^{(j)}$ is empty, then set $L = j$ and STOP.

(S2) Let $u_j \in V_{i'}$ be the i'-maximum vertex of $\mathcal{A}^{(j)}$.

(S3) If $u_j \in I$, then add j to the set S and let

$$\mathcal{G}_*^{(j+1)} := \mathcal{G}_*^{(j)} \cup \Big\{ E \setminus \{u_j\} : E \in \mathcal{A}^{(j)} \text{ and } u_j \in E \Big\}.$$

(S4) Let $\mathcal{A}^{(j+1)}$ be the hypergraph obtained from $\mathcal{A}^{(j)}$ by removing from it all edges E such that either of the following hold:

(a) $u_j \in E$,

(b) there exists a non-empty $T \subseteq E$, such that

$$T \in M_v^{x'}\big(\mathcal{G}_*^{(j+1)}\big)$$

for some $v \in \mathcal{V}(x')$.

Finally, set $\mathcal{A} := \mathcal{A}^{(L)}$ and $\mathcal{G}_* := \mathcal{G}_*^{(L)}$. Moreover, set

$$W := \{0, \dots, L-1\} \setminus S = \Big\{ j \in \{0, \dots, L-1\} : u_j \notin I \Big\}.$$

Observe that the algorithm always stops after at most $|V(\mathcal{G})|$ iterations of the main loop. Indeed, since all hyperedges E with $u_j \in E$ are removed from $\mathcal{A}^{(j+1)}$ in part (a) of step (S4), the vertex u_j cannot be the i'-maximum vertex of any $\mathcal{A}^{(j')}$ with $j' > j$ and hence the map $\{0, \dots, L-1\} \ni j \mapsto u_j \in V(\mathcal{G})$ is injective.

The analysis

We shall now establish some basic properties of the algorithm described in the previous subsection. To this end, let us fix some $x \in \mathcal{U}$, x' and i' as defined in (2.5). Moreover, suppose that $\mathcal{G}$ is an x-bounded hypergraph and that we have run the algorithm with input $I \in \mathcal{I}(\mathcal{G})$ and obtained the x'-bounded hypergraph $\mathcal{G}_*$, the integer L, the injective map $\{0, \dots, L-1\} \ni j \mapsto u_j \in V(\mathcal{G})$, and the partition of $\{0, \dots, L-1\}$ into S and W such that $u_j \in I$ if and only if $j \in S$. We first state two straightforward, but fundamental, properties of the algorithm.

Observation 2.12. *If $I \in \mathcal{I}(\mathcal{G})$, then $I \in \mathcal{I}(\mathcal{G}_*)$.*

Proof. Observe that $\mathcal{G}_*$ contains only edges of the form $E \setminus \{v\}$ where $v \in E \cap I$ and $E \in \mathcal{G}$, see (S3). Hence, if I contained the edge $E \setminus \{v\}$ it would also contain the edge E. ∎

The next observation says that if the algorithm applied to two sets I and I' outputs the same set $\{u_j : j \in S\}$, then the rest of the output is also the same.

Observation 2.13. *Fix the hypergraph $\mathcal{G}$ we input in the algorithm, suppose that the algorithm applied to $I' \in \mathcal{F}(\mathcal{G})$ outputs a hypergraph $\mathcal{G}'_*$, an integer L', a map $j \mapsto u'_j$, and a partition of $\{0, \ldots, L'-1\}$ into S' and W'. If $\{u_j : j \in S\} = \{u'_j : j \in S'\}$, then $\mathcal{G}_* = \mathcal{G}'_*$, $L = L'$, $u_j = u'_j$ for all j, and $W = W'$.*

Proof. The only step of the algorithm that depends on the input pair I is (S3). There, an index j is added to the set S if and only if $u_j \in I$. Therefore, the execution of the algorithm depends only on the set $\{u_j : j \in S\}$ and the hypergraph $\mathcal{G}$. ■

The next two lemmas will allow us to maintain suitable upper and lower bounds on the degrees and densities of the hypergraphs obtained by applying the algorithm iteratively. The first lemma, which is the easier of the two, states that if all the maximum degrees of $\mathcal{G}$ are appropriately bounded, then all the maximum degrees of $\mathcal{G}_*$ are also appropriately bounded.

Lemma 2.14. *Given $v \in \mathcal{V}(x)$ with $v_{i'} > 0$, let $v' = v - e_{i'}$. If $\Delta_v(\mathcal{G}) \leq \Delta_v^x$, then $\Delta_{v'}(\mathcal{G}_*) \leq \Delta_{v'}^{x'}$.*

Proof. Suppose (for a contradiction) that there exists a set T, with $|T \cap V_i| = v'_i$ for all i, such that $\deg_{\mathcal{G}_*}(T) > \Delta_{v'}^{x'}$. Let j be the smallest integer satisfying

$$\deg_{\mathcal{G}_*^{(j+1)}}(T) > \Delta_{v'}^{x'}/2$$

and note that $j \geq 0$, since $\mathcal{G}_*^{(0)}$ is empty. We claim first that

$$\deg_{\mathcal{G}_*}(T) = \deg_{\mathcal{G}_*^{(j+1)}}(T). \tag{2.6}$$

Indeed, observe that $T \in M_{v'}^{x'}(\mathcal{G}_*^{(j+1)})$, and therefore the algorithm removes from $\mathcal{A}^{(j)}$ (when forming $\mathcal{A}^{(j+1)}$ in step (S4)) all edges E such that $T \subset E$. As a consequence, no further edges E with $T \subseteq E$ are added to $\mathcal{G}_*$ in step (S3).

We next claim that

$$\deg_{\mathcal{G}_*^{(j+1)}}(T) - \deg_{\mathcal{G}_*^{(j)}}(T) \leq \Delta_v^x. \tag{2.7}$$

To see this, recall that when we extend $\mathcal{G}_*^{(j)}$ to $\mathcal{G}_*^{(j+1)}$ in step (S3), we only add edges $E \setminus \{u_j\}$ such that $E \in \mathcal{A}^{(j)} \subseteq \mathcal{G}$ and $u_j \in E$. Therefore, setting $T^* = T \cup \{u_j\}$, we have

$$\deg_{\mathcal{G}_*^{(j+1)}}(T) - \deg_{\mathcal{G}_*^{(j)}}(T) \leq \deg_{\mathcal{G}}(T^*) \leq \Delta_v(\mathcal{G}) \leq \Delta_v^x,$$

where the last inequality is by our assumption, as claimed.

Combining (2.6) and (2.7), it follows immediately that

$$\deg_{\mathcal{G}_*}(T) \leq \Delta_{v'}^{x'}/2 + \Delta_v^x \leq \Delta_{v'}^{x'},$$

where the final inequality holds by Definition 2.9. This contradicts our choice of T and therefore the lemma follows. ■

We are now ready for the final lemma, which is really the heart of the matter. We will show that if $\mathcal{G}$ has sufficiently many edges and all of the maximum degrees of $\mathcal{G}$ are appropriately bounded, then either the output hypergraph $\mathcal{G}_*$ has sufficiently many edges, or the output set W must be big. We remark that here we shall use the assumption that $|I \cap V_r| \geq |V_r| - m$.

Lemma 2.15. *Suppose that $|I \cap V_r| \geq |V_r| - m$ and let $\alpha > 0$. If*

(A1) $e(\mathcal{G}) \geq \alpha \prod_{i=1}^{r-1} \left(\frac{b}{|V_i|}\right)^{1-x_i} \left(\frac{b}{m}\right)^{r_0 - x_r} e(\mathcal{H})$ *and*

(A2) $\Delta_v(\mathcal{G}) \leq \Delta_v^x$ *for every* $v \in \mathcal{V}(x)$*,*

then at least one of the following statements is true:

(P1) $e(\mathcal{G}_*) \geq 2^{-|x|-x_r-1} \alpha \prod_{i=1}^{r-1} \left(\frac{b}{|V_i|}\right)^{1-x'_i} \left(\frac{b}{m}\right)^{r_0 - x'_r} e(\mathcal{H})$.

(P2) $i' < r$ *and* $|W| \geq 2^{-i'-1} R^{-1} \alpha |V_{i'}|$.

(P3) $i' = r$ *and* $|W| \geq 2^{-r-r_0-1} R^{-1} \alpha q$.

Proof. Observe that[‡]

$$e(\mathcal{G}_*) = \sum_{j \in S} \left(e(\mathcal{G}_*^{(j+1)}) - e(\mathcal{G}_*^{(j)}) \right) = \sum_{j \in S} \Delta_{e_{i'}}(\mathcal{A}^{(j)}), \tag{2.8}$$

since $e(\mathcal{G}_*^{(j+1)}) - e(\mathcal{G}_*^{(j)}) = d_{\mathcal{A}^{(j)}}(\{u_j\})$ and u_j is the i'-maximum vertex of $\mathcal{A}^{(j)}$ for each $j \in S$, and $\mathcal{G}_*^{(j+1)} = \mathcal{G}_*^{(j)}$ for each $j \notin S$. To bound the right-hand side of (2.8), we count the edges removed from $\mathcal{A}^{(j)}$ in (a) and (b) of step (S4), which gives

$$e(\mathcal{A}^{(j)}) - e(\mathcal{A}^{(j+1)}) \leq \Delta_{e_{i'}}(\mathcal{A}^{(j)}) + \sum_v \left| M_v^{x'}(\mathcal{G}_*^{(j+1)}) \setminus M_v^{x'}(\mathcal{G}_*^{(j)}) \right| \cdot \Delta_v(\mathcal{G}).$$

Summing over $j \in \{0, \dots, L-1\}$ it follows (using (2.8)) that

$$e(\mathcal{G}) - e(\mathcal{A}) \leq e(\mathcal{G}_*) + |W| \cdot \Delta_{e_{i'}}(\mathcal{G}) + \sum_v \left| M_v^{x'}(\mathcal{G}_*) \right| \cdot \Delta_v^x,$$

since $\mathcal{A} = \mathcal{A}^{(L)} \subseteq \cdots \subseteq \mathcal{A}^{(0)} = \mathcal{G}$ and $\Delta_v(\mathcal{G}) \leq \Delta_v^x$ by (A2). Furthermore,

$$\Delta_{e_{i'}}(\mathcal{A}) \leq \Delta_{e_{i'}}(\mathcal{A}^{(j)}) \leq \Delta_{e_{i'}}(\mathcal{G}) \leq \Delta_{e_{i'}}^x, \tag{2.9}$$

since $\mathcal{A} \subseteq \mathcal{A}^{(j)} \subseteq \mathcal{G}$ and $\mathcal{G}$ satisfies (A2), which implies

$$e(\mathcal{G}) - e(\mathcal{A}) \leq e(\mathcal{G}_*) + |W| \Delta_{e_{i'}}^x + \sum_v \left| M_v^{x'}(\mathcal{G}_*) \right| \Delta_v^x. \tag{2.10}$$

Combining (2.8) and (2.9), we get

$$e(\mathcal{G}_*) = \sum_{j \in S} \Delta_{e_{i'}}(\mathcal{A}^{(j)}) \geq |S| \Delta_{e_{i'}}(\mathcal{A}) = b \Delta_{e_{i'}}(\mathcal{A}), \tag{2.11}$$

where the equality is due to the fact that $|S| \neq b$ only when $\mathcal{A}$ is empty, see step (S1).

Next, to bound the sum in (2.10), observe that, by Definition 2.11, we have

$$\left| M_v^{x'}(\mathcal{G}_*) \right| \Delta_v^{x'}/2 \leq \sum_{T:\ |T \cap V_i| = v_i} \deg_{\mathcal{G}_*}(T) \leq \binom{x_r}{v_r} e(\mathcal{G}_*) \leq 2^{x_r} e(\mathcal{G}_*)$$

for each $v \in \mathcal{V}(x')$ and therefore

$$\begin{aligned}
\sum_{v \in \mathcal{V}(x')} \left| M_v^{x'}(\mathcal{G}_*) \right| \Delta_v^x &\leq 2^{x_r+1} \sum_v e(\mathcal{G}_*) \left(\Delta_v^x / \Delta_v^{x'} \right) \\
&\leq 2^{x_r+1} (2^{|x'|} - 1) e(\mathcal{G}_*) \max_v \left\{ \Delta_v^x / \Delta_v^{x'} \right\} \\
&\leq 2^{x_r+1} (2^{|x'|} - 1) e(\mathcal{G}_*) w_{i'} / b,
\end{aligned} \tag{2.12}$$

[‡] Recall that $\mathcal{G}_*$ (and $\mathcal{G}_*^{(j)}$ etc.) are multi-hypergraphs and that edges are counted with multiplicity.

where the last inequality follows from Definition 2.9.

Suppose first that $i' < r$ and observe that substituting (2.12) into (2.10) yields

$$e(\mathcal{G}) - e(\mathcal{A}) \le e(\mathcal{G}_*) + |W|\Delta^x_{e_{i'}} + 2^{x_r+1}(2^{|x'|}-1)e(\mathcal{G}_*)|V_{i'}|/b. \tag{2.13}$$

Moreover, by (2.11) we have

$$\frac{e(\mathcal{G}_*)}{b} \ge \Delta_{e_{i'}}(\mathcal{A}) \ge \frac{e(\mathcal{A})}{|V_{i'}|} \tag{2.14}$$

since the maximum degree of a hypergraph is at least as large as its average degree. Combining (2.13) and (2.14), we obtain

$$\begin{aligned} e(\mathcal{G}) &\le e(\mathcal{G}_*)\frac{|V_{i'}|}{b}\left(\frac{b}{|V_{i'}|} + 1 + 2^{x_r+|x'|+1} - 2\right) + |W|\Delta^x_{e_{i'}} \\ &\le e(\mathcal{G}_*)\frac{|V_{i'}|}{b}2^{x_r+|x|} + |W|\Delta^x_{e_{i'}}, \end{aligned} \tag{2.15}$$

since $b \le |V_{i'}|$. Now, if the first summand on the right-hand side of (2.15) exceeds $e(\mathcal{G})/2$, then (A1) implies (P1). Otherwise, the second summand is at least $e(\mathcal{G})/2$ and by (A1) and (2.3),

$$|W| \ge \frac{e(\mathcal{G})}{2\Delta^x_{e_{i'}}} \ge \frac{\alpha}{2^{i'+1}R}|V_{i'}|,$$

which is (P2).

Finally, suppose $i' = r$. Substituting (2.12) into (2.10) yields, using the bound $\Delta^x_v/\Delta^{x'}_v \le m/b$,

$$e(\mathcal{G}) - e(\mathcal{A}) \le e(\mathcal{G}_*) + |W|\Delta^x_{e_r} + \left(2^{x_r+|x|} - 2^{x_r+1}\right)e(\mathcal{G}_*)\frac{m}{b}. \tag{2.16}$$

We claim that

$$\frac{e(\mathcal{G}_*)}{b} \ge \Delta_{e_r}(\mathcal{A}) \ge \frac{e(\mathcal{A})}{m}. \tag{2.17}$$

The first inequality follows from (2.11), so we only need to prove the second inequality. To do so, since $I \in \mathcal{F}(\mathcal{G})$ is an independent set in $\mathcal{A}$ (and all edges of $\mathcal{A}$ are contained in V_r) then every edge in $\mathcal{A}$ must be incident to $V_r \setminus I$, which has size at most m by assumption. This shows that

$$\Delta_{e_r}(\mathcal{A}) \ge \frac{e(\mathcal{A})}{|V_r \setminus I|} \ge \frac{e(\mathcal{A})}{m}.$$

Combining (2.16) and (2.17), we obtain

$$\begin{aligned} e(\mathcal{G}) &\le e(\mathcal{G}_*)\frac{m}{b}\left(\frac{b}{m} + 1 + 2^{x_r+|x|} - 2^{x_r+1}\right) + |W|\Delta^x_{e_r} \\ &\le e(\mathcal{G}_*)\frac{m}{b}2^{x_r+|x|} + |W|\Delta^x_{e_r}, \end{aligned} \tag{2.18}$$

since $b \le m$. Now, if the first summand on the right-hand side of (2.15) exceeds $e(\mathcal{G})/2$, then (A1) implies (P1). Otherwise, the second summand is at least $e(\mathcal{G})/2$ and by (A1) and (2.4),

$$|W| \ge \frac{e(\mathcal{G})}{2\Delta^x_{e_r}} \ge \frac{\alpha}{2^{r_0+r+1}R}q,$$

which is (P3). ■

Construction of the container

In this subsection, we present the construction of containers for pairs in $\mathcal{I}_m(\mathcal{H})$ and analyse their properties, thus proving Theorem 2.8. For each $s \in \{0, \ldots, r_0 + r - 1\}$, define

$$\alpha_s = 2^{-s(2r_0+r)} \qquad \text{and} \qquad \beta_s = \alpha_s \prod_{j=1}^{\min\{r-1,s\}} \left(\frac{b}{|V_j|}\right) \left(\frac{b}{m}\right)^{\max\{0,s-r+1\}}.$$

Given an $I \in \mathcal{I}_m(\mathcal{H})$, we construct the container $(A_1, \ldots, A_r)$ for I using the following procedure.

Initialize $s = 0$, $x = (1, \ldots, 1, r_0)$, $\mathcal{H}^x = \mathcal{H}$ and $S_i = \emptyset$ for all $i \in [r]$.

(C1) Let i' and x' be defined from x as before.

(C2) Run the algorithm with $\mathcal{G} \leftarrow \mathcal{H}^x$ to obtain the x'-bounded hypergraph $\mathcal{G}_*$, the sequence $u_0, \ldots, u_{L-1} \in V(\mathcal{H})$, and the partition $\{0, 1, \ldots, L-1\} = S \cup W$.

(C3) Let $S_{i'} \leftarrow S_{i'} \cup \{u_j : j \in S\}$.

(C4) If $e(\mathcal{G}_*) < \beta_{s+1} e(\mathcal{H})$, then define $(A_1, \ldots, A_r)$, the container for I, by

$$A_{i'} = V_{i'} \setminus \{u_j : j \in W\}$$

and $A_j = V_j$ for $j \neq i'$ and STOP.

(C5) Otherwise, let $\mathcal{H}^x \leftarrow \mathcal{G}_*$, $x \leftarrow x'$ and $s \leftarrow s + 1$ and CONTINUE.

We will show that the above procedure indeed constructs containers for $\mathcal{I}_m(\mathcal{H})$ that have the desired properties. To this end, we first claim that for each $x \in \mathcal{U} \cup \{0\}$, the hypergraph $\mathcal{H}^x$, if it was defined, satisfies:

(*i*) $I \in \mathcal{I}(\mathcal{H}^x)$ and

(*ii*) $\Delta_v(\mathcal{H}^x) \leq \Delta_v^x$ for every $v \in \mathcal{V}(x)$.

Indeed, one may easily prove (*i*) and (*ii*) by induction on $|x| - |v|$. The base case is true by Definition 2.9, and the inductive step follows immediately from Observation 2.12 and Lemma 2.14.

Secondly, we claim that for each input $I \in \mathcal{I}_m(\mathcal{H})$, step (C4) is called for some s and hence the container $(A_1, \ldots, A_r)$ is defined. If this were not true, the condition in step (C5) would be met $r + r_0 - 1$ times and, consequently, we would finish with a non-empty $(0, \ldots, 0)$-bounded hypergraph $\mathcal{H}^0$, i.e., we would have $\emptyset \in E(\mathcal{H}^0)$. But this contradicts (*i*), since $\emptyset \subset I$, so it would not be independent.

Suppose, therefore, that step (C4) is executed when $\mathcal{G} = \mathcal{H}^x$ for some $x \in \mathcal{U}$. We claim that $e(\mathcal{H}^x) \geq \beta_s e(\mathcal{H})$. This is trivial if $s = 0$ since here $\mathcal{H}^x = \mathcal{H}$, and for $s > 0$ it holds since otherwise step (C4) would have been executed in the previous iteration. We therefore have

$$e(\mathcal{G}) = e(\mathcal{H}^x) \geq \beta_s e(\mathcal{H}) \qquad \text{and} \qquad e(\mathcal{G}_*) < \beta_{s+1} e(\mathcal{H}),$$

which, by Lemma 2.15 and (*ii*), implies that either (P2) or (P3) of Lemma 2.15 holds. Define $\delta = 2^{-(r_0+r-1)(2r_0+r)} R^{-1}$ and note that $\delta \leq \alpha_s R^{-1}$ for all $s \in [0, r_0 + r - 1]$. If $i' < r$, we see that (P2) implies

$$|W| \geq 2^{-r-1} R^{-1} \alpha_s |V_{i'}| \geq \alpha_r R^{-1} |V_{i'}| \geq \delta |V_{i'}|,$$

Similarly, if $i' = r$, then by (P3),

$$|W| \geq 2^{-r_0-r-1}R^{-1}\alpha_s q \geq \alpha_{r_0+r-1}R^{-1}q = \delta q.$$

This verifies that $(A_1, \ldots, A_r)$ satisfies property (ii) from the statement of Theorem 2.8.

Let $\mathcal{S}$ denote the set of all tuples $(S_1, \ldots, S_r)$ that were defined in (C3) when running the procedure for all $I \in \mathcal{I}_m(\mathcal{H})$. We define $g(I) = (S_1, \ldots, S_r)$ and $f(g(I)) = (A_1, \ldots, A_r)$, where $(A_1, \ldots, A_r)$ is the container tuple that was defined in (C4). Note that f is well-defined by Observation 2.13. This follows directly when $i' < r$, since here the set $S_{i'}$ is equivalent to the set S obtained in (C2). But then, in particular, everything will be the same the first time that $i' = r$, and hence r–maximum vertices will be considered at the same time.

Finally, we see that clearly the inclusion statements of properties (i) and (iii) hold by construction, and the second one in (iii) is true since every S_i starts empty and we stop as soon as (C4) is true for the first time. ■

2.2 Possible future applications of the asymmetric container method

In this section we will investigate some possible applications of Theorem 2.8. To start, note that one such application in the additive setting will be presented in Section 3.2 of Chapter 3. Here, the multipartite container lemma will be used to investigate the number and typical structure of pairs of sets A, B in the first n positive integers such that their sumset is small. The container family theorem, Theorem 3.28 that is used to prove these results actually holds in much more generality than is needed for this particular application, but we will postpone the discussion of further uses of this specific family until the end of the following chapter.

The most obvious path for further applications of Theorem 2.8 would be to follow the original motivation of Morris, Samotij and Saxton for developing their asymmetric container lemma in [77] and study graphs and hypergraphs that avoid multicolored cliques. Most of the research for this setting is situated in what is called *anti-Ramsey* or *rainbow* theory. Here, instead of studying monochromatic structures as is done in classical Ramsey problems, one instead wants to find structures such that no pair of its components has the same color. A classical result in graph anti-Ramsey theory concerns cycles of length r in r-edge-colorings of the complete graph K_n. It was first conjectured in 1975 by Erdős, Simonovits and Sós in [35] and after several partial results for different ranges of r fully proved by Montellano-Ballesteros and Neumann-Lara in 2005 [76]. Let $\operatorname{ar}(K_n, C_r)$ denote the maximum number r satisfying the following statement. There exists an r-edge-coloring of the complete graph K_n such that every copy of the cycle C_r of size r will contain two edges of the same color. Note that here, the coloring will be required to actually use every color at least once. Then Montellano-Ballesteros and Neumann-Lara proved the following.

Theorem 2.16 ([76])**.** *For all $n \geq r \geq 3$,*

$$\operatorname{ar}(K_n, C_r) = \left(\frac{k-2}{2} + \frac{1}{k-1}\right) n + O(1).$$

To be precise, what the authors actually proved was the upper bound, the corresponding lower bound was already shown to hold by Erdős, Simonovits and Sós in [35]. This setting, while definitely interesting and widely studied (see for instance the survey [43] by Fujita, Magnant and Ozeki) is not suitable for our method in the sense that it is not clear that Proposition 2.1, the classical container lemma, would not give results of equal strength. This stems from the following observation. One of the key refinements in Theorem 2.8, the multipartite container lemma, when compared to Proposition 2.1 is that a more precise notion of codegrees is considered. Specifically, in the classical approach the maximum codegrees are only indexed by the number of fixed vertices, while in the multipartite version, one considers the refinement where the actual components the fixed vertices lie in also matters. To see that this can make a difference, consider the original motivation of Morris, Samotij and Saxton in [77] of studying induced copies of C_4. Then by fixing two vertices corresponding two *non-edges* one has already fixed the complete copy of C_4, that is, the maximum codegree of this is 1. On the other hand, fixing two edges will often result in a codegree of n, so these two cases are widely different. Now coming back to the anti-Ramsey setting as described in Theorem 2.16, we see immediately that this distinction will not exist here. Since every color in a rainbow copy of C_r appears exactly once, fixing ℓ vertices in an edge of the corresponding hypergraph always corresponds to fixing ℓ colors, and hence in some sense the 1-colored and the rainbow versions behave similarly.

Consequently, the multicolored structures that lend themselves to be studied using Theorem 2.8 should satisfy that at least one color appears at least twice in them. Problems of this type have been studied in the literature. For instance, Balogh in [7] considers the special case $r = 2$ and determines the asymptotic number of 2-edge-colorings of K_n that avoid a specific fixed 2-edge-coloring of a smaller K_k. Note that as discussed before, this case is equivalent to studying the problem of induced subgraphs. In this paper, Balogh actually proved a stability result similar to Lemma 2.5, so it would be interesting to see if one can use it in conjunction with Theorem 2.8 to prove statements similar to Theorem 2.7 for other color patterns than the one describing induced cycles of length four. Other papers exploring the topic of subgraphs with specific color patterns are [60, 13].

In addition to studying copies of small cliques with prescribed colorings, one could also investigate similar problems in additive settings. Similar to the graph setting, historically most of the focus has been placed on anti-Ramsey results, initiated with the study of rainbow arithmetic progressions in [61] by Jungić, Fox, Mahdian, Nešetřil and Radoičić. A common theme in these types of results is that one assumes the colors to be distributed nicely. On the other hand, Balandraud in [4] used a combinatorial approach to prove results on the number of colored solutions to linear equations that only depend on the sizes of the color classes and not the color distribution. His main theorem also allows the number of variables in an equation to be larger than the number of colors and hence could be applicable for the problems that can be investigated using Theorem 2.8. Another result in this direction is [75] due to Montejano and Serra. While most of the previous results where situated in cyclic or more generally finite groups, this holds in the more general combinatorial setting of *orthogonal arrays*. Most importantly, Montejano and Serra prove supersaturation results as well, and hence it might be interesting to explore whether they can be applied to study the number of arrays that avoid some specific coloring scheme.

Chapter 3

Sets with bounded sumset

The main contributions of this chapter are partial sumset versions of classical results in additive number theory, namely Kneser's addition theorem and Freĭman's $3k-4$ theorem, as well as an approximate structure theorem for almost all pairs of integer sets whose sumset is of a given size. All original work presented in this chapter is based on [21] and was done jointly with Marcelo Campos, Matthew Coulson and Oriol Serra.

The study of the sumset

$$A+B=\{a+b : a\in A,\, b\in B\}$$

and its cardinality for sets A and B in a group G is a classical topic in additive combinatorics, with the broadest categories being those of *direct* and *inverse* results.

For the former, one asks the question of how small the sumset can possibly be relative to the size of the summands. This depends heavily on what kind of group G we are situated in. Probably the earliest result, easily provable by high school students is the following statement in the integers and more generally in any group that admits a total order.

Proposition 3.1 (Folklore)**.** *Let A and B be finite sets in a linearly ordered group G. Then*

$$|A+B|\geq |A|+|B|-1.$$

Essentially, this follows immediately from the observation that if $a_1<\cdots<a_k$ are the elements of A and $b_1<\cdots<b_\ell$ those of B, then the $k+\ell-1$ elements

$$a_1+b_1, a_2+b_1, \ldots, a_k+b_1, a_k+b_2, \ldots, a_k+b_\ell$$

are all pairwise distinct.

Inverse results then look at statements of this type and ask what structure the sets A and B must have to actually achieve such lower bounds. For Proposition 3.1, the corresponding inverse result will follow from the fact that in addition to the chain of $k+\ell-1$ elements described above, one can define many more by increasing different indices of a_i and b_j one at a time. Each of these chains is $k+\ell-1$ elements long, so if the sumset contains exactly $k+\ell-1$ elements, all of them must describe the same elements, and one arrives at the following inverse result.

Proposition 3.2 (Folklore)**.** *Let A and B be finite sets in a linearly ordered group G. If*

$$|A+B| = |A|+|B|-1,$$

then A and B are arithmetic progressions with the same common difference.

The most classical nontrivial direct result is the Cauchy-Davenport theorem, which is situated in cyclic groups of prime order and was first proved by Cauchy and then independently rediscovered by Davenport.

Theorem 3.3 ([22, 27])**.** *Let p be a prime number and A and B subsets of $\mathbb{Z}/p\mathbb{Z}$. Then*

$$|A+B| \geq \min\{p, |A|+|B|-1\}.$$

The corresponding inverse statement was proved by Vosper and states that except for edge cases, we are again in the situation of arithmetic progressions.

Theorem 3.4 ([108])**.** *Let p be a prime number and A and B subsets of $\mathbb{Z}/p\mathbb{Z}$ such that $|A|, |B| \geq 2$ and*

$$|A+B| = |A|+|B|-1 \leq p-2.$$

Then A and B are arithmetic progressions with the same common difference.

All of the above results study the very extremal case of small sumsets, and in the sequel we will refer to such results as *Kneser-like*, since – at least in the case of abelian groups – the theorem of Kneser provides essentially the finishing touch in this regard. We say that a set S in an abelian group G is *periodic* if there exists a nontrivial subgroup $H < G$ such that $S + H = S$, that is, S is a union of H-cosets. Then Kneser proved the following.

Theorem 3.5 ([64])**.** *Let G be an abelian group and let $A, B \subset G$ be finite subsets. If $H = \{g \in G : g + A + B = A + B\}$ is the stabilizer of $A + B$, then*

$$|A+B| \geq |A+H| + |B+H| - |H|.$$

In particular, if $|A+B| < |A|+|B|$ then $A+B$ is periodic.

The corresponding inverse theorem was proved by Kemperman in [62], but it involves a few too many subcases to state here in full. A lot of work has been done by many authors (see for instance [42, 30, 104, 56]) to find generalizations of Kneser's and Kemperman's theorems to the non-abelian setting. Olson in [81] showed that the straightforward analogue of Theorem 3.5 in general groups does not hold. Specifically he constructed an example of finite sets A, B in a non-abelian group G with small sumset – although in non-abelian settings more often referred to as *product set* – such that HAB, AHB and ABH are all strictly larger than AB for any non-trivial subgroup $H < G$. Recently DeVos in [28] has proved a very strong analogue of Kneser's and Kemperman's theorems in this setting, implying most of the previously known results.

Kneser-like statements – more precisely their inverse results – are powerful because they give very strong guarantees on the structure of the constituent sets, but they clearly come at the expense of asking for a very small sumset. At the other end of the spectrum in terms of inverse results are what we will refer to as *Freĭman-like*. Here, one tries to obtain structural results but only asks that the set under investigation grows in a linear fashion when taking the sumset. For positive integers $d, L_1, \dots, L_d$ and elements $x_0, \dots, x_d$ in an abelian group G, a *generalized arithmetic progression* (GAP) of dimension d and size $L_1 \cdots L_d$ is a set of the form

$$\{x_0 + x_1\ell_1 + \cdots + x_d\ell_d : 0 \leq \ell_i < L_i \text{ for all } i \in [d]\}.$$

A GAP can be understood geometrically in G^d as the box defined by d arithmetic progressions going in linearly independent directions, or alternatively as a translation of a sumset of d progressions

$$x_0 + P_1 + \cdots + P_d \quad \text{where} \quad P_i = \{0, x_i, 2x_i \dots, (L_i - 1)x_i\}.$$

Freĭman proved the following structure theorem for integer sets that grow in a linear fashion.

Theorem 3.6 ([40, 41]). *Let $A \subset \mathbb{Z}$ be a finite set. If $|A+A| \leq K|A|$, then there exists a generalized arithmetic progression P of size $s(K)|A|$ and dimension $d(K)$ such that $A \subset P$.*

It is not hard to see that generalized arithmetic progressions grow only linearly under taking sumsets, and hence Theorem 3.6 completely characterizes such sets. The factor K is usually referred to as the *doubling constant* of A. Interest in this line of investigation was rekindled when Ruzsa in [89] provided a new proof of Theorem 3.6 that also gave the first effective bounds on s and d. The current best bounds are due to Sanders [90]. Moreover, Ruzsa's proof was formulated to handle the case of two distinct sets with the same cardinality as well. Theorem 3.6 was later generalized by Green and Ruzsa to the setting of general abelian groups in [52], and by Tointon to nilpotent groups in [106].

Note that when comparing Theorem 3.6 to Theorem 3.4, not only does one receive a weaker structure – that is, generalized arithmetic progressions instead of "regular" ones – it is also only a statement about being covered by this structure. This is natural, since removing a few points from a GAP will not affect the cardinality of the sumset much. A sort of middle ground between Kneser and Freĭman-like statements is represented by Freĭman's $3k-4$ theorem. The general idea is that instead of a doubling like 2 as in the former or a general K as in the latter, we are now dealing with a doubling constant somewhere between 2 and 3. In turn, one can prove a structure that still uses a covering type result as in Theorem 3.6 but with a structure similar to Theorem 3.4. Specifically, Freĭman proved the following statement.

Theorem 3.7 ([41]). *Let $A \subset \mathbb{Z}$ be a finite set of integers such that $|A+A| \leq 3|A|-4$. Then there exists an arithmetic progression P of size at most $|A+A|-|A|+1$ such that $A \subset P$.*

Note that if we want to preserve arithmetic progressions as the covering structure we cannot hope for a larger doubling constant than 3. For instance, for any pair of positive integers k and $x > 2k-2$ we see that $A = [k-1] \cup \{x\}$ satisfies $|A+A| = 3|A|-3$ but cannot be covered by any short arithmetic progression. A version of Theorem 3.7 for two distinct sets was proved by Lev and Smeliansky in [71], although some care has to be taken here since in addition to the cardinalities of the now distinct constituent sets A and B, another important aspect is how spread out they are. The case of more than two sets was also studied by Lev in [68].

A notion that is closely related to the sumset cardinality $|A+B|$ is the *representation function*

$$r_{A,B} \colon G \to \mathbb{N}, \quad x \mapsto |A \cap (x-B)|,$$

which counts the number of distinct ways to write a group element $x \in G$ as a sum of elements in A and B. Note that

$$|A||B| = \sum_{x \in G} r_{A,B}(x) \quad \text{and} \quad |A+B| = \sum_{x \in G} \min\{r_{A,B}(x), 1\}.$$

The latter identity presents another possible way to phrase both direct and indirect theorems, that is, what happens if we replace the 1 by a general positive integer t? In the case of the integers modulo a prime p, this was solved by Pollard in 1974.

Theorem 3.8 ([82]). *Let p be a prime and $A, B \subset \mathbb{Z}/p\mathbb{Z}$. For any positive integer t, it holds that*

$$\sum_{x \in \mathbb{Z}/p\mathbb{Z}} \min\{r_{A,B}(x), t\} \geq t \min\{p, |A|+|B|-t\}.$$

Note that by setting $t = 1$, we recover Theorem 3.3, the Cauchy-Davenport theorem. Theorem 3.8 has been generalized to arbitrary abelian groups in slightly different ways by both Green and Ruzsa in [51] as well as by Hamidoune and Serra in [58]. The inverse theorem that studies the structure of A and B when equality is obtained has also been proved by Nazarewicz, O'Brien, O'Neill and C. Staples in [79]. Since we will require it in a later section, let us state a version of Hamidoune and Serra's extension that was recently proved by Campos in [19]. We need the following definition. For any abelian group G and finite subsets $A, B \subset G$, define

$$\alpha(A,B) = \max\{|B'| : B' \subset G,\ |B'| \leq |B|,\ |\langle B' \rangle| \leq |A| + |B| - |B'|\}.$$

Campos' version of the Hamidoune–Serra extension of Theorem 3.8 now states the following.

Proposition 3.9 (Theorem 3.2 in [19])**.** *Let G be an abelian group, t be a positive integer and $U, V \subset G$ satisfying $t \leq |V| \leq |U| < \infty$. Then*

$$\sum_{x \in U+V} \min(r_{U,V}(x), t) \geq t(|U| + |V| - t - \alpha), \tag{3.1}$$

where $\alpha = \alpha(U,V)$.

Having presented a general overview of the topic, the remainder of this chapter is structured as follows. In Section 3.1 we will investigate so-called *partial sumsets*. The main results will be versions in this setting of Theorems 3.5 and 3.7. In Section 3.2 we will switch settings from the deterministic to the probabilistic and show how one can use Theorem 2.8 – the asymmetric version of the container lemma presented in Chapter 2 – as well as the results from Section 3.1 to give an approximate result on the typical structure of pairs of integer sets with bounded sumset. Finally, we conclude this chapter by considering some open problems and further research directions in Section 3.3.

3.1 Partial sumsets

All of the results mentioned up until now concerned the full sumset $A + B$, but a more general notion is to study what is called the partial or restricted sumset

$$A \overset{\Gamma}{+} B = \{a + b : (a,b) \in \Gamma\}$$

for a set $\Gamma \subset A \times B$, in which case the previous statements handle the case $\Gamma = A \times B$. We can now ask the same questions as before, that is, what are lower bounds on the size of partial sumsets, and what structure will sets satisfy if there exists a large Γ such that the associated partial sumset is small? Clearly, if $|A + B|$ is small, then the same will be true for $|A \overset{\Gamma}{+} B|$ for every $\Gamma \subset A \times B$. One might naively hope that the converse would also hold as long as there exists a large Γ such that $|A \overset{\Gamma}{+} B|$ is small. If that were true, one could just use the inverse results mentioned in the beginning, but sadly this is not the case. Take for instance the case when $A = B \subset \mathbb{Z}$ is a union of an arithmetic P and a geometric progression G of equal sizes $|A|/2$. Then the geometric progression will result in $|A + A|$ having size at least $|G + G| \approx |A|^2/4$, while $\Gamma = P \times P$ is large but defines a partial sumset that is very small. The classical result that links a small partial sumset to a small full sumset is the Balog–Szemerédi–Gowers theorem, first proved by Balog and Szemerédi in [6], with a new proof by Gowers in [47] which resulted in effective quantitative bounds. The version below is taken from [105].

Theorem 3.10 (Theorem 2.29 in [105])**.** *Let A and B be finite sets in an abelian group G, and let $\Gamma \subset A \times B$ satisfy*

$$|\Gamma| \geq |A||B|/K \text{ and } |A \overset{\Gamma}{+} B| \leq K'|A|^{1/2}|B|^{1/2},$$

for some $K \geq 1$ and $K' > 0$. Then there exist subsets $A' \subset A$ and $B' \subset B$ of cardinalities

$$|A'| \geq \frac{|A|}{4\sqrt{2}K} \text{ and } |B'| \geq \frac{|B|}{4K}$$

satisfying

$$|A' + B'| \leq 2^{12}K^4(K')^3|A|^{1/2}|B|^{1/2}.$$

Note that in this form the theorem is mostly useful when A and B are close in size to each other, although versions exist for the unbalanced case, see for instance Theorem 2.35 in [105]. Another interesting aspect is to consider the *almost all* case, that is, what happens if K is of the form $1/(1-\epsilon)$? The hope would then be that there exists a small $\delta > 0$ such that the sets A' and B' also contain almost everything – that is all but a δ proportion – of A and B, respectively. This was indeed proved by Shao in [96] in the case $|A| = |B|$, although his arguments also work in the slightly more general setting when the difference in cardinalities is very small compared to the size of the sets. One can then combine Shao's version of the Balog–Szemerédi–Gowers theorem with Theorem 3.7 to immediately get a partial sumset version of this for distinct sets of the same (or at least very similar) cardinalities. Specifically he proved the following.

Theorem 3.11 ([96])**.** *Let $A, B \subset \mathbb{Z}$ be two sets of cardinality $|A| = |B| = k$. Let $\epsilon > 0$ and let $\delta > 0$ be sufficiently small in terms of ϵ. Let $\Gamma \subset A \times B$ be a subset with $|\Gamma| \geq (1-\delta)k^2$. If $|A \overset{\Gamma}{+} B| \leq (3-\epsilon)k - 4$, then there exist arithmetic progressions P, Q with the same common difference and sizes at most $|A \overset{\Gamma}{+} B| - (1-\epsilon)k + 1$ such that $|A \cap P|, |B \cap Q| \geq (1-\epsilon)k$.*

It is not clear how to translate his proof to the setting of distinct sets of very different cardinalities, but obtaining something like this would be very interesting since it could lead to a tool that when combined with "full sumset" inverse theorems would give the almost all partial sumset equivalent for free. It is still possible to prove what we will in the sequel refer to as *robust* versions without this tool: Lev in [70] did this for Theorem 3.5 for a single set A in not necessarily abelian groups. To state his theorem, we require a definition that measures the regularity of $\Gamma \subset A \times B$. We may think of $\Gamma \subset A \times B$ as a subgraph of the complete bipartite graph $K_{|A|,|B|}$ where the edges (a, b) are colored by the element $c = a + b \in G$. The language of graphs will be handy, and in this way $d_\Gamma(x)$ and $N_\Gamma(x)$ will denote the degree (resp. neighborhood) of some vertex x.

Definition 3.12. Let A, B be two finite sets in an additive group and K, s non-negative integers. A subset $\Gamma \subset A \times B$ is (K, s)*-regular* if the following two things are true:

(i) $d_\Gamma(a) \geq |B| - s$ for each $a \in A$ and $d_\Gamma(b) \geq |A| - s$ for each $b \in B$.

(ii) For any $c \in A + B$ with $r_{A,B}(c) \geq K$, it holds that $c \in A \overset{\Gamma}{+} B$.

We can now state Lev's version of Theorem 3.5. We will use additive notation, but note that the theorem holds in non-abelian groups as well.

Theorem 3.13 ([70])**.** *Let A be a finite subset of a group G satisfying $|A| \geq 2$, and let K, s be non-negative integers. If $\Gamma \subset A \times A$ is (K, s)-regular and $A \overset{\Gamma}{+} A \neq A + A$, then*

$$|A \overset{\Gamma}{+} A| \geq \phi|A| - K - 2s,$$

where $\phi = (1+\sqrt{5})/2$ denotes the golden ratio.

Note that the restriction that $A \overset{\Gamma}{+} A \neq A + A$ is quite natural, since otherwise we are back in the full sumset realm, and hence we should work with Theorem 3.5 instead. Specifically, this condition of missing some sums from the full sumset in conjunction with (K, s)-regularity avoids the subgroup structures that appear in Kneser's theorem. Lev then used this robust Kneser theorem to give a lower bound on the cardinality of partial sumsets $A \overset{\Gamma}{+} A$ in the integers. Recently Shao and Xu in [97] realized that the latter result can be leveraged to prove a robust version of Theorem 3.7. They also generalized the previous two results by Lev to the case of distinct sets of the same cardinality, so their robust version of Freĭman's $3k - 4$ theorem is proved for this setting as well. One caveat is that Shao and Xu's robust Kneser theorem is only proved for abelian groups, as compared to the general group setting of Lev. Let us state this version of Theorem 3.7 in order to compare it to Theorem 3.11.

Theorem 3.14 ([97])**.** *Let $\epsilon > 0$, and let $A, B \subset \mathbb{Z}$ be subsets with $|A| = |B| = k \geq \max\{3, 2\epsilon^{-1/2}\}$, and let $\Gamma \subset A \times B$ be a subset with $|\Gamma| \geq (1 - \epsilon)k^2$. If $|A \overset{\Gamma}{+} B| < (1 + \phi - 11\epsilon^{1/2})k$, then there exist arithmetic progressions P, Q with the same common difference and sizes at most $|A \overset{\Gamma}{+} B| - (1 - 5\epsilon^{1/2})k$, such that $|A \cap P|, |B \cap Q| \geq (1 - \epsilon^{1/2})k$.*

Here as in Theorem 3.13 before, ϕ denotes the golden ratio. We can now compare Theorem 3.11 to Theorem 3.14. The former essentially works up to a doubling constant of 3 like Freĭman's original $3k - 4$ theorem, but the quantitative dependencies between ϵ and δ will be worse, since they are related to the arithmetic removal lemma by Green proved in [50]. The latter improves these dependencies and makes them explicit, but at the cost of only holding up to a doubling of essentially $1 + \phi$. The appearance of ϕ here stems directly from the use of a variant of Theorem 3.13 and consequently proving a version of this that holds up to a constant of 2 instead of ϕ would also result in an improvement of Theorem 3.14.

In any case, it is evident that the study of two distinct sets of possibly very different sizes has been largely unexplored up until now. The main results presented in what follows are versions of Theorem 3.5 and Theorem 3.7 in exactly this kind of setting. The proof of the latter largely follows the road map that was used by Shao and Xu, but some additional ideas were needed. First and foremost was the mentioned robust version of Kneser's theorem, which just like in Lev's and Shao and Xu's approaches played a central part in the final proof, so let us state it now.

Theorem 3.15. *Let U, V be finite sets in an abelian group G with $|U| \leq |V|$ and let K, s be non-negative integers. If $\Gamma \subset U \times V$ is (K, s)-regular and $U \overset{\Gamma}{+} V \neq U + V$, then*

$$|U \overset{\Gamma}{+} V| \geq |V| + \frac{|U|}{2} - K - 2s.$$

Proof. Suppose the statement is false and take a counterexample that minimizes $|U|$, the cardinality of the smaller set. Note that we can assume that the graph Γ is *saturated*, meaning that if some color $\sigma \in U + V$ is contained in $U \overset{\Gamma}{+} V$, then in fact all edges $(u, v) \in U \times V$ with $u + v = \sigma$ are contained in Γ. We start by showing that for any $u, u' \in U$, the distance $u - u'$ has many representations in $V - V$. To do this, note that since Γ is (K, s)-regular, we have $|(u + V) \setminus (U \overset{\Gamma}{+} V)| \leq s$, and similarly if we replace u by u'. So

$$|u + V \cup u' + V| \leq |U \overset{\Gamma}{+} V| + 2s < |V| + \frac{|U|}{2} - K,$$

which implies

$$r_{V,-V}(u - u') = |u + V \cap u' + V| = 2|V| - |u + V \cup u' + V| > |V| - \frac{|U|}{2} + K. \tag{3.2}$$

Next, we will show that there are many popular colors in Γ. For this, define the set P by

$$P = \left\{ \sigma \in U \overset{\Gamma}{+} V : r_{U,V}(\sigma) \geq |U|/2 \right\}.$$

Note that Γ is saturated, so the number of representations in $U \overset{\Gamma}{+} V$ and $U + V$ is identical for every color that actually appears. By (K,s)-regularity and the assumed upper bound on $|U \overset{\Gamma}{+} V|$, we have

$$\begin{aligned} |U|(|V|-s) &\leq |\Gamma| \\ &= \sum_{\sigma\in P} r_{U,V}(\sigma) + \sum_{\sigma\notin P} r_{U,V}(\sigma) \\ &< |P||U| + \left(|U \overset{\Gamma}{+} V| - |P|\right)\frac{|U|}{2} \\ &< |P|\frac{|U|}{2} + \left(|V| + \frac{|U|}{2} - K - 2s\right)\frac{|U|}{2}, \end{aligned}$$

which can be rearranged to get

$$|P| > |V| - \frac{|U|}{2} + K. \tag{3.3}$$

Next, we will show that for every $v \in V$ with $U + v \cap P \neq \emptyset$, we in fact have

$$U + v \subset U \overset{\Gamma}{+} V. \tag{3.4}$$

To see this, suppose $u_0 + v \in P$. By the definition of P, there is a set $\mathcal{P}_0 \in U \times V$ with $|\mathcal{P}_0| \geq |U|/2$ such that $u' + v' = u_0 + v$ for each $(u', v') \in \mathcal{P}_0$. Note that since $u_0 + v$ was fixed, the second components of these tuples are all pairwise distinct. Let $u \in U$ be chosen arbitrarily. It follows from (3.2) that there is a set $\mathcal{P}_1 \in V \times V$ with $|\mathcal{P}_1| \geq |V| - \frac{|U|}{2} + K$ such that $v'' - v' = u - u_0$ for each $(v', v'') \in \mathcal{P}_1$. Again, u and u_0 are fixed, and hence the first components of these tuples are pairwise distinct as well. Hence by pigeonhole, there are at least K pairs in $\mathcal{P}_0$ whose second coordinate coincides with the first coordinate of some pair in $\mathcal{P}_1$. Each two such pairs $((u', v'), (v', v'')) \in \mathcal{P}_0 \times \mathcal{P}_1$ define the relation $(u' + v') + (v'' - v') = (u_0 + v) + (u - u_0) = v + u$, implying that

$$r_{U,V}(u + v) \geq K,$$

and so $u + v \in U \overset{\Gamma}{+} V$ by (K,s)-regularity. Since the choice of u is arbitrary, this proves (3.4).

Let $V' \subset V$ be the set of elements v such that $U + v \cap P = \emptyset$. Then

$$\Gamma \cap (U \times (V \setminus V')) = U \times (V \setminus V'),$$

so since $U \overset{\Gamma}{+} V \neq U + V$, we must have $U \overset{\Gamma'}{+} V' \neq U + V'$, where $\Gamma' = \Gamma \cap (U \times V')$ is the induced subgraph of Γ on $U \times V'$. Furthermore, Γ' is (K,s)-regular: Firstly, it is clear that at most s edges are missing in the neighborhood of every vertex, since this was the case for Γ. Secondly, suppose $x \in U + V'$ is an element such that $r_{U,V'}(x) \geq K$. Then since $V' \subset V$, $r_{U,V}(x) \geq K$, and so $x \in U \overset{\Gamma}{+} V$. Since Γ was saturated, *every* edge that represented x was included, and hence $x \in U \overset{\Gamma'}{+} V'$.

Next we will show that $|U| > |V'|$. To see this, first note that we have the trivial lower bound

$$|U \overset{\Gamma'}{+} V'| \geq |V'| - s,$$

by using (K,s)-regularity and looking at the neighborhood of a single vertex of U in Γ'. On the other hand, every color in $U \overset{\Gamma'}{+} V'$ is contained in $U \overset{\Gamma}{+} V \setminus P$, and by (3.3) we thus have

$$|U \overset{\Gamma'}{+} V'| \le |U \overset{\Gamma}{+} V| - |P| < |V| + \frac{|U|}{2} - 2s - K - \left(|V| - \frac{|U|}{2} + K\right) = |U| - 2s - 2K.$$

Combining these inequalities implies

$$|U| > |V'| + s + 2K,$$

so in particular $|U| > |V'|$. Since U and V represented a counterexample that minimized the cardinality of the smaller set, we must have

$$|U \overset{\Gamma'}{+} V'| \ge |U| + \frac{|V'|}{2} - 2s - K. \tag{3.5}$$

But then by combining (3.3) and (3.5),

$$|U \overset{\Gamma}{+} V| \ge |P| + |U \overset{\Gamma'}{+} V'| > |V| + \frac{|U|}{2} - 2s + \frac{|V'|}{2} > |V| + \frac{|U|}{2} - 2s - K,$$

a contradiction. ■

When comparing Theorem 3.15 to Theorem 3.13 and the equivalent result proved by Shao and Xu in [97], we note that when the sets are of equal cardinality, our result only holds for a doubling constant of up to 3/2, so it is strictly weaker in this case. There is an additional interesting connection of Theorem 3.15 to Theorem 3.5. In [55] Hamidoune uses the so-called *isoperimetric method* to give a new proof of Theorem 3.5, in which an intermediate weaker lower bound that looks quite similar to that obtained in Theorem 3.15 is used repeatedly to eliminate certain sub-cases and arrive at the full statement in the end. It would be interesting to see whether something similar can be done here to prove the following strong statement.

Conjecture 3.16. *Let A, B be finite sets in an abelian group G and let K, s be non-negative integers. If $\Gamma \subset A \times B$ is (K,s)-regular and $A \overset{\Gamma}{+} B \neq A + B$, then*

$$|A \overset{\Gamma}{+} B| \ge |A| + |B| - K - 2s.$$

Note that it is unlikely to do better when it comes to the dependencies on K and s, as can be seen by some discussions in [70] and [69].

Next, we are going to prove a statement that can be understood as a robust version of a more fine grained equivalent to Proposition 3.1 that also takes into account the length of the sets in addition to their cardinalities. Note that some kind of more careful analysis is needed here, since even Kneser's theorem does not give any better lower bounds in the integers than the very easy to prove Proposition 3.1. For a finite set of integers A we denote its convex hull by $[A] = [\min(A), \max(A)]$.

Proposition 3.17. *Let U, V be two finite sets of integers. Assume that $\gcd(U \cup V) = 1$ and that $[U] = [0, \ell], [V] = [0, \ell']$, where $\ell' \le \ell$. Let $n = \min\{|U|, |V|\}$. Let $K \ge 2$, $s \ge 0$ and let $\Gamma \subset U \times V$ be (K,s)-regular. Then,*

$$|U \overset{\Gamma}{+} V| \ge \begin{cases} \ell + |V| - 2s, & \ell \le |U| + |V| - 2K - 2 \\ |U| + |V| + \frac{n}{2} - 4s - 2K - 2, & \ell > |U| + |V| - 2K - 2. \end{cases}$$

Before going to the proof, let us discuss some history regarding this result. The corresponding statements regarding full sumsets were proved by Freĭman [39] in the case of $A+A$ and by Lev and Smeliansky [71] and Stanchescu [101] for distinct sets. As mentioned before, other robust versions were also already established by Lev [70] in the case of $A+A$ and by Shao and Xu [97] for distinct summands of the same cardinality. The proof of the latter result was essentially identical to that of Lev, and this will also be true in the case of the proof of Proposition 3.17, with the biggest factor always being which robust version of Kneser's theorem is used.

Proof of Proposition 3.17. Let $f : \mathbb{Z} \to \mathbb{Z}/\ell\mathbb{Z}$ be the canonical projection. We write $f(x) = \tilde{x}$ and a similar notation for images of sets. From U, V, Γ we build the modular version $\tilde{U}, \tilde{V}, \tilde{\Gamma}$ in $\mathbb{Z}/\ell\mathbb{Z}$. We have $|\tilde{U}| \geq |U| - 1$ and $|\tilde{V}| \geq |V| - 1$.

Claim. $\tilde{\Gamma}$ *is* $(2K, s)$*-regular.*

Proof. The s missing edges incident to each vertex in Γ produce at most s missing edges incident to $\tilde{x}$ in $\tilde{\Gamma}$. On the other hand, since $\ell = \max(U) \geq \max(V)$, the preimage of each color different from zero in $\tilde{U} \times \tilde{V}$ produces at most two colors in $U \times V$. Hence, every nonzero color appearing at least $2K$ times in $\tilde{U} \times \tilde{V}$ must be present in Γ and therefore it must also be present in $\tilde{\Gamma}$. If $\tilde{0}$ appears more than $2K$ times in $\tilde{U} \times \tilde{V}$, since 0 and 2ℓ appear at most one time in $U \times V$ and $K \geq 2$, then ℓ must appear at least K times and the color is in Γ (and hence in $\tilde{\Gamma}$). Thus $\tilde{\Gamma}$ is $(2K, s)$–regular. ■

Now note that for every element $c \in N_\Gamma(0) \cap N_\Gamma(\ell) \subset V$, c and $c + \ell$ are distinct elements in $U \overset{\Gamma}{+} V$, but are mapped to the same element in $\mathbb{Z}/\ell\mathbb{Z}$. Since $|N_\Gamma(u)| \geq |V| - s$ for any $u \in U$ by (K, s)-regularity, using inclusion exclusion we see that

$$|U \overset{\Gamma}{+} V| \geq |\tilde{U} \overset{\tilde{\Gamma}}{+} \tilde{V}| + |N_\Gamma(0) \cap N_\Gamma(\ell)| \geq |\tilde{U} \overset{\tilde{\Gamma}}{+} \tilde{V}| + |V| - 2s. \tag{3.6}$$

Suppose first that $\ell \leq |U| + |V| - 2K - 2$. Then every $\tilde{x} \in \mathbb{Z}/\ell\mathbb{Z}$ appears in $\tilde{U} \times \tilde{V}$ at least

$$\begin{aligned} |\tilde{U} \cap (\tilde{x} - \tilde{V})| &= |\tilde{U}| + |\tilde{V}| - |\tilde{U} \cup (\tilde{x} - \tilde{V})| \\ &\geq |\tilde{U}| + |\tilde{V}| - \ell \\ &\geq |U| + |V| - 2 - \ell \\ &\geq 2K, \end{aligned}$$

times, and hence it appears in $\tilde{\Gamma}$ by $(2K, s)$-regularity. Therefore, $\tilde{U} \overset{\tilde{\Gamma}}{+} \tilde{V} = \mathbb{Z}/\ell Z$ and (3.6) gives

$$|U \overset{\Gamma}{+} V| \geq \ell + |V| - 2s,$$

as claimed.

Suppose now that $\ell > |U| + |V| - 2K - 2$. If $\tilde{U} \overset{\tilde{\Gamma}}{+} \tilde{V} \neq \tilde{U} + \tilde{V}$ then applying Theorem 3.15 gives

$$|\tilde{U} \overset{\tilde{\Gamma}}{+} \tilde{V}| \geq \max\{|\tilde{U}|, |\tilde{V}|\} + \frac{n-1}{2} - 2K - 2s \geq |U| + \frac{n}{2} - 2K - 2s - \frac{3}{2},$$

and (3.6) yields

$$|U \overset{\Gamma}{+} V| \geq |U| + |V| + \frac{n}{2} - 2K - 4s - \frac{3}{2},$$

as claimed.

Now suppose $\tilde{U} \overset{\tilde{\Gamma}}{+} \tilde{V} = \tilde{U} + \tilde{V}$. If $|\tilde{U} + \tilde{V}| \geq |\tilde{U}| + (n-1)/2$, then again (3.6) yields

$$|U \overset{\Gamma}{+} V| \geq |\tilde{U} + \tilde{V}| + |V| - 2s \geq |U| + |V| + \frac{n}{2} - 2s - \frac{3}{2}.$$

Suppose that $|\tilde{U} + \tilde{V}| < |\tilde{U}| + (n-1)/2$. Then Kneser's theorem implies that there is a nonzero subgroup $H \leq G$ such that

$$|\tilde{U} + \tilde{V}| = |\tilde{U} + \tilde{V} + H| = |\tilde{U} + H| + |\tilde{V} + H| - |H|.$$

If $H = \mathbb{Z}/\ell Z$, then (3.6) with our current hypothesis $\ell \geq |U| + |V| - 2K - 2$ gives the conclusion with room to spare. Suppose that H is a proper subgroup. We now repeat the adaptation by Shao and Xu of Lev's argument. Let

$$\begin{aligned}
C_1 &= \{c \in U \overset{\Gamma}{+} V : \tilde{c} \in \tilde{V}\}, \\
C_2 &= \{c \in U \overset{\Gamma}{+} V : \tilde{c} \in (\tilde{V} + H) \setminus \tilde{V}\}, \text{ and} \\
C_3 &= \{c \in U \overset{\Gamma}{+} V : \tilde{c} \in (\tilde{U} + \tilde{V}) \setminus (\tilde{V} + H)\},
\end{aligned}$$

which are pairwise disjoint by definition.

Using the same argument that was used to justify (3.6), we have

$$|C_1| \geq |\tilde{V}| + |V| - 2s.$$

For C_2, note that since $\tilde{U} \overset{\tilde{\Gamma}}{+} \tilde{V} = \tilde{U} + \tilde{V}$, every element in $(\tilde{V} + H) \setminus \tilde{V} \subset \tilde{U} + \tilde{V}$ has a preimage in $U \overset{\Gamma}{+} V$, so that

$$|C_2| \geq |\tilde{V} + H| - |\tilde{V}|.$$

Since $\gcd(U \cup V) = 1$ and $0 \in U \cap V$, it cannot happen that both $\tilde{U}$ and $\tilde{V}$ are contained in a single coset of H. If $\tilde{U} + H = H$, then $|\tilde{V} + H| \geq 2|H| \geq |\tilde{U}| + n - 1$, so that

$$|U \overset{\Gamma}{+} V| \geq |C_1| + |C_2| \geq |\tilde{V} + H| + |V| - 2s \geq |U| + |V| + n - 2s - 2$$

and we are done. Assume $\tilde{U} + H \neq H$, and let $N = |(\tilde{U} + \tilde{V}) \setminus (\tilde{V} + H)|/|H|$ be the number of cosets of H outside $\tilde{V} + H$. By Kneser's theorem there is at least one such coset, say $\tilde{u} + \tilde{v} + H$ with $\tilde{u} \notin H$. Let $U' = f^{-1}(\tilde{u} + H) \cap U$ and $V' = f^{-1}(\tilde{v} + H) \cap V$. Then

$$|f^{-1}(\tilde{u} + \tilde{v} + H) \cap (U \overset{\Gamma}{+} V)| \geq |U'| + |V'| - 2s - 1, \tag{3.7}$$

where the last inequality comes from considering the $|U'| + |V'| - 1$ different sums $(\min(U') + V') \cup (U' + \max(V'))$ from which at most $2s$ are missing. By inserting in (3.7) the estimates

$$|H| - |U'| \leq |(\tilde{U} + H) \setminus \tilde{U}|, \;\; |H| - |V'| \leq |(\tilde{V} + H) \setminus \tilde{V}|,$$

we obtain

$$\begin{aligned}
|C_3| &\geq N(|U'| + |V'| - 2s - 1) \\
&\geq N(2|H| - |(\tilde{U} + H) \setminus \tilde{U}| - |(\tilde{V} + H) \setminus \tilde{V}| - 2s - 1) \\
&\geq N(|H| + |\tilde{U}| + |\tilde{V}| - |\tilde{U} + \tilde{V}| - 2s - 1) \\
&\geq N|H| + (|\tilde{U}| + |\tilde{V}| - |\tilde{U} + \tilde{V}| - 2s - 1) \\
&= |(\tilde{U} + \tilde{V}) \setminus (\tilde{V} + H)| + (|\tilde{U}| + |\tilde{V}| - |\tilde{U} + \tilde{V}| - 2s - 1) \\
&= |\tilde{U}| + |\tilde{V}| - |\tilde{V} + H| - 2s - 1,
\end{aligned}$$

where in the third inequality we applied Kneser's theorem. Finally,

$$|U \overset{\Gamma}{+} V| \geq |C_1| + |C_2| + |C_3| \geq |\tilde{U}| + |\tilde{V}| + |V| - 4s - 1 \geq |U| + |V| + n - 4s - 3.$$

This completes the proof. ■

We are now ready to state and prove a robust version of Theorem 3.7 for distinct sets. As it will use Proposition 3.17, we will inherit the quantitative aspects and so when applied to sets of the same size it will give a weaker result than Theorem 3.14.

Theorem 3.18. *Let $0 < \epsilon < 1/2$ and let U, V be finite subsets of $\mathbb{Z}$ with $N = \min\{|U|, |V|\} \geq 3$ and $M = \max\{|U|, |V|\} \geq 2/\sqrt{\epsilon}$. Let $\Gamma \subset U \times V$ with $|\Gamma| \geq (1-\epsilon)|U||V|$ and*

$$|U \overset{\Gamma}{+} V| = |U| + |V| + r.$$

If

$$r < \frac{N}{2} - 13\sqrt{\epsilon}M,$$

then there are arithmetic progressions P and Q with the same common difference and lengths

$$|P| \leq |U| + r + 5\sqrt{\epsilon}M \quad \text{and} \quad |Q| \leq |V| + r + 5\sqrt{\epsilon}M$$

such that $|P \cap U| \geq (1-\sqrt{\epsilon})|U|$ and $|Q \cap V| \geq (1-\sqrt{\epsilon})|V|$.

Proof. Let

$$U' = \{u \in U : d_\Gamma(u) \geq (1-\sqrt{\epsilon})|V|\},$$

and observe that since $(1-\epsilon)|U||V| \leq |\Gamma| \leq |U'||V| + (1-\sqrt{\epsilon})(|U| - |U'|)|V|$, it holds that $|U'| \geq (1-\sqrt{\epsilon})|U|$. Similarly, if V' is the set of $v \in V$ with $d_\Gamma(v) \geq (1-\sqrt{\epsilon})|U|$, we have $|V'| \geq (1-\sqrt{\epsilon})|V|$. If $\Gamma_1 = \Gamma \cap (U' \times V')$ is the restriction of Γ, then for every $u \in U'$ we see that

$$d_{\Gamma_1}(u) = |N_\Gamma(u) \cap V'| \geq d_\Gamma(u) + |V'| - |V| \geq |V'| - \sqrt{\epsilon}|V|.$$

Similarly, $d_{\Gamma_1}(v) \geq |U'| - \sqrt{\epsilon}|U|$ for every $v \in V'$. We may assume that $[U'] = [0, \ell(U')]$, $[V'] = [0, \ell(V')]$ and $\gcd(U' \cup V') = 1$. Furthermore, without loss of generality assume that

$$\ell(U') \geq \ell(V').$$

For a set X with $[X] = [0, \ell(X)]$, denote by $h(X) = \ell(X) - |X| + 1$ the number of holes of X. We consider two cases.
Case 1. $h(U') > h(V')$. Set $K = s = \sqrt{\epsilon}M$, and define

$$\Gamma' = \Gamma_1 \cup \{(u, v) \in U' \times V' : r_{U',V'}(u+v) \geq K\}.$$

Note that by doing this we have added at most

$$\frac{(U' \times V') \setminus \Gamma_1}{K} \leq \frac{(U \times V) \setminus \Gamma}{K} \leq \sqrt{\epsilon}N$$

elements $x \in (U' \overset{\Gamma'}{+} V') \setminus (U' \overset{\Gamma_1}{+} V')$, and hence

$$|U \overset{\Gamma}{+} V| \geq |U' \overset{\Gamma'}{+} V'| - \sqrt{\epsilon}N. \tag{3.8}$$

Since Γ' is (K, s)-regular by construction, we can apply Proposition 3.17 to get

$$|U' \overset{\Gamma'}{+} V'| \ge \begin{cases} \ell(U') + |V'| - 2s, & \ell(U') \le |U'| + |V'| - 2K - 2 \\ |U'| + |V'| + \frac{\min\{|U'|,|V'|\}}{2} - 4s - 2K - 2, & \ell(U') > |U'| + |V'| - 2K - 2. \end{cases} \tag{3.9}$$

Note that by our lower bounds on $|U'|, |V'|$ and (3.8), the second line of (3.9) would imply

$$\begin{aligned} |U \overset{\Gamma}{+} V| &\ge |U' \overset{\Gamma'}{+} V'| - \sqrt{\epsilon}N \\ &\ge (1 - \sqrt{\epsilon})(|U| + |V| + N/2) - 6\sqrt{\epsilon}M - 2 - \sqrt{\epsilon}N \\ &\ge |U| + |V| + \frac{N}{2} - 11\sqrt{\epsilon}M, \end{aligned}$$

which violates our initial assumption on $|U \overset{\Gamma}{+} V|$, so the first case must hold. In particular,

$$\begin{aligned} \ell(U') &\le |U' \overset{\Gamma'}{+} V'| - |V'| + 2\sqrt{\epsilon}M \\ &\le |U \overset{\Gamma}{+} V| - (1 - \sqrt{\epsilon})|V| + 2\sqrt{\epsilon}M + \sqrt{\epsilon}N \\ &< |U| + r + 4\sqrt{\epsilon}M, \end{aligned}$$

and similarly using $h(U') > h(V')$,

$$\ell(V') < \ell(U') + |V'| - |U'| < |V| + r + 4\sqrt{\epsilon}M.$$

Case 2. $h(U') \le h(V')$. Define $U_1' = U' \cap [0, \ell(V')]$, $U_2' = U' \setminus U_1'$ and $\Gamma_1' = \Gamma \cap (U_1' \times V')$. We see that

$$|U \overset{\Gamma}{+} V| \ge |U' \overset{\Gamma_1}{+} V'| \ge |U_1' \overset{\Gamma_1'}{+} V'| + |U_2'| - \sqrt{\epsilon}|U|. \tag{3.10}$$

Since $h(U_1') \le h(U') \le h(V')$ and $\ell(U_1') = \ell(V')$ by construction, we have $|U_1'| \ge |V'|$. Furthermore, by definition of U', for every $u \in U_1'$ it holds that $d_{\Gamma_1'}(u) \ge |V'| - \sqrt{\epsilon}|V|$, and similarly, $d_{\Gamma_1'}(v) \ge |U_1'| - \sqrt{\epsilon}|U|$ for every $v \in V'$. Setting $K = s = \sqrt{\epsilon}M$ and defining

$$\Gamma'' = \Gamma_1' \cup \{(u, v) \in U_1' \times V' : r_{U_1', V'}(u + v) \ge K\},$$

we again see that

$$|U_1' \overset{\Gamma_1'}{+} V'| \ge |U_1' \overset{\Gamma''}{+} V'| - \sqrt{\epsilon}N. \tag{3.11}$$

Again, Γ'' is (K, s)-regular by construction, and so applying Proposition 3.17 we get

$$|U_1' \overset{\Gamma''}{+} V'| \ge \begin{cases} \ell(V') + |U_1'| - 2s, & \ell(V') \le |U'| + |V'| - 2K - 2 \\ \frac{3}{2}|V'| + |U_1'| - 4s - 2K - 2, & \ell(V') > |U'| + |V'| - 2K - 2. \end{cases} \tag{3.12}$$

Putting together (3.11) and (3.10), the second line of this would imply

$$\begin{aligned} |U \overset{\Gamma}{+} V| &\ge |U_1' \overset{\Gamma''}{+} V'| + |U_2'| - \sqrt{\epsilon}|U| - \sqrt{\epsilon}N \\ &\ge \frac{3}{2}|V'| + |U'| - 10\sqrt{\epsilon}M \\ &\ge |U| + |V| + \frac{N}{2} - 13\sqrt{\epsilon}M, \end{aligned}$$

a contradiction to our initial assumption. Hence the first case of (3.12) must hold, which implies

$$\begin{aligned}\ell(V') &\leq |U_1' \overset{\Gamma''}{+} V'| - |U_1'| + 2\sqrt{\epsilon}M \\ &\leq |U \overset{\Gamma}{+} V| - |U'| + 4\sqrt{\epsilon}M \\ &\leq |V| + r + 5\sqrt{\epsilon}M.\end{aligned}$$

Since $h(U') \leq h(V')$, we also have

$$\ell(U') \leq \ell(V') - |V'| + |U'| \leq |U| + r + 5\sqrt{\epsilon}M.$$

This completes the proof. ■

As discussed before, the factor of $1/2$ in front of the smaller set in Theorem 3.18 stems entirely from Theorem 3.15, hence a proof of Conjecture 3.16 would lead to an essentially best possible robust version of Freĭman's $3k-4$ theorem, since it would have the strength of Theorem 3.11 of holding for a doubling of up to 3, while also enjoying the good quantitative dependencies of Theorems 3.18 and 3.14.

3.2 The number and typical structure of sets with bounded sumset

The results in the introductory section suggest that any pair of finite sets A and B in an abelian group G that have a small sumset must be highly structured. In particular, if $F \subset G$ is a sufficiently large finite set, we therefore imagine that if we choose a pair $(A, B) \in 2^F \times 2^F$ of prescribed sizes s and t uniformly at random, the probability that their sumset is small should be close to 0. If we want to know more precisely how this probability looks like, in this finitary setting we are left with the equivalent problem of counting all such pairs. Another motivation to study this particular counting task is the *Cameron-Erdős conjecture* solved independently by Green [48] and Sapozhenko [91]. Here the problem was to count the number of sum-free subsets of the first n integers, that is, sets free of solutions to the equation $x + y = z$. In [2] Alon, Balogh, Morris and Samotij considered the sparse refinement of this problem, that is, what happens if we do not want to count the number of *all* subsets of $[n]$ that are sum-free, but only those of a specific size s? A key tool in their proof was an upper bound on the number of subsets of $[n]$ of prescribed size s and doubling K, where K was a fixed constant. They proved the following.

Theorem 3.19 ([2]). *Let $\delta > 0$, and suppose that $s \in \mathbb{N}$ is sufficiently large and that $\ell \leq s^2/\delta$. Then for each $\lambda \geq 2$, there are at most*

$$2^{\delta s}\left(\frac{(4\lambda - 3)e}{6}\right)^s$$

sets $A \subset \mathbb{N}$ with $|A| = s$, $\sum_{a \in A} a = \ell$, and $|A + A| \leq \lambda s$.

They also conjectured the following stronger result.

Conjecture 3.20 ([2]). *For every $\delta > 0$, there exists $C > 0$ such that the following holds. If $s \geq C\max(CK, \log n)$, then there are at most*

$$2^{\delta s}\binom{Ks/2}{s}$$

sets $S \subset [n]$ with $|A| = s$ and $|A + A| \leq Ks$.

Note that Conjecture 3.20 is essentially as strong as it could possibly be, since $|A+A| \leq Ks$ for every s-subset $A \subset [Ks/2]$. This conjecture was slightly too optimistic, since there is a simple counter-example when $K = \Omega(s/\log n)$ for certain values of s, pointed out by Campos in [19]. Specifically, let n and s be positive integers, and let $K, \epsilon > 0$ and $C \geq 2$ satisfy

$$\min(s, n^{1/2-\epsilon}) \geq K \geq 4\frac{\log(24C)s}{\epsilon \log n}.$$

We can now construct the counter-example as follows. Let $P \subset [n]$ be an arithmetic progression of size $Ks/8$. Furthermore, define a set $A = A_0 \cup A_1$ such that $A_0 \subset P$ is an arbitrary subset of P of size $s - K/4$, while $A_1 \subset [n] \setminus P$ is an arbitrary set outside of P of size $K/4$. Then

$$\begin{aligned}|A+A| &\leq |A_0 + A_0| + |A_0 + A_1| + |A_1 + A_1| \\ &\leq 2|P| + |A_0||A_1| + |A_1|^2 \\ &\leq \frac{Ks}{4} + \frac{Ks}{4} + \frac{K^2}{16} \leq Ks,\end{aligned}$$

the last inequality following from the bounds on K. Now using the assumed bound

$$\log(n/K^2) \geq \epsilon \log n$$

and

$$\binom{b}{d}\binom{a}{c-d} \geq \left(\frac{bc}{4ad}^{d}\right)\binom{a}{c} \quad \text{and} \quad a\binom{b}{c} \geq \binom{a^{1/c}b/e}{c}$$

valid for any positive integers a, b, c, d such that $4d \leq c$, the number of choices for A is at least

$$\binom{n/2}{K/4}\binom{Ks/8}{s-K/4} \geq \left(\frac{n}{K^2}\right)^{K/4}\binom{Ks/8}{s} \geq \binom{\exp(\epsilon K \log n/4s)Ks/8\epsilon}{s} \geq \binom{CKs}{s}.$$

On the other hand, using his variant of the method of hypergraph containers discussed in Chapter 2, Campos [19] did manage to prove a very strong version of Conjecture 3.20 almost up to the doubling of this counter-example.

Theorem 3.21 ([19])**.** *Let n be a sufficiently large integer. Then for every integer s and K satisfying $2 \leq K = o(s/(\log n)^3)$ the number of sets $A \subset [n]$ such that $|A| = s$ and $|A+A| \leq Ks$ is at most $2^{o(s)}\binom{Ks/2}{s}$.*

Note that in order for Theorem 3.21 to be non-vacuous, we require $s = \omega(\log(n)^3)$, so the case of very small s is not handled here. Campos actually proved this statement in a more general setting, that of arbitrary abelian groups. In order to state his result, we require the following definition. For an abelian group G and a real number t, define

$$\beta(t) = \max\{|H| : H \leq G,\ |H| \leq t\},$$

then Campos proved the following counting result.

Theorem 3.22 ([19])**.** *Let G be an abelian group, n a sufficiently large integer and $F \subset G$ a finite subset of size $|F| = n$. Then for any integers s and K satisfying $2 \leq K = o(s/\log(n)^3)$, the number of sets $A \subset F$ such that $|A| = s$ and $|A+A| \leq Ks$ is at most $2^{o(s)}\binom{Ks(1+\beta)}{s}$, where $\beta = \beta((1+o(1))Ks)$.*

It is not difficult to see that this statement implies Theorem 3.21 since the integers do not contain any finite non-trivial subgroups. In particular, the integer result holds in the same way for any torsion free group.

As already stated in the beginning of this section, we know that sets with small doubling will be highly structured. Specifically, in the integer setting, Theorem 3.6 gives us a precise characterization in that a set $A \subset [n]$ of size s satisfies $|A+A| \leq Ks$ if and only if it is contained in a small generalized arithmetic progression. Clearly the more complicated notion of generalized arithmetic progressions is needed to capture the structure of *any* set with a small sumset, but a natural question to ask is what happens if we only care about the *typical* case, that is, suppose we pick a set $A \subset [n]$ of prescribed size s and doubling K uniformly at random, what will it look like? One motivation for this is the following observation. Suppose for any positive integers s and $K \geq 2$ we fix a single arithmetic progression $P \subset [n]$ of size $|P| = Ks/2$. Then clearly any subset $A \subset P$ of size s will satisfy

$$|A+A| \leq |P+P| \approx Ks.$$

Since there are $\binom{Ks/2}{s}$ such subsets, Theorem 3.21 suggests that maybe almost all sets of size s and doubling K will be of this structure. In the same paper as mentioned before, Campos managed to prove a slightly weaker result.

Theorem 3.23 ([19])**.** *Let n be a sufficiently large integer and let s and K be integers satisfying*

$$2 \leq K = o(s/(\log n)^3).$$

Then for almost all sets $A \subset [n]$ of size $|A| = s$ and doubling $|A+A|/s \leq K$ there exists an arithmetic progression P of size $|P| = (1+o(1))Ks/2$ such that $A \setminus P = o(s)$.

Note that while this does not achieve a strict containment, it is quantitatively extremely strong when comparing it to the *all* case in Theorem 3.6. Currently the best bounds, due to Sanders [90], for the dimension $d(K)$ and the size $f(K)$ are of the form $d(K) = O((\log K)^{3+o(1)})$ and $f(K) = \exp(O((\log K)^{3+o(1)}))|A|$. Hence, if K is chosen to be close to the upper limit stated in Theorem 3.23, these bounds are a lot worse. This holds in particular for the dimension $d(K)$, which in the approximate structure theorem above is always 1, while even the best possible value in Theorem 3.6 would be of the order $\Omega(\log K)$. This lower bound is conjectured to represent reality, and this is usually referred to as the *polynomial Freĭman-Ruzsa conjecture*. On the other hand, for fixed K, the fact that one only achieves approximate containment in Theorem 3.23 is somewhat unsatisfying, and so it was a natural question to ask whether it is possible to do better when only sticking to such small values of K. It turns out that this is indeed the case, as was proved in follow-up work by Campos, Collares, Morris, Morrison and Souza in [20], where they proved the following structural result for sets with small doubling.

Theorem 3.24 ([20])**.** *Let n be a sufficiently large integer and let $\epsilon > 0$ and $K \geq 3$ be fixed. Then for every $s \geq (\log n)^4$ it holds that for all but an ϵ proportion of sets $A \subset [n]$ of size $|A| = s$ and doubling $|A+A|/s \leq K$ there exists an arithmetic progression P of size*

$$|P| \leq Ks/2 + c(K, \epsilon)$$

such that $A \subset P$. Furthermore, it holds that

$$c_1 K^2 \log(1/\epsilon)) \leq c(K, \epsilon) \leq c_2 K^2 \log(1/\epsilon) \log K,$$

where c_1, c_2 are absolute constants.

Note here that they only stated this theorem for $K \geq 3$ for technical reasons, specifically because as K tends to 2 the constant $c(K, \epsilon)$ will tend to infinity. Hence it should hold with the weaker assumption of $K \geq 2 + \alpha$ for some fixed $\alpha > 0$ as well, and since the case $K = 2 + o(1)$ is covered by Theorem 3.7 already, this structure should be valid for the full range of fixed K.

The main results of this chapter will be version of Theorems 3.22 and 3.23 for the case of two distinct sets A and B of sizes that are not too far apart. We start by stating the analogue of the counting result, Theorem 3.22.

Theorem 3.25. *Let G be an abelian group. Let $n \geq s_2 \geq s_1 = \Omega(s_2)$ be integers and m an integer satisfying $s_1 + s_2 \leq m = o(s_2^2(\log s_2)^{-4}(\log n)^{-3})$. Then for any $F_1, F_2 \subset G$ with $|F_i| = n$, the number of pairs of sets $(X_1, X_2) \in 2^{F_1} \times 2^{F_2}$ such that $|X_i| = s_i$ and $|X_1 + X_2| \leq m$ is at most*

$$2^{o(s_2)} \binom{\frac{s_1}{s_1+s_2}(m+\beta)}{s_1} \binom{\frac{s_2}{s_1+s_2}(m+\beta)}{s_2},$$

where $\beta = \beta((1+o(1))m)$.

In groups that allow a result similar to Theorem 3.7 one can get rid of the $\log s_2$ term in the upper bound of m in Theorem 3.25. Specific examples would be the integers $G = \mathbb{Z}$ or the integers modulo some prime p, that is, $G = \mathbb{Z}/p\mathbb{Z}$. Let us also remark that the previous example for the case $A = B$ that showed that (in the integers) this counting result does not hold in general for $K = \Omega(s/\log n)$ can easily be adapted to the case of two distinct sets, so we are in the same situation as before in that some power of $\log n$ in the denominator will be necessary. We now state our analogue of Theorem 3.23 about the structure of integer subsets with bounded sumset.

Theorem 3.26. *Let $n \geq s_2 \geq s_1 = \Omega(s_2)$ be integers and m an integer satisfying*

$$s_1 + s_2 \leq m = o(s_2^2/(\log n)^3).$$

Then for almost all sets $X_1, X_2 \subset [n]$ such that $|X_i| = s_i$ and $|X_1 + X_2| \leq m$, there exist arithmetic progressions P_1 and P_2 with the same common difference of size $|P_i| = (1+o(1))s_i m/(s_1+s_2)$ and $|X_i \setminus P_i| = o(s_i)$.

The strength of the statement can be illustrated by considering the case where $s_2 = n^\alpha$ for some $0 < \alpha \leq 1/2$. Then $m = n^{\alpha(2-\epsilon)} = s_2^{2-\epsilon}$ is a valid choice for any fixed $\epsilon > 0$, and hence Theorem 3.26 states that for almost every pair of sets A, B of size $\Theta(n^\alpha)$ such that $|A+B| \leq s_2^{2-\epsilon}$, both A and B are (up to scaling and translating) almost contained in an interval of size $O(n^{\alpha(2-\epsilon)}) = o(n)$. The proof requires that the cardinalities of the two sets X_1, X_2 are not arbitrarily far apart, but this seems natural: Even if the condition $s_1 = \Omega(s_2)$ can be weakened, it is not clear that a nontrivial structural result should hold when s_1 is much smaller than s_2. Note that in contrast to the counting result, it is not clear that any power of $\log n$ should be necessary in the upper bound stated in Theorem 3.26 and an interesting question would be to investigate whether it might be true for any $m = o(s_2^2)$.

Before proceeding to the proofs of Theorems 3.25 and 3.26, let us quickly discuss how they relate to Theorems 3.22 and 3.23 proved by Campos in [19]. Firstly, note that our results as stated do not imply Campos'. This is rather obvious for the counting result, but is also the case when considering the structural one, since Theorem 3.26 is an *almost all* result concerning tuples of possibly distinct sets (A, B). Since in comparison the number of tuples (A, A) is negligible, our theorem is not able to make any statement on the structure of them. On the other hand, this is somewhat of a technical consequence of how the theorem was stated, and it is possible to formulate versions of both the counting and the structural result that supersede Campos' at least in a qualitative sense. For instance, such a version of Theorem 3.26 is the following.

Corollary 3.27. *Let $0 \leq \lambda_1, \lambda_2 \leq 2$ be integers such that $\lambda_1 + \lambda_2 = 2$. Furthermore, let $n \geq s_2 \geq s_1 = \Omega(s_2)$ and m be positive integers satisfying*

$$\lambda_1 s_1 + \lambda_2 s_2 \leq m = o(s_1^{\lambda_1} s_2^{\lambda_2} / (\log n)^3).$$

Then for almost all sets $X_1, X_2 \subset [n]$ such that $|X_i| = s_i$ if $\lambda_i > 0$ and $|X_i| = 0$ otherwise, and satisfying $|\lambda_1 X_1 + \lambda_2 X_2| \leq m$, there exist arithmetic progressions P_1 and P_2 with the same common difference of size

$$|P_i| = (1 + o(1)) s_i m / (\lambda_1 s_1 + \lambda_2 s_2),$$

such that $|X_i \setminus P_i| = o(s_i)$.

Proof. The case $\lambda_1 = \lambda_2 = 1$ follows from Theorem 3.26, while the cases $\lambda_i = 2, \lambda_{3-i} = 0$ with $i \in [2]$ follow from Theorem 3.23 applied with $K = m/s_i$. ■

While we refrained from stating our results in this manner, further investigations should probably be conducted in this way, since as mentioned before, the proofs of any combination of for instance three summand sets would likely be very similar. This leads to the second point of comparison, namely the quantitative aspects that are hidden in Theorems 3.25 and 3.26 but can be seen when looking at their technical versions, Theorems 3.31 and 3.33. Here, our counting result is essentially quantitatively identical to that of Campos. On the other hand, the structural result has slightly worse bounds, which is due to the fact that an additional case has to be considered when dealing with a pair of potentially distinct sets. Having mentioned this, we can proceed to the proofs.

3.2.1 The proofs of Theorems 3.25 and 3.26

Both Theorem 3.25 and 3.26 will follow from more technical versions that will be stated later. The main tool used in proving these statements is the method of hypergraph containers, specifically the r-partite version presented in Chapter 2. As already discussed there, the usual structure of this method is to essentially have two separate ingredients, the first being a result about independent sets in hypergraphs following certain degree conditions, and the second being supersaturation and stability results. Our proof will also follow along these lines. We start by showing how Theorem 2.8 can be used to construct a small family of containers that is suitable for our needs.

The container family

We will be making use of the following hypergraph construction. For a group G and finite subsets $F_1, \ldots, F_r \subset G$, define the r-partite and $(1, \ldots, 1)$-bounded hypergraph $\mathcal{H}(F_1, \ldots, F_r)$ in the following way. The vertex set is $\bigsqcup_{i \in r} F_i$ and $\{f_1, \ldots, f_r\}$ is a hyperedge if $f_i \in F_i$ for all $i \in [r]$ and $f_r = f_1 \cdots f_{r-1}$. Note that the sets F_i need not actually be disjoint. We can now state what is usually called the *container theorem*.

Theorem 3.28. *Let G be a group, $h \geq 2$ an integer and $\epsilon > 0$. Suppose $n, m, s_1, \ldots, s_h$ are integers such that $\log n \leq \max s_i \leq m \leq \log n (\min s_i)^h$, and let $F_1, \ldots, F_h$ be subsets of G of cardinality $|F_i| = n$ with product set $F = F_1 F_2 \cdots F_h$. Then there exists a family $\mathcal{A} \subset \prod_{i \in [h]} 2^{F_i} \times 2^F$ of $(h+1)$–tuples $(A_1, \ldots, A_h, B)$ of size*

$$|\mathcal{A}| \leq \exp\left(2^{(h+1)(h+5)} \epsilon^{-h} m^{1/h} (\log n)^{(2h-1)/h}\right) \tag{3.13}$$

such that the following two things are true:

(a) *For all $X_i \subset F_i$, $Y \subset F$ with $|X_i| = s_i$, $X_1 X_2 \cdots X_h \subseteq Y$ and $|Y| \leq m$, there exists a tuple $(A_1, \dots, A_h, B) \in \mathcal{A}$ such that $B \subset Y$ and $X_i \subset A_i$ for all $i \in [h]$.*

(b) *For every $(A_1, \dots, A_h, B) \in \mathcal{A}$ it holds that $|B| \leq m$ and either $\max_i |A_i| < m/\log n$ or there are at most $\epsilon^h \prod |A_i|$ tuples $(a_1, \dots, a_h) \in \prod A_i$ such that $a_1 a_2 \cdots a_h \notin B$.*

Proof. We will construct a rooted tree $\mathcal{T}$ with root $\mathcal{H}(F_1, \dots, F_h, F)$ and leaves $\mathcal{H}(A_1, \dots, A_{h+1})$ such that one of the following properties holds:

(i) $|A_i| < s_i$ for some $i \leq h$,

(ii) $\max\{|A_1|, \dots, |A_h|\} < m/\log n$,

(iii) $|A_{h+1}| < |F| - m$, or

(iv) $\mathcal{H}(A_1, \dots, A_{h+1})$ has less than $\epsilon^h \prod_{i \in [h]} |A_i|$ hyperedges.

The end-goal is to essentially have $\mathcal{A}$ be the subset of the leaves of $\mathcal{T}$ that correspond to properties (ii) and (iv).

We construct $\mathcal{T}$ in the following way. Given a vertex $\mathcal{H} = \mathcal{H}(V_1, \dots, V_{h+1})$ of $\mathcal{T}$ with $\max_{i \in [h]} |V_i| \geq m/\log n$, $|V_i| \geq s_i$ for all $i \in [h]$, $|V_{h+1}| \geq |F| - m$ and $e(\mathcal{H}) \geq \epsilon^h \prod_{i \in [h]} |V_i|$, we apply Theorem 2.8 with parameters $R = \epsilon^{-h}$, $q = m/\log n$ and $b = q^{1/h}$. Note that $b \leq \min s_i$ because of our upper bound on m. Let us show that these choices indeed satisfy the codegree conditions of the container lemma. Let $v = (v_1, \dots, v_{h+1}) \in \{0,1\}^{h+1}$. The edges of the hypergraph are defined by a linear relation, hence if we fix $|v|$ entries of an edge according to the vector v, at most $h - |v|$ of the $h + 1 - |v|$ remaining components can be chosen freely, and hence the maximum v–degree of $\mathcal{H}$ can be upper bounded by the product of cardinalities of the $h - |v|$ smallest open sets. Expressing this as formula, we see

$$\Delta_v(\mathcal{H}) \leq \min_{\substack{w=(w_1,\dots,w_{h+1}) \\ w_i \in \{0,1-v_i\} \\ |w| = h-|v|}} \prod_{i=1}^{h} |V_i|^{w_i} |(V_1 V_2 \cdots V_h) \cap V_{h+1}|^{w_{h+1}}. \tag{3.14}$$

Because of our lower bound on $e(\mathcal{H})$, in order to prove that (2.2) holds, it suffices that the parameters are chosen such that the right-hand side of (3.14) can be upper bounded by

$$\epsilon^h R q^{-v_{h+1}} \prod_{i=1}^{h} |V_i|^{1-v_i} b^{|v|-1}.$$

We begin by looking at a special case, $v = (1, \dots, 1)$. It is easy to see that the codegree $\Delta_v(\mathcal{H})$ is 1, and hence we get the condition

$$\epsilon^h R q^{-1} b^h \geq \Delta_v(\mathcal{H}) = 1. \tag{3.15}$$

We now consider v such that $v_{h+1} = 1$ and $|v| \leq h$. Let $j \in [h]$ be the index that corresponds to the largest V_i. Since there clearly is a v such that $v_j = 0$, equation (3.14) implies that we need

$$\epsilon^h R b^{|v|-1} \max_{i \in [h]} |V_i| \geq q,$$

so in particular if $v = (0, \dots, 0, 1)$ we get the most restrictive one (any sensible b will be positive), namely

$$\epsilon^h R \max_{i \in [h]} |V_i| \geq q. \tag{3.16}$$

Finally we consider v such that $v_{h+1} = 0$ and note $|v| \leq h$. It suffices to be larger than any one of the expressions on the right-hand side of equation (3.14), so we can ignore the product set expression and have $w_i = 1$ for the $h - v$ smallest available V_i with $i \in [h]$, and hence

$$\Delta_v(\mathcal{H}) \leq \prod_{j=1}^{h} |V_i|^{1-v_i}.$$

We can ignore the case $|v| = h$, since this will lead to a weaker restriction than (3.15). For v with $|v| < h$, the above implies $\epsilon^h R b^{|v|-1} \geq 1$, which in its most restrictive form with $|v| = 1$ means

$$\epsilon^h \geq R. \tag{3.17}$$

So we see that $R = \epsilon^{-h}$, $q = m/\log n$ and $b = q^{1/h}$ are indeed valid choices. Hence, by Theorem 2.8, there exists a family $\mathcal{C} \subset \prod 2^{V_i}$ of size at most

$$|\mathcal{C}| \leq \prod_{i=1}^{h+1} \binom{|V_i|}{\leq b} \leq b^{h+1} \binom{n^h}{b} \binom{n}{b}^h \leq n^{2hb} = \exp\left(2hm^{1/h}(\log n)^{(h-1)/h}\right), \tag{3.18}$$

such that for each $I \in \mathcal{I}_m(\mathcal{H})$ there exist $(A_1, \ldots, A_{h+1}) \in \mathcal{C}$ with $I \cap V_i \subset A_i$ for all $i \in [h+1]$, and either $|A_i| \leq (1-\delta)|V_i|$ for some $i \in [h]$, or $|A_{h+1}| \leq |V_{h+1}| - \delta q$, with $\delta = \epsilon^h 2^{-(h+1)(h+3)}$. For each $(A_1, \ldots, A_{h+1}) \in \mathcal{C}$, add $\mathcal{H}(A_1, \ldots, A_{h+1})$ as a child of $\mathcal{H}$ in $\mathcal{T}$. In order to bound the number of leaves of $\mathcal{T}$, we will first bound its height.

Claim. *The tree $\mathcal{T}$ has height at most $d = 2^{(h+1)(h+4)}\epsilon^{-h} \log n$.*

Proof. Suppose $\mathcal{H}(A_1, \ldots, A_{h+1})$ is a vertex of $\mathcal{T}$ of depth d. Recall that after each application of Theorem 2.8, one component shrunk, hence after d applications one of them shrunk at least $d/(h+1)$ times. Since we started at $\mathcal{H}(F_1, \ldots, F_h, F)$ and $\delta = 2^{-(h+1)(h+3)}\epsilon^h$, either

$$|A_{h+1}| \leq |F| - \frac{d\delta q}{h+1} = |F| - \frac{d\epsilon^h m}{(h+1)2^{(h+1)(h+3)} \log n} < |F| - m,$$

or for one $i \in [h]$,

$$|A_i| \leq (1-\delta)^{d/(h+1)} n \leq \exp(-\delta d/(h+1)) n \leq 1,$$

and so this vertex has no children. ∎

We will now define the family $\mathcal{A}$ formally. If $\mathcal{L}$ is the set of leaves of $\mathcal{T}$, let

$$\mathcal{A} = \left\{ (A_1, \ldots, A_h, B) : \begin{array}{c} \mathcal{H}(A_1, \ldots, A_h, F \setminus B) \in \mathcal{L}, \\ |B| \leq m \text{ and } |A_i| \geq s_i \text{ for all } i \in [h] \end{array} \right\}.$$

Since every tuple $(A_1, \ldots, A_h, B)$ in this family corresponds to a leaf with $|A_i| \geq s_i$ for every $i \in [h]$ and $|F \setminus B| \geq |F| - m$, we must have either $\max_{i \in [h]} |A_i| < m/\log n$ or the corresponding hypergraph must have less than $\epsilon^h \prod_{i \in [h]} |A_i|$ edges, that is, there are less than $\epsilon^h \prod_{i \in [h]} |A_i|$ h-tuples $(a_1, \ldots, a_h) \in \prod_{i \in [h]} A_i$ such that $a_1 a_2 \cdots a_h \notin B$. In any case, (b) holds. The size of $\mathcal{A}$ is at most the dth power of the maximal number of children of a vertex in $\mathcal{T}$, and so by (3.18) we see that

$$|\mathcal{A}| \leq \exp\left(d2hm^{1/h}(\log n)^{(h-1)/h}\right) \leq \exp\left(2^{(h+1)(h+5)}\epsilon^{-h} m^{1/h} (\log n)^{(2h-1)/h}\right),$$

and so (3.13) holds. Finally, property (a) holds since for all $X_1, \ldots, X_h$ with $X_i \subset F_i$, $Y \subset F$ satisfying $|X_i| = s_i$ for $i \in [h]$, $|Y| \leq m$ and $X_1 X_2 \cdots X_h \subset Y$, we see that $\bigcup_{i \in [h]} X_i \cup (F \setminus Y)$ is contained in $\mathcal{I}_m(\mathcal{H}(F_1, \ldots, F_h, F))$, and so by the properties of the containers there is a path in $\mathcal{T}$ from the root to a leaf $\mathcal{H}(A_1, \ldots, A_h, F \setminus B)$ such that $X_i \subset A_i$ and $B \subset Y$. By the size bounds for X_i and Y it is clear that this leaf must correspond to an $(h+1)$-tuple in $\mathcal{A}$. ∎

Supersaturation and stability results

The next step will be to prove supersaturation and stability results that we can then apply to the containers obtained from Theorem 3.28. These will make use of the results proved in Section 3.1. We begin by stating the supersaturation result, which is a corollary of Proposition 3.9, Campos' version of the Hamidoune-Serra generalization of Pollard's theorem.

Proposition 3.29. *Let G be an abelian group, $A_1, A_2, B \subset G$ be finite and non-empty subsets of G and $0 < \epsilon < 1/2$, and denote $\beta = \beta((1+4\epsilon)|B|)$. If*

$$|A_1| + |A_2| \geq (1+2\epsilon)(|B| + \beta),$$

then there are at least $\epsilon^2|A_1||A_2|$ pairs $(a_1, a_2) \in A_1 \times A_2$ such that $a_1 + a_2 \notin B$.

Proof. Without loss of generality we can assume $|A_2| \geq |A_1|$. If $|B| \leq (1-\epsilon^2)|A_2|$, then since $r_{A_1,A_2}(b) \leq |A_1|$ for every fixed $b \in B$, we have at least

$$|A_1||A_2| - |B||A_1| \geq \epsilon^2|A_1||A_2|$$

pairs $(a_1, a_2) \in A_1 \times A_2$ such that $a_1 + a_2 \notin B$. So we can assume $|B| > (1-\epsilon^2)|A_2|$, which also implies $|A_1| \geq \epsilon|A_2|$. Now applying Proposition 3.9 with $t = \epsilon|A_2|$, $U = A_2$ and $V = A_1$ gives us

$$\sum_{x \in A_1 + A_2} \min(r_{A_1,A_2}(x), \epsilon|A_2|) \geq \epsilon|A_2|(|A_1| + (1-\epsilon)|A_2| - \alpha)$$

and hence

$$\sum_{x \in (A_1+A_2)\setminus B} \min(r_{A_1,A_2}(x), \epsilon|A_2|) \geq \epsilon|A_2|(|A_1| + (1-\epsilon)|A_2| - |B| - \alpha). \tag{3.19}$$

We will show that

$$\alpha \leq \max(\beta, |A_1| + |A_2| - (1+4\epsilon)|B|).$$

Indeed, suppose $A' \subset G$ satisfies $|A'| \leq |A_1|$ and $|\langle A' \rangle| \leq |A_1| + |A_2| - |A'|$. If $|A'| \leq \beta$ we are done, so suppose $|A'| > \beta$, and hence $|\langle A' \rangle| \geq (1+4\epsilon)|B|$ by definition of β. So A' satisfies

$$|A'| \leq |A_1| + |A_2| - |\langle A' \rangle| \leq |A_1| + |A_2| - (1+4\epsilon)|B|,$$

which is what we wanted to show. Now note that since $\epsilon < 1/2$ and $|B| \geq (1-\epsilon^2)|A_2|$ we have $|B| \geq 2|A_2|$ and hence

$$|A_1| + (1-\epsilon)|A_2| - |B| - (|A_1| + |A_2| - (1-4\epsilon)|B|) = 4\epsilon|B| - \epsilon|A_2| \geq \epsilon|A_2| \geq \epsilon|A_1|.$$

Similarly, since $|A_1| + |A_2| \geq (1+2\epsilon)(|B| + \beta)$ and $\epsilon < 1/2$ we have

$$\begin{aligned} |A_1| + (1-\epsilon)|A_2| - |B| - \beta &\geq |A_1| + (1-\epsilon)|A_2| - \frac{|A_1| + |A_2|}{1+2\epsilon} \\ &> |A_1| + (1-\epsilon)|A_2| - (1-\epsilon)(|A_1| + |A_2|) \\ &= \epsilon|A_1|. \end{aligned}$$

Hence (3.19) implies

$$\sum_{x \in (A_1+A_2)\setminus B} r_{A_1,A_2}(x) \geq \epsilon^2|A_1||A_2|.$$

■

The stability result that we will use follows from Theorem 3.18, the robust version of Freĭman's $3k-4$ theorem proved in the preceding section.

Corollary 3.30. *Let $s_1 \le s_2$ be positive integers, and $0 < \epsilon \le 2^{-8}\left(\frac{s_1}{s_1+s_2}\right)^2$. If $A_1, A_2, B \subset \mathbb{Z}$, such that $(1-\epsilon)|B| \le |A_1| + |A_2|$ and $|A_i| \le \left(\frac{s_i}{s_1+s_2} + 2\sqrt{\epsilon}\right)|B|$ for $i = 1,2$, then one of the following holds:*

(a) There are at least $\epsilon^2|A_1||A_2|$ pairs $(a_1, a_2) \in A_1 \times A_2$ such that $a_1 + a_2 \notin B$.

(b) There are arithmetic progressions P_1, P_2 of length $|P_i| \le \frac{s_i}{s_1+s_2}|B| + 4\sqrt{\epsilon}|B|$ with the same common difference such that P_i contains all but at most $\epsilon|A_i|$ points of A_i.

Proof. Let $\Gamma = \{(a_1,a_2) \in A_1 \times A_2 : a_1 + a_2 \in B\}$. If $|\Gamma| < (1-\epsilon^2)|A_1||A_2|$ case (a) holds, so assume the converse. It is a straightforward computation that for $\epsilon \le 2^{-8}\left(\frac{s_i}{s_1+s_2}\right)^2$,

$$\frac{3}{2}\left(\frac{s_i}{s_1+s_2} - 2\sqrt{\epsilon} - \epsilon\right) + (1-13\epsilon)\left(1 - \frac{s_i}{s_1+s_2} - 2\sqrt{\epsilon} - \epsilon\right) > 1,$$

and so since

$$|A_i| \ge (1-\epsilon)|B| - \left(\frac{s_{3-i}}{s_1+s_2} + 2\sqrt{\epsilon}\right)|B| = \left(\frac{s_i}{s_1+s_2} - \epsilon - 2\sqrt{\epsilon}\right)|B|,$$

it holds that

$$\begin{aligned}|A_1 \overset{\Gamma}{+} A_2| &\le |B| \\ &\le \frac{3}{2}\left(\frac{s_i}{s_1+s_2} - 2\sqrt{\epsilon} - \epsilon\right)|B| + (1-13\epsilon)\left(1 - \frac{s_i}{s_1+s_2} - 2\sqrt{\epsilon} - \epsilon\right)|B| \\ &\le \frac{3}{2}|A_i| + (1-13\epsilon)|A_{3-i}|.\end{aligned}$$

We can thus apply Theorem 3.18 with ϵ^2 in place of ϵ. Note that since $s_2 \ge s_1$, both $|A_1|$ and $|A_2|$ are upper bounded by $(\frac{s_2}{s_1+s_2} + 2\sqrt{\epsilon})|B|$, and so the theorem implies that there exist arithmetic progressions P_1, P_2 with the same common difference of length

$$\begin{aligned}|P_i| &\le |A_1 \overset{\Gamma}{+} A_2| - |A_{3-i}| + 5\epsilon\left(\frac{s_2}{s_1+s_2} + 2\sqrt{\epsilon}\right)|B| \\ &\le |B| - \left(\frac{s_i}{s_1+s_2} - \epsilon - 2\sqrt{\epsilon}\right)|B| + 5\epsilon\left(\frac{s_2}{s_1+s_2} + 2\sqrt{\epsilon}\right)|B| \\ &\le \frac{s_i}{s_1+s_2}|B| + 4\sqrt{\epsilon}|B|\end{aligned}$$

such that

$$|A_i \setminus P_i| = |A_i| - |A_i \cap P_i| \le \epsilon|A_i|,$$

so case (b) holds. ∎

Putting everything together

The goal now will be to combine Theorem 3.28 with Propositions 3.29 and 3.30 in order to prove more technical versions of Theorems 3.26 and 3.25.

We begin by stating the precise structural result.

Theorem 3.31. *Let s_1, s_2, n be integers and $\alpha > 0$ a fixed real number satisfying*

$$s_2 \geq s_1 \geq 2^{10}\alpha^{-1}(s_1 + s_2)^{11/12}(\log n)^{1/4},$$

and let m be an integer such that

$$(1+\alpha)(s_1+s_2) \leq m < 2^{-108}\alpha^{12}s_1^{12}(s_1+s_2)^{-10}(\log n)^{-3}.$$

Suppose $X_1, X_2 \subset [n]$ are two uniformly chosen random sets with $|X_1| = s_1$, $|X_2| = s_2$ and $|X_1 + X_2| \leq m$. With probability at least $1 - \exp(-2^5 m^{1/6}(s_1+s_2)^{2/3}\sqrt{\log n})$ the following holds: there are sets $T_i \subset X_i$ of size $|T_i| \leq 2^{11}\alpha^{-1}m^{1/6}(s_1+s_2)^{2/3}\sqrt{\log n}$, such that $X_i \setminus T_i$ is contained in an arithmetic progression P_i of size

$$\frac{s_i m}{s_1+s_2} + 2^6 m^{13/12}(s_1+s_2)^{-1/6}(\log n)^{1/4},$$

where P_1 and P_2 have the same common difference.

Let us first see that this indeed implies Theorem 3.26.

Proof of Theorem 3.26. If $m = (1+o(1))(s_1+s_2) = s_1+s_2+o(s_1)$, we can apply an asymmetric version of Freĭman's $3k-4$ theorem directly and see that any sets X_1, X_2 satisfying the theorem hypotheses are contained in arithmetic progressions P_1 and P_2 with the same common difference of size

$$|P_i| = (1+o(1))s_i = (1+o(1))s_i m/(s_1+s_2).$$

If on the other hand there exists some absolute constant $\alpha > 0$ such that $m \geq (1+\alpha)(s_1+s_2)$, we can apply Theorem 3.31 instead. ∎

In order to prove Theorem 3.31, we need the following technical bound on the product of two specific binomial coefficients.

Lemma 3.32. *Let m, s and t be positive integers and let $1 \geq \alpha > 0$ such that $m \geq (1+\alpha)(s+t)$ and $s+t \geq 2^5\alpha^{-1}$. If $\epsilon > 0$ satisfies*

$$\frac{2^{10}\min(s^2,t^2)}{(s+t)^2m^2} \leq \epsilon \leq \frac{\alpha^2\min(s^2,t^2)}{2^{10}(s+t)^2},$$

then

$$\binom{\left(\frac{t}{s+t} - 2\sqrt{\epsilon} + 2\epsilon\right)m}{t}\binom{\left(\frac{s}{s+t} + 2\sqrt{\epsilon}\right)m}{s} \leq e^{-\epsilon(s+t)}\binom{\frac{sm}{s+t}}{s}\binom{\frac{tm}{s+t}}{t}. \tag{3.20}$$

Proof. Dividing by the binomial coefficients on the right hand side of (3.20) and taking the logarithm, we need to prove

$$\sum_{i=0}^{t-1}\log\left(1 - \frac{(2\sqrt{\epsilon}-2\epsilon)m}{\frac{tm}{s+t} - i}\right) + \sum_{j=0}^{s-1}\log\left(1 + \frac{2\sqrt{\epsilon}m}{\frac{sm}{s+t} - j}\right) \leq -\epsilon(s+t). \tag{3.21}$$

By using the bound $\log(1+x) \leq x - \frac{x^2}{2} + \frac{x^3}{3}$ valid on the interval $(-1,\infty)$, it suffices to prove the upper bound in (3.21) for the expression

$$\begin{aligned} &-2(\sqrt{\epsilon}-\epsilon)\sum_{i=0}^{t-1}\left(\frac{t}{s+t} - \frac{i}{m}\right)^{-1} - \frac{(2\sqrt{\epsilon}-2\epsilon)^2}{2}\sum_{i=0}^{t-1}\left(\frac{t}{s+t} - \frac{i}{m}\right)^{-2} \\ &+ 2\sqrt{\epsilon}\sum_{j=0}^{s-1}\left(\frac{s}{s+t} - \frac{j}{m}\right)^{-1} - 2\epsilon\sum_{i=0}^{s-1}\left(\frac{s}{s+t} - \frac{j}{m}\right)^{-2} + 4\epsilon^{3/2}\sum_{i=0}^{s-1}\left(\frac{s}{s+t} - \frac{j}{m}\right)^{-3}. \end{aligned} \tag{3.22}$$

Next, we are going to approximate the sums in (3.22) by integrals, using the bounds

$$\int_a^b f(x) - \frac{2(b-a)||f||_\infty}{n} \leq \frac{1}{n}\sum_{i=na}^{nb} f\left(\frac{i}{n}\right) \leq \int_a^b f(x)$$

which hold for any continuous, non-decreasing function f on the interval $[a,b]$. We start with the linear terms. Defining $K > (1+\alpha)$ by $m = K(s+t)$, we get

$$\begin{aligned} -2(\sqrt{\epsilon}-\epsilon)\sum_{i=0}^{t-1}\frac{1}{\frac{t}{s+t}-\frac{i}{m}} &\leq -2(\sqrt{\epsilon}-\epsilon)\left(m\int_0^{t/m}\frac{1}{\frac{t}{s+t}-x}dx - 2\frac{t}{m}\left(\frac{t}{s+t}-\frac{t}{m}\right)^{-1}\right) \\ &= -2(\sqrt{\epsilon}-\epsilon)\left(-m\log\left(1-\frac{s+t}{m}\right)-\frac{2}{K-1}\right) \\ &\leq -2m(\sqrt{\epsilon}-\epsilon)\log\left(\frac{K}{K-1}\right)+\frac{4\sqrt{\epsilon}}{K-1}, \end{aligned} \tag{3.23}$$

and similarly

$$2\sqrt{\epsilon}\sum_{j=0}^{s-1}\left(\frac{s}{s+t}-\frac{j}{m}\right)^{-1} \leq 2m\sqrt{\epsilon}\log\left(\frac{K}{K-1}\right). \tag{3.24}$$

For the quadratic terms, we see that

$$\begin{aligned} -\frac{(2\sqrt{\epsilon}-2\epsilon)^2}{2}\sum_{i=0}^{t-1}\left(\frac{t}{s+t}-\frac{i}{m}\right)^{-2} &\leq -\frac{(2\sqrt{\epsilon}-2\epsilon)^2}{2}\left(\frac{K(s+t)^2}{t(K-1)}-\frac{2K(s+t)}{t(K-1)^2}\right) \\ &\leq -\frac{2\epsilon K(s+t)^2}{t(K-1)}+\frac{4\epsilon^{3/2}K(s+t)^2}{t(K-1)}+\frac{8\epsilon K(s+t)}{t(K-1)^2}, \end{aligned} \tag{3.25}$$

and similarly for the one involving s,

$$-2\epsilon\sum_{i=0}^{s-1}\left(\frac{s}{s+t}-\frac{j}{m}\right)^{-2} \leq -\frac{2\epsilon K(s+t)^2}{s(K-1)}+\frac{4\epsilon K(s+t)}{s(K-1)^2}. \tag{3.26}$$

Finally, for the cubic term we see that

$$4\epsilon^{3/2}\sum_{i=0}^{s-1}\left(\frac{s}{s+t}-\frac{j}{m}\right)^{-3} \leq \frac{2\epsilon^{3/2}K(2K-1)(s+t)^3}{s^2(K-1)^2}. \tag{3.27}$$

Note that the $2m\sqrt{\epsilon}$ parts of the linear terms cancel out, while

$$2\epsilon K(s+t)\log\left(\frac{K}{K-1}\right) \leq \frac{2\epsilon K(s+t)^2}{\max(s,t)(K-1)},$$

which follows from the fact that $x\log(1+x^{-1}) \leq 1$ for all $x > 0$. On the other hand, because of $(s+t) \geq 2^5\alpha^{-1} \geq 2^5(K-1)^{-1}$ and the bounds on ϵ, the sum of all remaining positive term in Equations (3.23)–(3.27) can be upper bounded by $\frac{\epsilon K(s+t)^2}{\min(s,t)(K-1)}$, and hence we see that

$$(3.22) \leq -\frac{\epsilon K(s+t)^2}{\min(s,t)(K-1)} \leq -\epsilon(s+t),$$

which implies (3.21) and hence proves the statement. ■

We are now ready to prove Theorem 3.31.

Proof of Theorem 3.31. The upper bound on m in particular implies $m \leq s_1^2 \log n$, so let $\mathcal{A}$ be the family obtained from Theorem 3.28 applied with $G = \mathbb{Z}$, $h = 2$, $F_1 = F_2 = [n]$ and $2^{-10}\alpha^2 s_1^2/(s_1+s_2)^2 > \epsilon > 2^{10}s_1^2(s_1+s_2)^{-2}m^{-2}$ to be specified later. We claim that one of the following holds for every triple $(A_1, A_2, B) \in \mathcal{A}$:

(a) $|A_1| + |A_2| \leq (1-\epsilon)m$,

(b) $|A_i| > \frac{s_i m}{s_1+s_2} + 2\sqrt{\epsilon}m$ for some $i \in \{1,2\}$, or

(c) There are arithmetic progressions P_1, P_2 with the same common difference and sets T_1, T_2 such that $|P_i| \leq \frac{s_i m}{s_1+s_2} + 4\sqrt{\epsilon}m$, $|T_i| \leq \epsilon|A_i|$ and $A_i \setminus T_i \subseteq P_i$ for $i = 1, 2$.

Note first that we always have $|A_1| + |A_2| \leq (1+2\epsilon)m$ since by Theorem 3.28(b) applied with $G = \mathbb{Z}$, either there are at most $\epsilon^2|A_1||A_2|$ pairs $(a_1, a_2) \in A_1 \times A_2$ with $a_1 + a_2 \notin B$, and hence Proposition 3.29 together with $|B| \leq m$ gives the required upper bound on $|A_1| + |A_2|$, or $\max\{|A_1|, |A_2|\} < m/\log n$. Suppose neither (a) nor (b) hold, then by Proposition 3.30(b) we see that (c) holds.

We will now count the number of pairs of sets X_1, X_2 of size s_1 and s_2 respectively, satisfying $|X_1 + X_2| \leq m$ that do not have large intersections with arithmetic progressions in the sense of the theorem. To do this, recall that by Theorem 3.28(a), for any such pair, there exists a container triple $(A_1, A_2, B) \in \mathcal{A}$ such that $X_i \subset A_i$. We begin by giving an upper bound on the number of X_1, X_2 with containers satisfying property (a), that is, $|A_1| + |A_2| \leq (1-\epsilon)m$. Clearly there are at most $\sum_{\mathcal{A}}^{(a)} \binom{|A_1|}{s_1}\binom{|A_2|}{s_2}$ of these. By comparing $\binom{a-b}{c}\binom{b}{d}$ and $\binom{a-b-1}{c}\binom{b+1}{d}$ it is easy to check that an expression of this form has its maximum at $\binom{ca/(c+d)}{c}\binom{da/(c+d)}{d}$. So choosing $\epsilon = 2^8 m^{1/6}(s_1+s_2)^{-1/3}\sqrt{\log n} < 2^{-10}\alpha^2 s_1^2(s_1+s_2)^{-2}$ and using (3.13), we see that

$$\begin{aligned}\sum_{\mathcal{A}}^{(a)} \binom{|A_1|}{s_1}\binom{|A_2|}{s_2} &\leq |\mathcal{A}| \binom{(1-\epsilon)\frac{s_1 m}{s_1+s_2}}{s_1}\binom{(1-\epsilon)\frac{s_2 m}{s_1+s_2}}{s_2} \\ &\leq \exp(2^{21}\sqrt{m}\epsilon^{-2}(\log n)^{3/2} - \epsilon(s_1+s_2))\binom{\frac{s_1 m}{s_1+s_2}}{s_1}\binom{\frac{s_2 m}{s_1+s_2}}{s_2} \qquad (3.28) \\ &\leq \exp(-2^7 m^{1/6}(s_1+s_2)^{2/3}\sqrt{\log n})\binom{\frac{s_1 m}{s_1+s_2}}{s_1}\binom{\frac{s_2 m}{s_1+s_2}}{s_2}.\end{aligned}$$

We will now count pairs coming from containers of type (b). We will not make use of the fact that $s_2 \geq s_1$ so suppose without loss of generality that (b) holds for $i = 2$. Similar to the previous case, it suffices to give an upper bound for

$$\sum_{\mathcal{A}}^{(b)} \binom{|A_1|}{s_1}\binom{|A_2|}{s_2} \leq \sum_{\mathcal{A}}^{(b)} \binom{\left(\frac{s_1}{s_1+s_2} + 2\epsilon - 2\sqrt{\epsilon}\right)m}{s_1}\binom{\left(\frac{s_2}{s_1+s_2} + 2\sqrt{\epsilon}\right)m}{s_2}.$$

Noting that $\epsilon = 2^8 m^{1/6}(s_1+s_2)^{-1/3}\sqrt{\log n} > 2^{10}s_1^2(s_1+s_2)^{-2}m^{-2}$ we can apply Lemma 3.32 and see that

$$\binom{\left(\frac{s_1}{s_1+s_2} + 2\epsilon - 2\sqrt{\epsilon}\right)m}{s_1}\binom{\left(\frac{s_2}{s_1+s_2} + 2\sqrt{\epsilon}\right)m}{s_2} \leq e^{-\epsilon(s_1+s_2)}\binom{\frac{s_1 m}{s_1+s_2}}{s_1}\binom{\frac{s_2 m}{s_1+s_2}}{s_2},$$

and hence

$$\sum_{\mathcal{A}}^{(b)} \binom{|A_1|}{s_1}\binom{|A_2|}{s_2} \leq \exp(-2^7 m^{1/6}(s_1+s_2)^{2/3}\sqrt{\log n})\binom{\frac{s_1 m}{s_1+s_2}}{s_1}\binom{\frac{s_2 m}{s_1+s_2}}{s_2}. \qquad (3.29)$$

Finally, it remains to count the relevant X_1, X_2 with containers satisfying property (c). Observe that there are at most

$$\sum_{i=1}^{2} \sum_{s_i'=8\alpha^{-1}\epsilon(s_1+s_2)}^{s_i} \binom{|A_i|}{s_i - s_i'} \binom{\epsilon|A_i|}{s_i'} \binom{|A_{3-i}|}{s_{3-i}} \tag{3.30}$$

pairs of sets $X_i \subset A_i$ with $|X_i| = s_i$ that violate the theorem statement, since for at least one $\delta \in \{1,2\}$ there must be at least s_δ' elements in T_δ for some $s_\delta' \geq 8\alpha^{-1}\epsilon(s_1+s_2)$. Indeed, otherwise $X_i \setminus T_i \subset P_i$ with $|P_i| \leq \frac{s_i m}{s_1+s_2} + 4\sqrt{\epsilon}m$ and $|X_i \cap T_i| \leq 8\alpha^{-1}\epsilon(s_1+s_2)$ for both i. For any $d \leq c \leq a/4$, it holds that

$$\binom{a}{c-d}\binom{b}{d} \leq \binom{a}{c}\left(\frac{4bc}{ad}\right)^{d},$$

so applying this to each innermost summand of (3.30) gives

$$\begin{aligned}
\binom{|A_i|}{s_i - s_i'}\binom{\epsilon|A_i|}{s_i'}\binom{|A_{3-i}|}{s_{3-i}} &\leq \left(\frac{4\epsilon s_i}{s_i'}\right)^{s_i'} \binom{|A_1|}{s_1}\binom{|A_2|}{s_2} \\
&\leq \left(\frac{4\epsilon s_i}{s_i'}\right)^{s_i'} \binom{(1+2\epsilon)\frac{s_1 m}{s_1+s_2}}{s_1}\binom{(1+2\epsilon)\frac{s_2 m}{s_1+s_2}}{s_2} \\
&\leq \left(\frac{4\epsilon s_i}{s_i'}\right)^{s_i'} \left(1+4\alpha^{-1}\epsilon\right)^{s_1+s_2} \binom{\frac{s_1 m}{s_1+s_2}}{s_1}\binom{\frac{s_2 m}{s_1+s_2}}{s_2}
\end{aligned}$$

for every $i \in \{1,2\}$ and $s_i \geq s_i' \geq 8\alpha^{-1}\epsilon(s_1+s_2)$. Here, for the last inequality we used the bound $\binom{a}{c} \leq \left(\frac{a-c}{b-c}\right)^c \binom{b}{c}$ valid for any $a \geq b \geq c \geq 0$, as well as the upper bound $\alpha \leq 1$. Note that, by our choice of ϵ, we have $\max\{|\mathcal{A}|, s_1+s_2\} \leq \exp(\epsilon(s_1+s_2))$, hence summing (3.30) over all triples $(A_1, A_2, B) \in \mathcal{A}$ we obtain

$$\begin{aligned}
&\sum_{\mathcal{A}}^{(c)} \sum_{i=1}^{2} \sum_{s_i'=8\alpha^{-1}\epsilon(s_1+s_2)}^{s_i} \binom{|A_i|}{s_i - s_i'}\binom{\epsilon|A_i|}{s_i'}\binom{|A_{3-i}|}{s_{3-i}} \\
&\leq |\mathcal{A}|(1+4\alpha^{-1}\epsilon)^{s_1+s_2} \binom{\frac{s_1 m}{s_1+s_2}}{s_1}\binom{\frac{s_2 m}{s_1+s_2}}{s_2} \sum_{i=1}^{2} s_i \max_{s_i' \geq 8\alpha^{-1}\epsilon(s_1+s_2)} \left(\frac{4\epsilon s_i}{s_i'}\right)^{s_i'} \\
&\leq \exp\left(6\alpha^{-1}\epsilon(s_1+s_2)\right) 2^{-16\alpha^{-1}\epsilon(s_1+s_2)} \binom{\frac{s_1 m}{s_1+s_2}}{s_1}\binom{\frac{s_2 m}{s_1+s_2}}{s_2} \\
&\leq \exp(-4\alpha^{-1}\epsilon(s_1+s_2)) \binom{\frac{s_1 m}{s_1+s_2}}{s_1}\binom{\frac{s_2 m}{s_1+s_2}}{s_2} \\
&= \exp(-2^{10} m^{1/6}(s_1+s_2)^{2/3}\sqrt{\log n}) \binom{\frac{s_1 m}{s_1+s_2}}{s_1}\binom{\frac{s_2 m}{s_1+s_2}}{s_2}.
\end{aligned} \tag{3.31}$$

To conclude, note that bounds (3.28), (3.29) and (3.31) imply the probability we claimed in the statement since we can fix a single pair of disjoint arithmetic progressions of length $\frac{s_i m}{s_1+s_2}$ respectively with the same common difference and see that any of the $\prod \binom{s_i m/(s_1+s_2)}{s_i}$ pairs of s_i-subsets will have a sumset of size at most m. ■

We now turn to the technical version of Theorem 3.25.

Theorem 3.33. *Let G be an abelian group. Let s_1, s_2, n be integers satisfying*

$$s_2 \geq s_1 \geq \max\left(\sqrt{(s_1+s_2)\log n}, 2^{48}(\log n)^3 - s_2\right),$$

and let m be an integer such that

$$s_1 + s_2 \le m \le \min\left(\frac{s_1^2}{\log n}, \frac{(s_1+s_2)^2}{2^{48}(\log n)^3}\right).$$

Then for any $F_1, F_2 \subset G$ with $|F_i| = n$, it holds that the number of pairs of sets $(X_1, X_2) \in 2^{F_1} \times 2^{F_2}$ with $|X_i| = s_i$ and $|X_1 + X_2| \le m$ is at most

$$\exp\left(2^{10} m^{1/6}(s_1+s_2)^{2/3}\lambda^{2/3}\sqrt{\log n}\right)\binom{s_1(m+\beta)/(s_1+s_2)}{s_1}\binom{s_2(m+\beta)/(s_1+s_2)}{s_2},$$

where $\lambda = \min\left(\frac{m}{m-s_1-s_2}, \log(s_1+s_2)\right)$ and $\beta = \beta(m + 2^8 m^{7/6}(s_1+s_2)^{-1/3}\lambda^{-1/3}\sqrt{\log n})$.

Proof. We can apply Theorem 3.28 with s_1, s_2, m, n and $1/4 > \epsilon > 0$ to be specified later, let $\mathcal{A}$ be the family obtained this way. So for every pair of sets $(X_1, X_2) \in 2^{F_1} \times 2^{F_2}$ there exists a container triple $(A_1, A_2, B) \in \mathcal{A}$ such that $X_i \subset A_1$ and $B \subset X_1 + X_2$. Note that if we define $\beta = \beta(m + 4\epsilon m)$, it holds true that $|A_1| + |A_2| \le (1+2\epsilon)(m+\beta)$ for any pair (A_1, A_2) appearing in a container triple in $\mathcal{A}$. Indeed, by Theorem 3.28(b), we either have

$$|A_1| + |A_2| \le 2\max|A_i| < 2m/\log n \le (1+2\epsilon)(m+\beta),$$

or there are at most $\epsilon^2|A_1||A_2|$ pairs (a_1, a_2) such that $a_1 + a_2 \notin B$, and hence Proposition 3.29 gives the required bound. Hence the number of pairs (X_1, X_2) satisfying the theorem hypotheses is at most

$$\begin{aligned} &|\mathcal{A}| \max_{(A_1,A_2,B)\in\mathcal{A}} \binom{|A_1|}{s_1}\binom{|A_2|}{s_2} \\ &\le \exp\left(\frac{2^{21}\sqrt{m(\log n)^3}}{\epsilon^2}\right)\binom{\frac{s_1(1+2\epsilon)(m+\beta)}{s_1+s_2}}{s_1}\binom{\frac{s_2(1+2\epsilon)(m+\beta)}{s_1+s_2}}{s_2}. \end{aligned} \tag{3.32}$$

If $m/(m-s_1-s_2) \le \log(s_1+s_2)$ we can again apply the bound $\binom{a}{c} \le \left(\frac{a-c}{b-c}\right)^c\binom{b}{c}$ valid for any $a \ge b \ge c \ge 0$ to both binomials in (3.32) separately and see that it is at most

$$\exp\left(2^{21}\epsilon^{-2}\sqrt{m}(\log n)^{3/2} + 2\epsilon\lambda(s_1+s_2)\right)\binom{\frac{s_1(m+\beta)}{s_1+s_2}}{s_1}\binom{\frac{s_2(m+\beta)}{s_1+s_2}}{s_2}.$$

Suppose now that $m/(m-s_1-s_2) \ge \log(s_1+s_2)$, and note that this implies in particular $m = s_1 + s_2 + o(1)$. We compute

$$\begin{aligned} \log\left(\binom{(1+\delta)a}{b}\binom{a}{b}^{-1}\right) &= \sum_{i=0}^{b-1}\log\left(1 + \frac{\delta a}{a-i}\right) \\ &\le \delta a\int_0^b (a-x)^{-1}dx \\ &\le \delta a \log a. \end{aligned}$$

Applying this with $\delta = 2\epsilon$, $a = s_i(m+\beta)/(s_1+s_2)$ and $b = s_i$ and noting that $\beta \le 2m$, we can upper bound (3.32) by

$$\exp\left(2^{21}\epsilon^{-2}\sqrt{m}(\log n)^{3/2} + 2^3\epsilon\lambda(s_1+s_2)\right)\binom{\frac{s_1(m+\beta)}{s_1+s_2}}{s_1}\binom{\frac{s_2(m+\beta)}{s_1+s_2}}{s_2}.$$

Hence setting $\epsilon = 2^6 m^{1/6}(s_1+s_2)^{-1/3}\lambda^{-1/3}\sqrt{\log n} < 1/4$ implies

$$\beta(m + 4\epsilon m) = \beta(m + 2^8 m^{7/6}(s_1+s_2)^{-1/3}\lambda^{-1/3}\sqrt{\log n})$$

and the number of pairs (X_1, X_2) satisfying the theorem hypotheses is at most

$$\exp\left(2^{10}m^{1/6}(s_1+s_2)^{2/3}\lambda^{2/3}\sqrt{\log n}\right)\binom{\frac{s_1(m+\beta)}{s_1+s_2}}{s_1}\binom{\frac{s_2(m+\beta)}{s_1+s_2}}{s_2}.$$

■

3.3 Open problems and further outlook

Let us begin by discussing some further directions to explore connected to the results presented in the previous section. The container family constructed in Theorem 3.28 is applicable in settings that are a lot more general than the one used to prove Theorems 3.31 and 3.33, since it works for r-fold sumsets in arbitrary groups. So a natural question would be to ask whether something akin to these results holds in that framework as well. As already mentioned in the discussion at the end of Section 3.2, both theorems can be stated in a way to cover the semi-degenerate cases, as was shown in Corollary 3.27. That is, in the case of $r = 3$ for instance, it is possible to formulate statements in such a way that they are effective both in the case of tuples of three in general distinct sets (A, B, C), as well as for tuples of the form (A, A, B) for instance. The generalization to more than two summands seems very natural, so that the following two conjectures seem like realistic goals to prove. For a non-negative integer r and a subset A of some group G, we will write $rA = \{a_1 + \cdots + a_r : a_i \in A, i \in [r]\}$ as an abbreviation for the r-fold sumset of A, and use the convention $0A = \{0_G\}$, where 0_G denotes the identity element of G.

Conjecture 3.34. *Let G be an abelian group, $r \geq 2$ a fixed integer and $\lambda_1, \ldots, \lambda_r$ non-negative integers satisfying $\sum \lambda_i = r$. Furthermore, let n be a sufficiently large integer and $F_1, F_2, \ldots, F_r \subset G$ subsets of cardinality $|F_i| = n$. Then for all positive integers $s_r \geq s_{r-1} \geq \cdots \geq s_1 = \Omega(s_r)$ and m satisfying*

$$\lambda_1 s_1 + \cdots + \lambda_r s_r \leq m = o\left(s_1^{\lambda_1}\cdots s_r^{\lambda_r}/(\log n)^{O(1)}\right),$$

the number of r-tuples of sets $(X_1, X_2, \ldots, X_r)$ in $2^{F_1} \times 2^{F_2} \times \cdots \times 2^{F_r}$ such that $X_i = F_i$ if $\lambda_i = 0$ and $|X_i| = s_i$ otherwise, satisfying $|\lambda_1 X_1 + \cdots + \lambda_r X_r| \leq m$ is at most

$$2^{o(\lambda_1 s_1+\cdots+\lambda_r s_r)}\prod_{i=1}^{r}\binom{\frac{s_i}{\lambda_1 s_1+\cdots+\lambda_r s_r}(m+\beta)}{\delta_i s_i},$$

where $\beta = \beta((1 + o(1))m)$ and $\delta_i = 1$ whenever $\lambda_i > 0$ and 0 otherwise.

Conjecture 3.35. *Let $r \geq 2$ be a fixed integer and $\lambda_1, \ldots, \lambda_r$ non-negative integers satisfying $\sum \lambda_i = r$. Furthermore, let $n \geq s_r \geq s_{r-1} \geq \cdots \geq s_1 = \Omega(s_r)$ and m be positive integers satisfying*

$$\lambda_1 s_1 + \cdots + \lambda_r s_r \leq m = o\left(s_1^{\lambda_1}\cdots s_r^{\lambda_r}/(\log n)^{O(1)}\right).$$

Then for almost all sets $X_1, X_2, \ldots, X_r \subset [n]$ such that $|X_i| = s_i$ if $\lambda_i > 0$ and $|X_i| = n$ otherwise, and satisfying $|\lambda_1 X_1 + \cdots + \lambda_r X_r| \leq m$, there exist arithmetic progressions $P_1, \ldots, P_r$ with the same common difference of size

$$|P_i| = (1 + o(1))s_i m/(\lambda_1 s_1 + \cdots + \lambda_r s_r),$$

such that $|X_i \setminus P_i| = o(s_i)$ for all $i \in [r]$ satisfying $\lambda_i > 0$.

Note that as mentioned before, increasing the number of distinct summands will in general also increase the difficulty of proving the statement. Hence, if establishing Conjectures 3.34 and 3.35 proves to be too difficult, an intermediate step could be to look at sub-cases with r-tuples of $t < r$ distinct sets.

One potential path to prove these conjectures is clear, namely proving r-fold analogues of Proposition 3.9 and Theorem 3.18 which could then be used to get supersaturation and stability results similar to Propositions 3.29 and 3.30 for the r-fold case. Actually obtaining these analogues is something that is currently largely unexplored in the literature, but there are some possible directions. Notably, there does exist a version of the Balog-Szemerédi-Gowers (BSG) theorem, Theorem 3.10, for more than two sets, proved by Sudakov, Szemerédi and Vu in [102]. Since the proof of Shao's almost all version of the original BSG theorem in [96] consisted of essentially using the standard proof of Theorem 3.10 combined with Green's arithmetic removal lemma, it would be interesting to see if something similar can be established for the case of r possibly distinct sets that are extremely close in cardinality, using the r set version by Sudakov, Szemerédi and Vu. Alternatively, one could try to follow the exact route that was taken in Section 3.1. Relevant to this, DeVos, Goddyn, and Mohar in [29] proved a multiple set addition version of Kneser's theorem, and similarly, Lev [68] proved a version in the full-sumset setting of Proposition 3.17 for the special case rA. But one would still need to obtain a robust version of these similar to Theorem 3.15 and Proposition 3.17, which seems more difficult than in the setting of two summands.

Another possible further direction to explore is to translate the structural result, Theorem 3.26, even in the case of only two sets, into the setting of groups other than the integers. The structural description given by Green and Ruzsa [52] in their extension of Freĭman's theorem to general abelian groups provides a guideline for the structure of typical sets with bounded sumset in general abelian groups. As explored in Section 3.2 and previously in [19] and [20], the typical structure of sets with bounded sumset in the integers, in this case plain arithmetic progressions, is a lot simpler than the general one given by Freĭman's theorem, so one would expect the same to be the case in more general groups. Here Shao's version of the BSG theorem (in the case of two sets) could be combined directly with an appropriate version of Theorem 3.7 for groups other than the integers to immediately get a robust version that could then be used to obtain a stability result. Since Proposition 3.29 is already stated for arbitrary abelian groups, a supersaturation theorem can already be obtained directly, although subject to the β parameter defined in (3.2). An important natural case to explore is that of groups of prime order, $G = \mathbb{Z}/p\mathbb{Z}$, where sufficiently strong analogues of the Freĭman $3k-4$ theorem are already available and an equivalent statement of our main result might require less work to prove, since like in the integers no small non-trivial subgroups exist. An even less explored direction is to also translate the structural result to general groups, not necessarily abelian. Perhaps the more appropriate quest in this setting is to ask for the typical structure of approximate groups in the light of its structural characterization by Breuillard, Green and Tao in [16].

Finally, there are more specific problems that consider sumsets of distinct, but not independent sets. A classical example of this is to study pairs of sets A, B having small sumset when

$$B = \lambda * A = \{\lambda a : a \in A\}$$

is the *dilation* of A by some scaling factor $\lambda \in \mathbb{Z}$. The associated direct statement, that is, the answer to the question of how small a sum of dilations can be, was proved by Bukh in [17]. He established the following lower bound.

Theorem 3.36 ([17]). *Let $r \geq 2$ be an integer, $A \subset \mathbb{Z}$ a finite set of integers and $\lambda_1, \ldots, \lambda_r \in \mathbb{Z}$ be* *relatively prime integers. Then*

$$|\lambda_1 * A + \cdots + \lambda_r * A| \geq (\lambda_1 + \cdots + \lambda_r)|A| - o(|A|).$$

It then was a natural question to ask whether it is possible to improve the lower order term to something of the form $C_{\lambda_1,\ldots,\lambda_r}$, a constant only depending on the dilation factors. Initial research on establishing such a bound was focused on the specific case of $r = 2$ summands. Specifically, after partial progress by Cilleruelo, Hamidoune and Serra [23], Du, Cao, and Sun [31], Hamidoune and Rué [57] and Ljujić [73] that investigated specific instances where λ_1 was small and fixed (in fact, either 1 or 2) and λ_2 was either prime or a prime power, Balog and Shakan in [5] managed to prove the following general explicit result.

Theorem 3.37 ([5])**.** *Let* $1 \leq p < q$ *be relatively prime positive integers and let* $A \subset \mathbb{Z}$ *be a finite integer subset. Then*

$$|p * A + q * A| \geq (p+q)|A| - (pq)^{(p+q-3)(p+q)+1}$$

Later, Shakan in [95] was able to generalize this to the case of arbitrary r by proving the following.

Theorem 3.38 ([95])**.** *Let* $r \geq 2$ *be an integer,* $A \subset \mathbb{Z}$ *a finite set of integers and* $\lambda_1, \ldots, \lambda_r \in \mathbb{Z}$ *be relatively prime integers. Then there exists a constant* C *depending only on* $\lambda_1, \ldots, \lambda_r$ *such that*

$$|\lambda_1 * A + \cdots + \lambda_r * A| \geq (\lambda_1 + \cdots + \lambda_r)|A| - C.$$

Moreover, one can take $C = \frac{1}{3}\binom{r+1}{2}(|\lambda_1| \cdots |\lambda_r|)^{(r-1)(|\lambda_1|+\cdots+|\lambda_r|)^2+r-1}$.

Less is known about the relevant inverse statements, that is, the question of what is the structure of sets that obtain or are close to the lower bounds stated in Theorems 3.37 and 3.38. In the language of Theorem 3.36, Cilleruelo, Hamidoune and Serra managed to prove such an inverse result in the case $r = 2$, $\lambda_1 = 1$ and $\lambda_2 = p$ for a prime p when A is large enough with respect to p. Specifically, they proved the following.

Theorem 3.39 ([23])**.** *Let* p *be an odd prime and let* A *be a finite subset of the integers satisfying* $|A| \geq 3(p-1)^2(p-1)!$ *and*

$$|A + p * A| = (p+1)|A| - \lceil p(p+2)/4 \rceil.$$

Then, up ot affine transformations, it holds that there exists an integer n *such that*

$$A = p * \{0, 1, \ldots, n\} + \{0, 1, \ldots, (p-1)/2\}.$$

Can one prove an approximate structure result similar in scope to Theorem 3.26 for this specific problem? Since we clearly have $|A| = |p * A|$, a fruitful approach could be to investigate whether Theorem 3.39 yields a robust version using Shao's version of the Balog-Szemerédi-Gowers theorem. This could then be used to prove an approximate structure result similar to Theorem 3.26.

Chapter 4

Sidon set systems

The main contributions of this chapter are extensions of classical results on Sidon sets to a generalization for set systems. In particular, statements are proved on the maximal size of such a set system, as well as for what ranges of p a p-random set system will be one. All original work presented in this chapter is based on [25] and was done jointly with Javier Cilleruelo and Oriol Serra.

In this chapter we return to the topic of linear equations already introduced in Chapter 1. Specifically, we will have a closer look at the *Sidon equation*. Recall that a finite set A in an abelian group G is called a *Sidon set* if the equation $a + b = c + d$ has no nontrivial solutions in A, that is, this happens if and only if $\{a, b\} = \{c, d\}$. Note that as already mentioned in Chapter 1, a nontrivial solution to the Sidon equation does not need to be proper, meaning that we could have $a = b$. Looking at the sumset $A + A$, we also get the following equivalent definition: A is a Sidon set if and only if $|A + A| = \binom{|A|+1}{2}$. In this sense, Sidon sets somewhat represent the opposite of those investigated in Chapter 3. We will investigate the following generalization of Sidon sets to set systems.

Definition 4.1. Let I be some index set and let $\mathcal{A} = \{A_i : i \in I, A_i \subseteq G\}$ be a family of subsets of an abelian group G. We say that $\mathcal{A}$ is a *Sidon system* if

$$A_i + A_j = A_{i'} + A_{j'} \implies \{i, j\} = \{i', j'\}.$$

We will mostly restrict ourselves to uniform set systems of k-subsets, and in that sense one can recover "normal" Sidon sets by setting $k = 1$. A central problem in the investigation of Sidon sets in the integers is the following question: For an integer n, what is the largest cardinality of a Sidon set $A \subset [n]$? Let us introduce the following definition.

Definition 4.2. Given integers $n > k \geq 1$, we denote by $F_k(n)$ the largest cardinality of a Sidon system $\mathcal{A} \subseteq \binom{[n]}{k}$.

Before presenting new results regarding the case $k \geq 2$, let us recall some theorems regarding the classical case of $k = 1$, Sidon sets. It is not too difficult to prove the following upper bound, which was done initially by Erdős and Turán in [33], although the following proof is due to Lindström [72].

Theorem 4.3 ([33, 72])**.** *Let n be a positive integer and let $A \subset [n]$ be a Sidon set. Then*

$$|A| < n^{1/2} + n^{1/4} + 1.$$

Proof. Suppose $k = |A|$ and denote the elements of A in an ordered fashion as

$$1 \leq a_1 < a_2 < \cdots < a_k \leq n.$$

Note that for all $i \neq j \in [k]$, A being a Sidon set also implies that the difference $a_i - a_j$ is uniquely represented in $A - A$. Indeed, if $a_i - a_j = a_{i'} - a_{j'}$, then $a_i + a_{j'} = a_{i'} + a_j$ and since we assumed $i \neq j$, we must also have $i' \neq j'$, which results in a contradiction to A being a Sidon set. Using this observation, we see that for any integer $u < k$ the differences

$$\begin{gathered} a_2 - a_1, a_3 - a_2, \dots, a_k - a_{k-1} \\ a_3 - a_1, a_4 - a_2, \dots, a_k - a_{k-2} \\ \vdots \\ a_{u+1} - a_1, a_{u+2} - a_2, \dots, a_k - a_{k-u} \end{gathered} \tag{4.1}$$

are all pairwise distinct. The i-th row of this contains $k - i$ differences, and hence summing up over all u we get

$$\sum_{i=1}^{u} k - i = uk - u(u+1)/2$$

distinct elements. Note that since A consists of integers, a_i and a_{i+1} differ by at least 1, and hence for any $i, j \in [k]$ it will hold that $a_j - a_i \geq j - i$. In particular all elements in (4.1) are positive, and since they are pairwise distinct their sum must be at least

$$\sum_{i=1}^{uk-u(u+1)/2} i = (uk - u(u+1)/2)(uk + 1 - u(u+1)/2)/2.$$

Conversely, summing up only the differences in the i-th row results in the upper bound

$$\sum_{j=i+1}^{k} a_j - \sum_{j=1}^{k-i} a_j = \sum_{j=k-i+1}^{k} a_j - \sum_{j=1}^{i} a_j < in,$$

since $i \geq 1$ and so we get an upper bound for the whole sum of

$$\sum_{i=1}^{u} in = nu(u+1)/2.$$

Comparing the upper and lower bound with $u = n^{1/4}$ will result in $k < n^{1/2} + n^{1/4} + 1$ which is what we wanted to show. ■

In [33], Erdős and Turán were only able to construct Sidon sets of size $\Omega(n^{1/2})$, but Ruzsa [88], Bose [15] and Singer [98] showed that $n^{1/2}$ is indeed the correct value for the main term. The following construction is that of Ruzsa.

Theorem 4.4 ([88, 15, 98])**.** *There exist Sidon sets $A \subset [n]$ satisfying $|A| \geq (1 - o(1))n^{1/2}$.*

Proof. Let p be an odd prime. We will construct a Sidon set $A \subset [p(p-1) - 1]$ of cardinality $p - 1$. Since the ratio of consecutive primes tends to 1, this implies the theorem statement. Let g be a primitive root modulo p, then we define A to be the set containing all integers a_i, $i \in [p-1]$ that satisfy

$$a_i \in [p^2 - p - 1], \quad a_i \equiv i \pmod{p-1}, \quad \text{and} \quad a_i \equiv g^i \pmod{p},$$

which exist by the Chinese remainder theorem. Let us check that A is indeed a Sidon set, so suppose that for some fixed integer k there are indices i, j, s, t in $[p-1]$ such that $a_i + a_j = a_s + a_t = k$. Since the indices satisfy

$$s + t \equiv i + j \equiv a_i + a_j = k \pmod{p-1},$$

we see that using $g^{p-1} \equiv 1 \pmod p$ it will hold that $i+j$ and $s+t$ satisfy

$$g^{i+j} \equiv g^{s+t} \equiv g^k \pmod p.$$

Using this, we see that

$$(x-a_s)(x-a_t) \equiv x^2 - kx + g^k \equiv (x-a_i)(x-a_j) \pmod p,$$

and since factorization modulo p is unique, the elements must be congruent modulo p, for instance

$$a_i \equiv a_s \pmod p. \tag{4.2}$$

But by definition of A, this implies $g^i \equiv g^s \pmod p$, which in turn tells us that $i \equiv s \pmod{p-1}$. But again using the definition of A, we then get

$$a_i \equiv a_s \pmod{p-1}. \tag{4.3}$$

Since a_i and a_s are both contained in $[p^2-p-1]$, Equations (4.2) and (4.3) and the Chinese remainder theorem imply $a_i = a_s$, and $a_i + a_j = k = a_s + a_t$ then implies $a_j = a_t$, so A is a Sidon set. ■

While Theorems 4.3 and 4.4 settle the asymptotic behavior of $F_1(n)$, it is still an open problem (originally valued at \$500 by Erdős) whether the lower order term is bounded or not. A good reference for results on Sidon sets and some of their generalizations is O'Bryant's survey [80]. Ruzsa's construction can be combined with some manual ones for small k to get the following construction of Sidon sets of size k in $[2k^2]$ valid for any k, which we will need later.

Corollary 4.5. *For any $k \geq 1$, there exists a Sidon set $A \subset [2k^2]$ of size $|A| = k$.*

Proof. For $k \leq 8$, one can for instance take the powers of 2, $A = \{1, 2, \dots, 2^{k-1}\}$. We see that $2^{k-1} \leq 2k^2$ indeed holds here. For the interval $9 \leq k \leq 24$, it is easy to check manually that the smallest prime $p \geq k+1$ has size at most $\sqrt{2}k$, and hence $p(p-1) \leq 2k^2$. So any subset of Ruzsa's construction of size $k \leq p-1$ is a valid choice. For $k \geq 25$, results by Nagura [78] show that the next smallest prime is always less than $6k/5 < \sqrt{2}k$, and again one can use Ruzsa's construction. ■

Our first main results are upper and lower bounds for $F_k(n)$ with $k \geq 2$, that is, versions of Theorems 4.3 and 4.4 for Sidon systems of k-subsets of the first n integers.

Theorem 4.6. *For $2 \leq k < n$ we have*

$$F_k(n) \leq \binom{n-1}{k-1} + n - k.$$

Our lower bounds are slightly different depending on k. For $k = 2$ we are able to show that the upper bound in Theorem 4.6 is sharp, while for $k = 3$ the main term will indeed be of the form $n^{k-1}/(k-1)!$. For $k \geq 4$ we are only able to show that n^{k-1} is the correct order of magnitude.

Theorem 4.7. *It holds that*

$$\begin{aligned} &F_2(n) = 2n-3, \textit{ for } n > 2;\\ &F_3(n) \geq n^2/2 - O(n), \textit{ for } n > 3\\ &F_k(n) = \Omega_k(n^{k-1}), \textit{ for } n > k \geq 4. \end{aligned}$$

We suspect that $n^{k-1}/(k-1)!$ is indeed asymptotically the correct value for $F_k(n)$, and moreover we conjecture that one can take a very specific structure to achieve this. For integers $n > k \geq 2$, define the set system

$$\binom{[n]}{k}_0 = \{A \subset \{0, 1, \ldots, n\} : |A| = k, 0 \in A\}.$$

Then we pose the following conjecture on the nature of maxium cardinality Sidon systems of k-subsets for arbitrary but fixed $k \geq 3$.

Conjecture 4.8. *Let $n > k \geq 3$, and suppose $\mathcal{F} \subset \binom{[n]}{k}$ is any family of k-subsets of the first n integers such that for every $A \in \binom{[n]}{k}_0$ it holds that*

$$|\{x \in \mathbb{Z} : A + x \in \mathcal{F}\}| \leq 1.$$

Then one can remove $o(n^{k-1})$ sets from $\mathcal{F}$ to make it a Sidon system. In particular,

$$F_k(n) \sim n^{k-1}/(k-1)!.$$

Note that we will indeed prove this stronger type of result in order to establish the lower bound in the $k = 3$ case. We will also later discuss some progress on solving Conjecture 4.8. Before proceeding to the proofs of Theorems 4.3 and 4.4, let us discuss some related questions and results. While we will treat this problem very clearly as a generalization of Sidon sets in the integers, another interpretation is to consider it as a special case of the "normal" Sidon problem in the semigroup of integer subsets. Hence, one could abstract even further and forget about the specific semigroup. We would thus ask more generally when a finite set A in a general semigroup S is free of solutions to the Sidon, or any specific linear equation. Similarly, the topics presented in Chapter 3 could be reinterpreted in this more general way. This extension of additive problems in the integers or in additive groups to the monoid of sumsets has been considered in the literature. For instance, Cilleruelo, Hamidoune and Serra [24] proved analogues of Theorem 3.3, the Cauchy–Davenport theorem, and Theorem 3.4, Vosper's theorem, in this setting. Let us return to mentioning some related results in the original interpretation in the integers. As will become apparent in the proofs, an important connected question is whether a certain set can be expressed as a sumset in multiple ways. Alon [1] used probabilistic arguments combined with spectral techniques to improve earlier bounds by Green [49] on the maximal cardinality of subsets of a cyclic group of prime order that cannot be expressed as a sumset. Fan and Tringali [36] use tools from factorization theory to give (among other results) necessary and sufficient conditions for certain subsets of integers to be written as sumsets in more than one way. Selfridge and Straus [94] showed that the representation function $r_A(n) = |\{(a, a') \in A \times A : n = a + a'\}|$ of a subset A in a field of characteristic zero determines the set. They also considered the general case of h-fold sumsets hA and gave necessary conditions when the representation function of this sumset completely determines the set A. These results were later generalized by Gordon, Fraenkel and Straus [46] to torsion free abelian groups. For a more detailed look on these problems, see also the recent survey by Fomin [37]. In contrast to these results, in the asymmetric case, the representation function does not in general determine the summands, even in the case of twofold sumsets $A + B$.

The remainder of this chapter is structured as follows. Section 4.1 contains the proofs of Theorems 4.3 and 4.4, the upper and lower bounds on $F_k(n)$. In Section 4.2 we will consider the topic of Sidon systems through a probabilistic lense. The main result presented here determines the threshold probability p for when the binomial random set $\binom{[n]}{k}_p$ is a Sidon system almost surely. We will also prove a so-called *relative density* result (cf. the discussion at the beginning of Section 2.1) that is partially conditional on Conjecture 4.8. Finally, in Section 4.3 we will discuss some open problems on generalizations of Sidon sets, as well as present some partial results on them.

4.1 The proofs of Theorems 4.6 and 4.7

It will often be helpful to have an ordering of $\binom{[n]}{k}$. We denote by $\preceq$ the lexicographic order of k-subsets of $[n]$, namely, $A \preceq B$ if and only if $\min(A\Delta B) \in A$ or $A = B$, where $A\Delta B$ denotes the symmetric difference of A and B. Similarly, we define $\prec$ such that $A \prec B$ if and only if $A \preceq B$ and $A \neq B$. We will also use $\preceq$ (resp. $\prec$) to denote the induced lexicographic order on tuples of k-subsets.

The upper bound

We will start by proving Theorem 4.6, the upper bound on $F_k(n)$, with k and n as defined in the theorem statement. Let $\mathcal{A}$ be a Sidon system of k-subsets in $[n]$. For each set A in $\binom{[n-1]}{k}_0$, define

$$\mathcal{A}(A) = \{x \in [n] : x + A \in \mathcal{A}\}.$$

We have

$$|\mathcal{A}| = \sum_{A \in \binom{[n-1]}{k}_0} |\mathcal{A}(A)|. \tag{4.4}$$

Denoting by $Z_+ = Z \cap \mathbb{N}$ the set of positive numbers in a set Z of integers, we clearly have

$$|\mathcal{A}(A)| \leq |(\mathcal{A}(A) - \mathcal{A}(A))_+| + 1. \tag{4.5}$$

We observe that if $A \neq B$ are sets in $\binom{[n-1]}{k}_0$, then

$$(\mathcal{A}(A) - \mathcal{A}(A))_+ \cap (\mathcal{A}(B) - \mathcal{A}(B))_+ = \emptyset. \tag{4.6}$$

Indeed, suppose there are $x' > x$ in $\mathcal{A}(A)$ and $y > y'$ in $\mathcal{A}(B)$ such that $x' - x = y - y'$. Then this implies $(x + A) + (y + B) = (x' + A) + (y' + B)$, which is a violation to the Sidon property of $\mathcal{A}$. Equation 4.6 directly implies that

$$\left| \bigcup_{A \in \binom{[n-1]}{k}_0} (\mathcal{A}(A) - \mathcal{A}(A))_+ \right| = \sum_{A \in \binom{[n-1]}{k}_0} |(\mathcal{A}(A) - \mathcal{A}(A))_+|. \tag{4.7}$$

Since $\max(A) \geq k - 1$ for any set A in $\binom{[n-1]}{k}_0$ and $x + \max(A) \leq n$ for any $x \in \mathcal{A}(A)$, we clearly have $\mathcal{A}(A) \subseteq [n - k + 1]$, and hence $(\mathcal{A}(A) - \mathcal{A}(A))_+ \subseteq [n - k]$. Together with (4.6), this implies

$$\left| \bigcup_{A \in \binom{[n-1]}{k}_0} (\mathcal{A}(A) - \mathcal{A}(A))_+ \right| \leq n - k. \tag{4.8}$$

Combining Equations 4.4, 4.5, 4.7, and 4.8, we conclude

$$\begin{aligned}|\mathcal{A}| &\leq \binom{n-1}{k-1} + \sum_{A \in \binom{[n-1]}{k}_0} |(\mathcal{A}(A) - \mathcal{A}(A))_+| \\ &= \binom{n-1}{k-1} + \left| \bigcup_{A \in \binom{[n-1]}{k}_0} (\mathcal{A}(A) - \mathcal{A}(A))_+ \right| \\ &\leq \binom{n-1}{k-1} + n - k.\end{aligned}$$

■

The lower bounds

We will now prove Theorem 4.4, considering the different cases for k separately. For $k = 2$, one can quickly check that

$$\mathcal{A} = \{\{1, 1+i\} : i = 1, \ldots, n-1\} \cup \{\{n-i, n\} : i = 1, \ldots, n-2\}$$

is indeed a Sidon system whose cardinality matches the upper bound proved in the previous subsection. ■

In order to determine the precise asymptotics for $F_3(n)$, we employ a lengthy case analysis, so we postpone the presentation of this until the end of the section in order to showcase the following construction, valid for any $k \geq 3$, that leads to the lower bounds on $F_k(n)$ for $k \geq 4$ in Theorem 4.7. Let $k \geq 3$ be a fixed integer, and let $A = \{a_0, a_1, \ldots, a_{k-1}\}$ be a Sidon set such that

$$0 = a_0 < a_1 < \cdots < a_{k-1} \leq 2k^2,$$

which exists by Corollary 4.5. We can assume that $n \geq 2(2k^2 + 1)$. For $i = 1, \ldots, k-1$, denote by I_i the interval

$$I_i = \left(\frac{n}{a_{k-1} + 1} \cdot [a_i, a_i + 1/2) \right) \cap \mathbb{N},$$

with cardinality

$$|I_i| = \left\lfloor \frac{n}{2(a_{k-1} + 1)} \right\rfloor.$$

Let $I_0 = \{0\}$. Since A is a Sidon set, for all $i, j, i', j' \in [0, k-1]$ we have

$$(I_i + I_j) \cap (I_{i'} + I_{j'}) = \emptyset,$$

unless $\{i, j\} = \{i', j'\}$. Consider the family

$$\mathcal{B} = \{\{b_0, \ldots, b_{k-1}\} : b_i \in I_i\} \subseteq \binom{[n-1]}{k}_0,$$

which has cardinality

$$|\mathcal{B}| \geq \left\lfloor \frac{n}{2(a_{k-1} + 1)} \right\rfloor^{k-1} \geq \left(\frac{n}{2(2k^2 + 1)} \right)^{k-1}.$$

We now prove that $\mathcal{B}$ is a Sidon system, which implies the lower bound on $F_k(n)$.

Let $U = \{0 < u_1 < \cdots < u_{k-1}\}$ and $V = \{0 < v_1 < \cdots < v_{k-1}\}$ be sets in $\mathcal{B}$. We will show that $U + V$ determines univocally the sets U, V. We have

$$U + V \subseteq \bigcup_{0 \le i \le j \le k-1} (I_i + I_j).$$

Since the intervals $I_i + I_j$ are pairwise disjoint, each of them contains at most the two elements

$$u_i + v_j,\ u_j + v_i \in I_i + I_j$$

from $U + V$. In particular, for any $i \in [1, k-1]$, the set $(U + V) \cap I_i$ contains exactly one element if and only if $u_i = v_i$, and clearly, if this happens for all i, then $U = V$ and the set is determined univocally by the elements in $(U + V) \cap I_i$. So assume that this is not the case, and let i_0 denote the least positive integer such that $(U + V) \cap I_{i_0}$ contains two elements. Hence for all $0 \le i < i_0$, the elements $u_i = v_i$ are determined by $U + V$. Without loss of generality, we can assume that $U \preceq V$. Therefore

$$u_{i_0} = \min((U + V) \cap I_{i_0}) < \max((U + V) \cap I_{i_0}) = v_{i_0},$$

and so both u_{i_0} and v_{i_0} are determined by $U + V$. Let $i > i_0$. We have

$$(U + V) \cap I_i = \{u_i, v_i\} \text{ and } (U + V) \cap (I_{i_0} + I_i) = \{u_{i_0} + v_i, u_i + v_{i_0}\}.$$

So by subtracting u_{i_0} and v_{i_0} from the elements in $(U + V) \cap (I_{i_0} + I_i)$, we can determine u_i and v_i. This shows that every sumset in $\mathcal{B} + \mathcal{B}$ can be written uniquely as a sum of two sets from $\mathcal{B}$, and so $\mathcal{B}$ is a Sidon system. Finally, we can shift every set in $\mathcal{B}$ by 1 such that the resulting family $\mathcal{B}'$ is a Sidon system of k-subsets of $[n]$ with the same cardinality as $\mathcal{B}$. ■

Let us now present the case analysis that will establish the sharp asymptotics for $F_3(n)$. Essentially, we are going to prove the $k = 3$ case of Conjecture 4.8, that is, defining

$$Q := \left\{ (X, Y, V, W) \in \binom{[n-1]}{3}_0^4 : X + Y = V + W \text{ and } \{X, Y\} \neq \{V, W\} \right\},$$

we establish the upper bound $|Q| = O(n)$. Removing one set per such quadruple from $\binom{[n]}{3}_0$ and shifting everything by 1 will thus result in a large Sidon system of 3-sets. In order to reduce the number of configurations, we introduce the following notion. For any set $X = \{0 < x < x'\} \in \binom{[n]}{k}_0$, let $\overline{X}$ denote the set

$$\overline{X} = x' - X = \{0 < x' - x < x'\}.$$

We will refer to $\overline{X}$ as the *dual* of X. Let $X = \{0 < x < x'\}, Y, V, W \in \binom{[n-1]}{k}_0$ such that $X + Y = V + W$. Without loss of generality, we can assume that $X \preceq Y$, $X \preceq V$, and $V \preceq W$. Furthermore, $X = V$ only if $Y \prec W$. We begin by showing that the following orderings cannot occur among quadruples in Q:

$$\begin{array}{lll} X = Y \prec V = W, & X = Y = V \prec W, & X \prec Y = V = W, \\ X = Y \prec V \prec W, & X \prec Y \prec V = W, & X \prec V = W \prec Y, \\ X \prec V = Y \prec W, & X \prec V \prec Y = W, & X \prec Y \prec V \prec W. \end{array}$$

We will quickly go through all of them and see by contradiction that they cannot belong to Q.

$\boldsymbol{X = Y \prec V = W.}$ This implies $x = v$ as they are the smallest element on either side, and also $2x' = 2v'$ as the largest element on either side, which implies $x' = v'$, and hence $X = Y = V = W$, a contradiction.

$\boldsymbol{X = Y = V \prec W.}$ This implies $2x' = x' + w'$ and hence $x' = w'$. Then we must have $w > x$, but now $x' + w$ is strictly larger than $x' + x$ and hence cannot be matched by any element of $X + X$, a contradiction. Note that $w > x$ also implies $x' - x > w' - w$, that is, $\overline{W} \prec \overline{X}$, and so the case $X \prec Y = V = W$ is also not possible.

$\boldsymbol{X = Y \prec V \prec W.}$ We have $v \leq w$, and hence $x = v$ as the smallest element on either side. This implies $x' < v'$ which in turn means $x' > w'$ since $2x' = v' + w'$. But then $v' + w' > v' + w > v' > w' > w \geq v > 0$ is a chain in $V + W$ of at least 6 elements, while $|X + X| \leq 6$. Hence we must have $v = w$ which implies $w' > v'$, a contradiction.

$\boldsymbol{X \prec Y \prec V = W.}$ We have $x \leq y \leq v$ and hence $x = y = v$ because the smallest elements coincide. But then $x' < y' < v'$ which contradicts $x' + y' = 2v'$.

$\boldsymbol{X \prec V = W \prec Y.}$ We have $x \leq v \leq y$ and hence $x = v \leq y$, which implies $x' < v'$, and hence $y' > v' > x'$. But then $x + y'$ is strictly larger than $v + v'$ and hence cannot be matched by any element of $V + V$.

$\boldsymbol{X \prec V = Y \prec W.}$ Since $x' + y' = y' + w'$ we know that $x' = w'$. Furthermore, $x \leq y \leq w$, and hence $x = y$ as smallest elements. This implies $x' < y'$, which in turn means $w' < y'$, and hence $w > y$. But then $y' + w$ is strictly larger than $y' + x$ and $x' + y$ and hence cannot be matched by any element of $X + Y$.

$\boldsymbol{X \prec V \prec Y = W.}$ We have $x' + y' = v' + y'$ and hence $x' = v'$. Furthermore $x \leq v \leq y$, which implies $x = v$ and hence $x' < v'$, a contradiction.

$\boldsymbol{X \prec Y \prec V \prec W.}$ Since $x \leq y \leq v \leq w$, we have $x = y = v \leq w$ as smallest elements, which implies $x' < y' < v'$. But now $v' + w$ will be strictly larger than both $x' + y$ and $y' + x$ and hence cannot be matched by any element of $X + Y$.

This settles the cases without solutions. The remaining three cases are

$$X = V \prec Y \prec W, \quad X \prec V \prec Y \prec W, \quad X \prec V \prec W \prec Y,$$

each of which will actually have solutions, and the analysis is more involved than in the previous cases.

$\boldsymbol{X = V \prec Y \prec W.}$ Before splitting this case into further subcases, we make some general assertions. Since $x' + y' = x' + w'$, we know that $y' = w'$, and hence $w > y$ (in particular, $w > x$). This means that $x' + w > x' + y$, and hence $x' + w \in \{y', x + y'\}$. Since $w > x$, this implies in particular that $y' > x'$. It also tells us that $x + y' > x' + y$. Finally, with the same arguments we also see that $x + w \in \{x', y', x' + y\}$ and $w \in \{x', x + y, x' + y\}$. Let us now consider specific subcases.

Case a) $\boldsymbol{x' + w = y'}$. This implies that $y' > x' + y$. We get the following diagrams that represent the sumsets. In this and all following diagrams, we will always omit the unique greatest element $x' + y'$, as well as the unique least element 0.

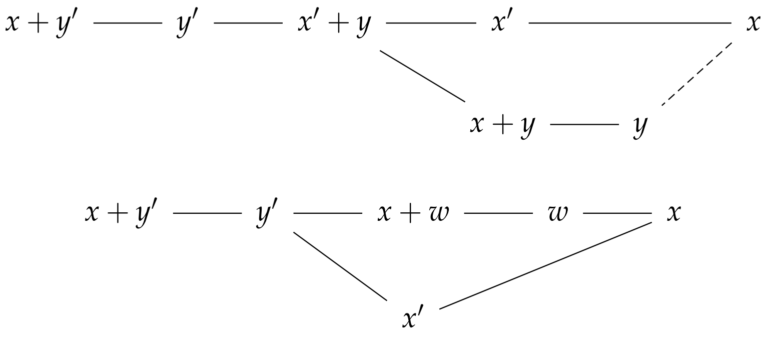

We see that there is at least one element strictly between y' and x' in the top diagram, and so it must hold that $x+w > x'$, in particular $x+w = x'+y$, so in particular $x'+y > w$, and hence $w \in \{x', x+y\}$. We also see that $w = x'$ if and only if $x = y$.

Case a)i. $w = x'$. Then $x = y$ and we have a chain of exactly 7 elements in the bottom diagram. This implies that $x' = 2x$ and we get the following solution: x can be freely chosen from $[1, (n-1)/4]$, $y = v = x$, $x' = v' = w = 2x$, and $y' = w' = 4x$.

Case a)ii. $w = x+y \neq x'$. In particular, we have $y > x$. We see that $x'+y = x+w = 2x+y$ and hence $x' = 2x$. Also, since $w > y > x$, we must have that $x' = y$, otherwise it cannot be matched. This leads to the solution $v = x$, $x' = y = v' = 2x$, $w = 3x$, and $y' = w' = 5x$, where x can be chosen freely from the positive integers smaller than $(n-1)/5$.

Case b) $x'+w = x+y'$. This does not give us enough new information, so we go directly into three further subcases, depending on the assignment of $x+w$. Recall that $x+w \in \{x', y', x'+y\}$, so there are three different cases to consider.

Case b)i. $x+w = x'$. This means that x' is strictly larger than $x+y$, but strictly smaller than y'. We get the following diagrams.

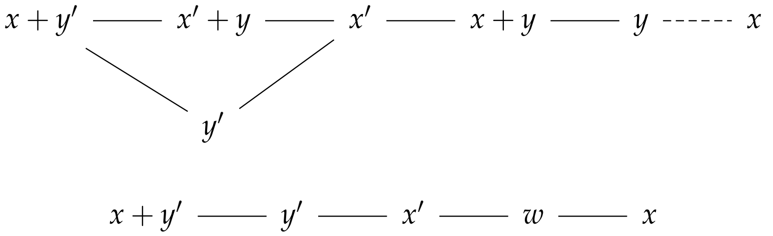

This directly implies that we must have $y = x$, since there is only one element strictly between x' and x in the bottom diagram. We thus have $w = x+y = 2x$ and $x'+y = y'$. Solving this, we get the solution $y = v = x$, $w = 2x$, $x' = v' = 3x$, and $y' = w' = 4x$, where x can be chosen freely among the positive integers smaller than $(n-1)/4$.

Case b)ii. $x+w = y'$. In particular, this means that $y' > x+y$. We get the following diagrams,

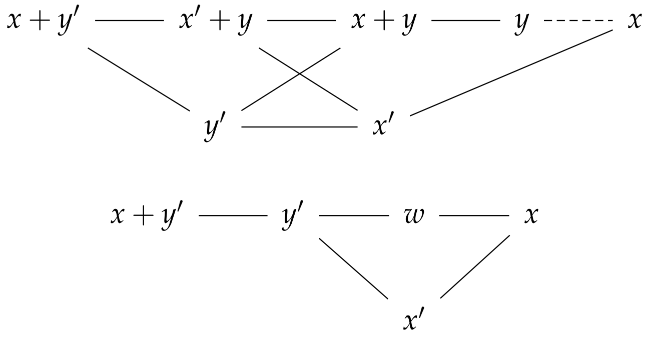

We see that there is no element strictly between y' and $x+y'$, and hence $y' \geq x'+y$.

Suppose $y' > x'+y$. This implies that there is an element strictly between x' and y', and hence we must have $w = x'+y > x'$. So the bottom diagram will now be a chain of 7 elements, and so we must have $y = x$ and $x+y = x'$. This simplifies to the solution $y = v = x$, $x' = v' = 2x$, $w = 3x$, and $y' = w' = 4x$, where x can be chosen freely among the positive integers smaller than $(n-1)/4$.

Suppose now that $y' = x'+y = x+w$. Since $y \geq x$, this implies that $w \geq x'$ with equality if and only if $x = y$. So we have two cases.

If $w = x'$ and $x = y$, then we must have $x' = x+y = 2x$, and we get the solution $y = v = x$, $x' = v' = w = 2x$, and $y' = w' = 3x$, where x can be chosen among the positive integers smaller than $(n-1)/3$.

If $w > x'$ and $y > x$, then this implies that $x+y = w$ and $y = x'$, which leads to the solution $v = x$, $x' = y = v' = 2x$, $w = 3x$, and $y' = w' = 4x$, with x able to be chosen freely from the positive integers smaller than $(n-1)/4$.

Case b)iii. $x+w = x'+y \neq y'$. Since $y \geq x$ we have $w \geq x'$ and $y = x$ if and only if $w = x'$. So since $x'+y \neq y'$, the bottom side of the diagram is now a chain of 7 or 8 elements (we don't know whether $w = x'$). Specifically we get the following diagrams.

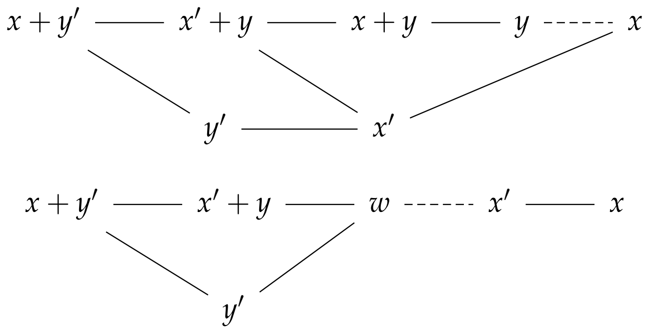

Now note that the cases $(x = y \wedge x+y = x')$ and $x < y$ lead to $x' = 2x$: The first is trivial, for the second note that $x < y$ implies $y = x'$ and $w = x+y$, and hence $2x+y = x+w = x'+y$. But then $y' = x'+w-x = w+x$, a contradiction. So we must have $x = y$ and $x+y = y' = 2x$, which implies $w = x'$. We get the solution $y = v = x$, $x' = v' = w = 3x/2$, and $y' = w' = 2x$, where x can be chosen freely among positive even integers smaller than $(n-1)/2$.

***$X \prec V \prec Y \prec W$*.** We start by observing some general relations. Since $x \leq v \leq y \leq w$, we must have $x = v \leq y \leq w$ as the smallest elements on each side. This implies $x' < v'$ and since $x'+y' = v'+w'$ we thus have $y' > w'$. But then we must have $w > y$, so in particular $w > x$. Then $v'+w$ is strictly larger than $x'+y$ and hence $v'+w \in \{y', x+y'\}$. Since $w > x$, this implies that $y' > v'$, and hence $w' > x'$. But then $x+w' < x+y'$, and hence we must have $v'+w = x+y'$, since otherwise $x+y'$ could not be matched.

This already determines the ordering of the duals. First, note that

$$x'-x = x'-(v'+w-y') = (v'+w'-y')-(v'+w-y') = w'-w,$$

and since $w' > x'$, this implies $\overline{X} \prec \overline{W}$. Furthermore

$$w'-w = w'-(x+y'-v') = (w'-y')+(v'-v),$$

and since $w' < y'$, this implies $v'-v > w'-w$ and hence $\overline{W} \prec \overline{V}$. Finally,

$$v'-v = y'+x-w-v = y'-w < y'-y,$$

and so $\overline{V} \prec \overline{Y}$. So by identifying X and Y with their own respective duals, $\overline{W}$ with V, $\overline{V}$ with W, this corresponds to a case of the form $X \prec V \prec W \prec Y$, and so it suffices to check these.

$X \prec V \prec W \prec Y$. We will again first state some universally true things and then make case distinctions. Since $x \le v \le w \le y$ we must actually have $x = v \le w \le y$ since x and v are the smallest element in their respective sumset. This implies $x' < v'$ and hence $y' > w'$. So $x + y' > x + w'$ and hence we know that $x + y' \in \{v', v' + w\}$.

But suppose that we have $x + y' = v'$. Then $v' > y'$ and hence we must have $w > x$ and $x' > w'$. At the same time, there has to be an element strictly between $x' + y'$ and $x + y'$ in $X + Y$, and so in particular we must have $x' + y > y' + x$ and $x' + y = v' + w$, which implies $y > w$. But then $x < w < \min(x', y)$ and hence it cannot be matched by any element in $X + Y$.

So we must have $x + y' = v' + w$, and so $y' \ge v'$ and hence $w' \ge x'$, in particular $y' > x'$. Furthermore, $v' + w > x + w'$ and $x + y' \ge x' + y$ since we know that $x + y'$ is the second largest element in $V + W$. Finally, we see that $y' \in \{v', x + w, x + w'\}$. But in fact, we cannot have $y' = v'$, since this would imply $x' = w'$ and hence $w > x$. But at the same time, we have $y' + x = v' + w = y' + w$ which implies $x = w$. So $y' \neq v'$ and hence in fact $y' > v'$, which also implies $w' > x'$ and $w > x$ because of $y' + x = v' + w$. In particular, $y > x$. We have to consider two different cases for the assignment of y'.

Case a) $y' = x + w \neq v'$. This tells us in particular that $y' \le x + y < x' + y$ and since $v' + w = y' + x = 2x + w$ we have $v' = 2x$. We get the following diagrams.

$$\begin{tikzcd}
x+y' \arrow[r, dashed, dash] & x'+y \arrow[r, dash] & x+y \arrow[r, dashed, dash] & y' \arrow[r, dash] \arrow[dr, dash] & x' \arrow[r, dash] & x \arrow[dl, dash] \\
 & & & & y &
\end{tikzcd}$$

$$\begin{tikzcd}
x+y' \arrow[r, dash] & x+w' \arrow[r, dash] & y' \arrow[r, dash] \arrow[dr, dash] & v' \arrow[rr, dash] & & x \arrow[dl, dash] \\
 & & & w' \arrow[r, dash] & w &
\end{tikzcd}$$

From the bottom diagram, we infer that there is only one element strictly between $x + y'$ and y', so exactly one of the inequalities in the top diagram has to be an equality. Furthermore, since w' is strictly between y' and x', we must have $w' = y$ and in particular $y' > y > x'$. This also implies $w = x'$ and $x + w' = x + y$, so we must have $x + y > y'$ and $x + y' = x' + y$. Finally, $x' + y = x + y' = v' + w = v' + x'$ implies $v' = y$. This leads to the solution $v = x$, $x' = w = 3x/2$, $y = v' = w' = 2x$, and $y' = 5x/2$ for any positive even integer x smaller than $2(n-1)/5$.

Case b) $y' = x + w'$. This tells us in particular that $x + w' > v'$. Furthermore, we have $x' + x + w' = x' + y' = v' + w'$ and hence $v' = x' + x$. This implies $x' + x + w = v' + w = x + y' = 2x + w'$, and hence $x' + y \ge x' + w = x + w' = y'$. Since y' is the third largest element in $V + W$, this means that $x' + y \in \{x' + y, y'\}$.

Case b)i. $x' + y = y'$. So $x' + y = x' + w$ and hence $w = y$. Since w' is strictly larger than $\max\{x', y\}$ but smaller than y', we must have $w' = x + y$. Now we must have $y = x'$, since otherwise x' could not be matched, and hence $v' = w'$. This leads to the solution $v = x$, $x' = y = w = 2x$, $v' = w' = 3x$, and $y' = 4x$ for a positive integer x smaller than $(n-1)/4$.

Case b)ii. $x' + y = x + y'$. We see that $x' + y = v' + w$ and hence $y > w$. Furthermore, we have that $y' < x' + y$, and since y' is the third largest element in $V + W$, we must have $y' \ge x + y$.

Case b)ii.α $y' = x + y$. Since v' and w' are strictly between y' and x', we must have $v' = w' = y$ and in particular $y > x'$, which also implies $w = x'$. Hence $x + w$ is strictly between x' and y', that is $x + w = y$. This leads to the solution $v = x$, $x' = w = 2x$, $y = v' = w' = 3x$, and $y' = 4x$, for any positive integer x smaller than $(n-1)/4$.

Case b)ii.β $y' = x + y$. We get following diagrams.

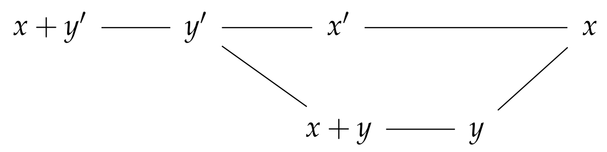

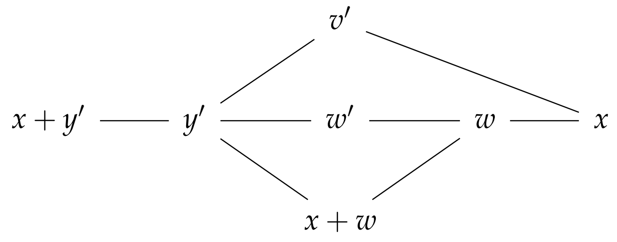

Since $w', v' > x'$, we must have that $x+y > x'$ and in particular $x+y = \max\{v', w'\}$. Furthermore, we also must have $v' > w$, since otherwise v' would have to be matched to $\min(y, x') \leq x'$. Since $w < y$, we in fact need to have $x' = w$ and in particular $x' < y$. Thus the top diagram is a chain of 8 elements, and hence there has to be exactly one equality on the bottom side. Since we now know that $x+y > x+w$, we must have $x+w = y$. Furthermore, $v' + x' = v' + w = x' + y$ and hence $v' = y$. So since we have to match $x+y$, we must have $w' = x+y > v'$. This leads to the solution $v = x$, $x' = w = 3x$, $y = v' = 4x$, $w' = 5x$, and $y' = 6x$, for any positive integer x smaller than $(n-1)/6$.

So there are only a constant number of different cases, each of which only has $O(n)$ different solutions. ■

Let us make two remarks concerning the $k = 3$ result. First, after also computing the relevant duals, the non-trivial sumset equations are exactly dilations of the sets in the following list:

$$\begin{aligned}
\{0,1,2\} + \{0,2,5\} &= \{0,1,2\} + \{0,3,5\} = \{0,1,2,3,4,5,6,7\}\\
\{0,1,3\} + \{0,1,5\} &= \{0,1,4\} + \{0,2,4\} = \{0,1,2,3,4,5,6,8\}\\
\{0,1,3\} + \{0,4,6\} &= \{0,1,4\} + \{0,3,5\} = \{0,1,3,4,5,6,7,9\}\\
\{0,2,3\} + \{0,4,5\} &= \{0,2,4\} + \{0,3,4\} = \{0,2,3,4,5,6,7,8\}\\
\{0,2,3\} + \{0,2,6\} &= \{0,2,5\} + \{0,3,4\} = \{0,2,3,4,5,6,8,9\}\\
\{0,1,2\} + \{0,1,4\} &= \{0,1,2\} + \{0,2,4\} = \{0,1,2,3,4,5,6\}\\
\{0,1,2\} + \{0,3,4\} &= \{0,1,3\} + \{0,2,3\} = \{0,1,2,3,4,5,6\}\\
\{0,1,3\} + \{0,1,4\} &= \{0,1,3\} + \{0,2,4\} = \{0,1,2,3,4,5,7\}\\
\{0,2,3\} + \{0,2,4\} &= \{0,2,3\} + \{0,3,4\} = \{0,2,3,4,5,6,7\}\\
\{0,1,2\} + \{0,1,3\} &= \{0,1,2\} + \{0,2,3\} = \{0,1,2,3,4,5\}.
\end{aligned}$$

In particular, since it thus suffices to remove all dilations of $\{0,1,2\}$ and $\{0,1,3\}$ from $\binom{[n-1]}{k}_0$, we get the precise lower bound

$$F_3(n) \geq \binom{n-1}{2} - \frac{5}{6}n.$$

Furthermore, when considering the equations above, note that there are no cases where the sumset has size 5 or 9. For the former, this comes from the fact that if $|A| = |B| = 3$, then $5 = |A| + |B| - 1$, and hence Proposition 3.2 tells us that A and B must be arithmetic progressions with the same common difference. In the case of $A, B \subset \binom{[n-1]}{3}_0$, they thus must both be dilations of $\{0,1,2\}$ by a common factor.

Similarly, $|A+B| = 9$ represents the $k = 3$ case of $|A+B| = |A||B|$, and we can present a short proof that the sumset uniquely determines the summand sets here. Suppose

$$A = \{0 < a_1 < a_2\} \quad \text{and} \quad B = \{0 < b_1 < b_2\},$$

and assume without loss of generality that $a_1 < b_1$.[*] If all sums in $A+B$ are distinct, then the equations

$$\begin{aligned} a_1 &= \min\left((A+B) \setminus \{0\}\right) \\ a_2 + b_2 &= \max(A+B) \\ 3(a_1+b_1+a_2+b_2) &= \sum_{z \in A+B} z \end{aligned}$$

determine a_1, b_1, and $a_2 + b_2$ from $A+B$. Moreover, $\{a_1+b_2, a_2+b_1\}$ are the second and third largest elements, say $s' < s$, in $A+B$. If $a_1 + b_2 = s$ then $b_2 = s - a_1$ and $a_2 = \max(A+B) - b_2$ are the two points distinct from $\{0, a_1, b_1, a_1+b_1, s, s', a_2+b_2\}$ in $A+B$, otherwise these points are $a_2 = s - b_1$ and $b_2 = \max(A+B) - a_2$, and only one of these two possibilities can occur. Therefore, if $|A+B| = 9$ then the sumset $A+B$ determines the sets A, B. As we will see later in Section 4.3, the analogous statement for $|A+B| = k^2$ will not be true anymore for $k \geq 4$.

4.2 Random Sidon systems

This section will cover two results on Sidon systems in the binomial random set $\binom{[n]}{k}_p$ in which every k-subset of $[n]$ is included independently with probability p. The first will answer the question of when we expect the binomial random set to be a Sidon system. Before stating our result, let us mention the equivalent theorem for Sidon sets, proved by Godbole, Janson, Locantore and Rapoport in [45].

Theorem 4.9 ([45]). *Let n be an integer and $p = p(n) \in [0,1]$. Then*

$$\lim_{n\to\infty} \mathrm{P}([n]_p \text{ is a Sidon set}) = \begin{cases} 1, & \text{if } p = o(n^{-3/4}) \\ 0, & \text{if } p = \omega(n^{-3/4}) \end{cases}.$$

We say that $n^{-3/4}$ is the *threshold probability* for $[n]_p$ to be a Sidon set. They also investigated the behavior at the threshold, that is, what happens if $p = Cn^{-3/4}$, and in fact for both questions regarded the more general case of the h-fold generalization of Sidon sets, called B_h sets. Our generalization to the setting of Sidon systems of k-subsets of the first n integers will be consistent with Theorem 4.9 when interpreted as the $k = 1$ case, although we will see later that this is actually more of a coincidence.

Theorem 4.10. *Let $k \geq 2$ be a fixed integer. Then for an integer n and $p = p(n) \in [0,1]$, it holds that*

$$\lim_{n\to\infty} \mathrm{P}\left(\binom{[n]}{k}_p \text{ is Sidon}\right) = \begin{cases} 1, & \text{if } p = o(n^{-(2k+1)/4}) \\ 0, & \text{if } p = \omega(n^{-(2k+1)/4}) \end{cases}.$$

Another interesting question is to study the *sparse random* analogue of determining bounds on $F_k(n)$. That is, instead of investigating the size of the largest Sidon system in $\binom{[n]}{k}$, what happens if we do this in $\binom{[n]}{k}_p$? The Sidon set equivalent of this was answered by Kohayakawa,

[*] Note that we cannot have equality here, since otherwise $a_1 + 0 = 0 + b_1$ are two distinct representations of the same element in $A+B$.

Lee, Rödl and Samotij in [65] and they determined an interesting phase transition. Essentially, as long as $p = o(n^{-1/3})$, the expected number of Sidon quadruples is negligible when compared to the expected size of the random set, and hence standard concentration bounds tell us that the size of the largest Sidon subset will be the same as the size of the random set. For p in the range between $n^{-1/3}$ and constant, the situation is similar to that in $[n]$, that is, the size of the largest Sidon subset is approximately the square root of np, the size of the random set. This range can be seen as an example of the *transference principle* that says that results in the dense setting can be moved to the sparse random one in certain cases. As already discussed in Chapter 1, this principle was studied in a very general way by Conlon and Gowers [26] and Schacht [93]. The transference principle also played a fundamental role in the proof of the existence of arbitrarily long arithmetic progressions in the primes due to Green and Tao [53]. Returning to the problem of the largest Sidon subset in the binomial random set $[n]_p$, we see that since our problem is clearly monotone in nature, the situation when $n^{-2/3} \leq p \leq n^{-1/3}$ is that the largest Sidon subset must stay constant in the exponent at approximately $n^{1/3}$. Let us summarize.

Theorem 4.11 ([65])**.** *Let $0 \leq a \leq 1$ be a fixed constant. Suppose $p = p(n) = (1 + o(1))n^{-a}$. There exists a constant $b = b(a)$ such that almost surely the largest Sidon subset of $[n]_p$ has size $n^{b+o(1)}$. Furthermore,*

$$b(a) = \begin{cases} 1 - a, & \text{if } 2/3 \leq a \leq 1, \\ 1/3, & \text{if } 1/3 \leq a \leq 2/3, \\ (1-a)/2, & \text{if } 0 \leq a \leq 1/3. \end{cases}$$

Our second result on Sidon systems in this chapter will imply a somewhat less nuanced version of Theorem 4.11. It will be helpful to change the language from the absence to the appearance of additive structures.

Definition 4.12. Let G be an abelian group and suppose $A, B, C, D \subset G$ are subsets. We say that (A, B, C, D) forms an *additive quadruple* if $A + B = C + D$, and furthermore, it is called *nontrivial* if $\{A, B\} \neq \{C, D\}$.

Hence, a Sidon system is a family that does not contain any nontrivial additive quadruples. We can now define a relative version of this concept.

Definition 4.13. Let G be an abelian group and $\delta > 0$. Then a finite family of subsets $\mathcal{F} \subset 2^G$ is called *δ-additive* if every subfamily $\mathcal{G} \subseteq \mathcal{F}$ with $|\mathcal{G}| \geq \delta|\mathcal{F}|$ contains a nontrivial additive quadruple.

We are now ready to state the second result of this chapter, which determines the threshold probability for when $\binom{[n]}{k}_p$ is δ-additive. Unfortunately, we are only able to prove this conditional on Conjecture 4.8 being true.

Theorem 4.14. *Let $k \geq 2$ be a fixed integer and $1 > \delta > 0$ and suppose Conjecture 4.8 holds for k. Then there exist constants C, c that only depend on k, δ such that*

$$\lim_{n\to\infty} \mathrm{P}\left(\binom{[n]}{k}_p \text{ is } \delta\text{-additive} \right) = \begin{cases} 1, & \text{if } p \geq c/n \\ 0, & \text{if } p \leq C/n \end{cases}.$$

Moreover, if $p = o(1/n)$, then asymptotically almost surely, $\binom{[n]}{k}_p$ is not δ-additive even for $\delta = 1$.

As will become clear in the proof, only the 0-statement of Theorem 4.14 is conditional on the conjecture. Recalling that $F_k(n) \leq O_k(n^{k-1})$ by Theorem 4.6, this immediately gives us the following analogue of Theorem 4.11.

Corollary 4.15. *Let $k \geq 2$ be a fixed integer such that Conjecture 4.8 holds. Then there exist constants C, c that only depend on k such that asymptotically almost surely, the largest Sidon system $\mathcal{F} \subset \binom{[n]}{k}_p$ has size*

$$|\mathcal{F}| = \begin{cases} \Theta(n^{k-1}), & \text{if } p \geq C/n \\ \Theta(n^k p), & \text{if } p \leq c/n \end{cases}.$$

Moreover, if $p = o(1/n)$ this can be strengthened to $|\mathcal{F}| \sim \left|\binom{[n]}{k}_p\right|$.

In other words, we are essentially always in the regime that one can remove a negligible number of k-subsets in order to transform the random family into a Sidon system comparable to the $p = o(n^{-2/3})$ case for Sidon sets.

4.2.1 The proofs of Theorems 4.10 and 4.14

We start by introducing some notation. If A is a k-subset of the first n positive integers, we can always write it in the form $\min(A) + A'$ with $A' \in \binom{[n-1]}{k}_0$. If not stated explicitly otherwise, we will use the respective lower case letter for a set's minimum element, for instance $a = \min(A)$, and a dashed letter for the translation $A' = A - a$. The set A' will be referred to as the *distance set* of A.

Let us now make an easy but helpful observation that was already used implicitly in the proof of Theorem 4.6 and the general construction in the proof of Theorem 4.7. If A, B, U, V are k-subsets of $[n]$, then the equality

$$A + B = U + V$$

occurs if and only if

$$a + b = u + v \quad \text{and} \quad A' + B' = U' + V',$$

so we can often restrict ourselves to elements of $\binom{[n-1]}{k}_0$ and consider translations separately.

The 1-statement of Theorem 4.10

Let n, k, p be defined as in the theorem statement, and abbreviate $\mathcal{A} = \binom{[n]}{k}_p$. We denote by $\mathcal{B}$ the family

$$\mathcal{B} = \left\{ A = (A_1, A_2) : A_1, A_2 \in \binom{[n]}{k},\ A_1 \preceq A_2 \right\},$$

of ordered pairs of k-subsets of $[n]$, and by $\mathcal{C}$ the family

$$\mathcal{C} = \{(A, B) \in \mathcal{B} \times \mathcal{B} : A \prec B,\ A_1 + A_2 = B_1 + B_2\}$$

of nontrivial ordered additive quadruples. For $(A, B) \in \mathcal{C}$, let $I_{A,B}$ denote the indicator variable that $\{A_1, A_2, B_1, B_2\}$ belongs to the binomial random system $\mathcal{A}$, and define

$$X = \sum_{(A,B) \in \mathcal{C}} I_{A,B}.$$

Therefore,

$$\mathrm{P}(\mathcal{A} \text{ is a Sidon system}) = \mathrm{P}(X = 0).$$

Finally, for $2 \leq \ell \leq 4$, define

$$\mathcal{C}(\ell) = \{(A, B) \in \mathcal{C} : |\{A_1, A_2, B_1, B_2\}| = \ell\}.$$

Every k-subset of $[n]$ is contained in $\mathcal{A}$ independently with probability p, hence for each $(A, B) \in \mathcal{C}(\ell)$, we have that

$$\mathbb{E}(I_{A,B}) = p^{\ell}.$$

In order to prove the 1-statement of Theorem 4.10, we begin by giving upper bounds for the cardinalities of the $\mathcal{C}(\ell)$.

Lemma 4.16. *For all* $2 \le \ell \le 4$, *we have that*

$$|\mathcal{C}(\ell)| = O_k(n^{\ell(2k+1)/4}).$$

Proof. We will use two slightly different approaches. The first one uses the equivalence described before, namely that two sumsets are equal if and only if the sums of the minimal elements and the sumsets of the distance sets are equal.

Suppose we have four pairwise distinct k-sets $A, B, C, D \subset [n]$ such that $A + B = C + D$, which by the previously mentioned equivalence means that we also have $a + b = c + d$ and $A' + B' = C' + D'$[†]. Now, since 0 will be contained in each of the sets A', B', C' and D', we see that all four of them will be contained in $A' + B' = C' + D'$. Hence, after fixing A' and B', the potential elements of C' and D' must be chosen from the at most k^2 elements in $A' + B'$, and hence there are at most $O_k(n^{2k-2})$ choices for the quadruple (A', B', C', D'). Since we have the relation $a + b = c + d$, at most three of these elements can be chosen freely, which results in

$$|\mathcal{C}(4)| = O_k(n^{2k+1}).$$

A similar argument works for the case $\ell = 2$, by noting that this implies a sumset equality of the form $A + A = A + B$ or $A + A = B + B$, with $A \neq B$, and so the elements of B' have to be chosen from $A' + A'$. So there are $O_k(n^{k-1})$ choices for the pair (A', B'), and since the minimal elements satisfy the equality $2a = a + b$ or $2a = 2b$, and hence $a = b$, only one of them can be chosen freely, which results in the upper bound

$$|\mathcal{C}(2)| = O_k(n^k).$$

For $\ell = 3$, there are two possible types of sumset equality, namely

$$A + A = B + C \quad \text{or} \quad A + B = A + C.$$

The first case can be handled the same way as before, noting that one can choose at most two of the three minimal elements a, b, c freely, and so there are at most $O_k(n^{k+1})$ such bad solutions. The second type requires us to make a slightly different argument. Since $B \neq C$, we see that $B \setminus C \neq \emptyset$. Let us for now assume that there is a unique $b \in B \setminus C$, and let $A = \{a_1, \ldots, a_k\}$ and $C = \{c_1, \ldots, c_k\}$. Since $A + B = A + C$, we have that $A + b \subseteq A + C$, and hence there exist functions $\pi, \tau \colon [k] \to [k]$ such that for all $i \in [k]$,

$$a_i + b = a_{\pi(i)} + c_{\tau(i)}.$$

Furthermore, since $b \notin C$, we see that $a_i \neq a_{\pi(i)}$ for all $i \in [k]$. Write this linear system of equations in matrix form as

$$M \cdot (a_1, \ldots, a_k, b, c_1, \ldots, c_k) = 0, \tag{4.9}$$

then the i-th row of M will have 1s in the i-th and $(k+1)$-st column, as well as -1s in the $\pi(i)$-th and $(k + 1 + \tau(i))$-th one, and 0s everywhere else. We will show that M has rank at

[†] Recall that for a set A, a denotes its minimal element, and A' the set $A - a$.

least $\lceil k/2 \rceil + 1$. Let us ignore the last k columns and only focus on those corresponding to A and b. We start with row 1, which is clearly nonzero. Call any row that has its 1 entry in a column in which a previously picked row has a -1 entry *closed*. For example, at the start, only row $\pi(1)$ is closed. Proceed by successively picking a row among the non-closed ones. Since every row contains only a single -1 entry, it is not hard to see that at the end of this process, we have at least $\lceil k/2 \rceil$ linearly independent rows. Finally, because the $(k+1)$-st column in every row vector is 1, we can pick a single arbitrary closed row and end up with $\lceil k/2 \rceil + 1$ linearly independent rows, which implies the lower bound on the rank of M. Now note that if $|B \setminus C| > 1$, we can do the same, but add more columns for the remaining elements in this set. Since we get at least one linearly independent row per new column as well, the general lower bound will hence be

$$\operatorname{rk}(M) \geq \lceil k/2 \rceil + |B \setminus C|.$$

Hence, at most

$$2k + |B \setminus C| - \left(\left\lceil \frac{k}{2} \right\rceil + |B \setminus C| \right) = k + \left\lfloor \frac{k}{2} \right\rfloor \leq 3k/2$$

of the elements in A, C and $B \setminus C$ can be chosen freely to satisfy the equation (4.9). Since elements from $B \cap C$ have to be in C, the same is true (up to maybe some factor only depending on k) for A, B, C. Since there are only $O_k(1)$ possible matrices M which may lead to an equation of the form $A + B = A + C$, we conclude that

$$|\mathcal{C}(3)| = O_k(n^{3k/2}).$$

This completes the proof. ∎

We can now use this to give a short proof of the 1-statement. By linearity of expectation and Markov's inequality, we see that

$$\mathrm{P}(X \geq 1) \leq \mathbb{E}(X) = \sum_{(A,B) \in \mathcal{C}} \mathbb{E}(I_{A,B}) = \sum_{\ell=2}^{4} \sum_{(A,B) \in \mathcal{C}(\ell)} \mathbb{E}(I_{A,B}) = \sum_{\ell=2}^{4} |\mathcal{C}(\ell)| p^{\ell}. \tag{4.10}$$

If $p = o(n^{-(2k+1)/4})$, then since $|\mathcal{C}(\ell)| = O_k(N^{\ell(2k+1)/4})$ by Lemma 4.16 it follows that

$$\mathrm{P}(X \geq 1) = o(1),$$

proving the 1-statement of Thereom 4.10. ∎

The 0-statement of Theorem 4.10

Keeping all the definitions from the previous subsection, define

$$\mathcal{C}' = \{(A, B) = ((A_1, A_2), (B_1, B_2)) \in \mathcal{C}(4) : A_1' \neq A_2',\, A_1' = B_1' \text{ and } A_2' = B_2'\}.$$

In particular, for any $(A, B) \in \mathcal{C}'$, the minimal elements a_i, b_i satisfy $a_i \neq b_i$ for $i = 1, 2$. Let

$$Y = \sum_{(A,B) \in \mathcal{C}'} I_{A,B},$$

then clearly

$$\mathrm{P}(X = 0) \leq \mathrm{P}(Y = 0),$$

and so it suffices to show the 0-statement for $\mathcal{C}'$. Let $\mathcal{D}$ denote the family

$$\mathcal{D} = \left\{ ((A, B), (C, D)) \in \mathcal{C}' \times \mathcal{C}' : \begin{array}{c} (A, B) \neq (C, D) \text{ and} \\ \{A_1, A_2, B_1, B_2\} \cap \{C_1, C_2, D_1, D_2\} \neq \emptyset. \end{array} \right\}.$$

That is, it contains the pairs of distinct elements in $\mathcal{C}'$ which share at least one k-set. For $4 \leq \ell \leq 7$ we define the families

$$\mathcal{D}(\ell) = \{((A,B),(C,D)) \in \mathcal{D} : |\{A_1, A_2, B_1, B_2, C_1, C_2, D_1, D_2\}| = \ell\}.$$

We first give a lower bound for $|\mathcal{C}'|$.

Lemma 4.17. $|\mathcal{C}'| = \Omega_k(n^{2k+1})$.

Proof. We can choose $\Omega_k(n^{2(k-1)})$ different pairs $A_1' \neq A_2'$ of sets in $\binom{[n]}{k}_0$ and, for every such pair, we can choose $\Omega(n^3)$ elements $a_1 < b_1 < b_2 < a_2 \in [n]$ with $a_1 + a_2 = b_1 + b_2$ and $(a_1 + A_1'), (a_2 + A_2'), (b_1 + A_1'), (b_2 + A_2') \in \binom{[n]}{k}$ which form elements in $\mathcal{C}'$, and so the statement follows. ■

Next we give upper bounds for $|\mathcal{D}(\ell)|$ along the same lines as for the upper bounds of $|\mathcal{C}(\ell)|$ in Lemma 4.16.

Lemma 4.18. *For $4 \leq \ell \leq 7$, we have that*

$$|\mathcal{D}(\ell)| = \begin{cases} 0, & \text{if } \ell = 4,5 \\ O_k(n^{3k+1}), & \text{if } \ell = 6 \\ O_k(n^{3k+2}), & \text{if } \ell = 7 \end{cases}$$

Proof. We first prove that $\mathcal{D}(4)$ and $\mathcal{D}(5)$ have to be empty. Indeed, for $\ell = 5$, three of the sets in (C,D) will be fixed by (A,B), and so, by the definition of $\mathcal{C}'$, the last one will be as well.

For $\ell = 4$, in order to get a nontrivial solution, (C,D) must define an equation of the form $A_1 + B_2 = B_1 + A_2$ and hence $a_1 + b_2 = b_1 + a_2$, which together with $a_1 + a_2 = b_1 + b_2$ implies $a_1 = b_1$ and $a_2 = b_2$, a contradiction to the definition of $\mathcal{C}'$.

Suppose $\ell = 7$. For each pair (A,B) there is a fixed set in the pair (C,D), say e.g. C_1. Since $C_1', C_2', D_1', D_2' \in \binom{[n]}{k}_0$, there are $O_k(n^{k-1})$ choices for C_2' and only $O_k(1)$ choices of D_1' and D_2' afterwards, since they have to be chosen from the elements of $C_1' + C_2'$. Since c_1 is fixed and we have $c_1 + c_2 = d_1 + d_2$, there are $O(n^2)$ choices for c_2, d_1, d_2. Summarizing,

$$|\mathcal{D}(7)| = |\mathcal{C}'| O_k(n^{k+1}) = O_k(n^{3k+2}).$$

For $\ell = 6$, two of the sets in the pair (C,D) are fixed by the pair (A,B). By repeating the above reasoning, we have at most $O(n)$ choices for c_1, c_2, d_1, d_2 giving

$$|\mathcal{D}(6)| = |\mathcal{C}'| O_k(n^k) = O_k(n^{3k+1}).$$

This completes the proof. ■

We are now ready to prove the 0-statement. By the Janson inequality (see e.g. Theorem 1.1 in Chapter 8 of [3]) we obtain

$$\mathrm{P}(Y = 0) \leq \prod_{(A,B) \in \mathcal{C}'} \mathrm{P}(I_{A,B} = 0) \cdot \exp(\Delta) = (1 - p^4)^{|\mathcal{C}'|} \exp(\Delta), \tag{4.11}$$

where

$$\begin{aligned} \Delta &= \sum_{((A,B),(C,D)) \in \mathcal{D}} \mathrm{P}(I_{A,B} I_{C,D} = 1) \\ &= \sum_{\ell=5}^{7} \sum_{((A,B),(C,D)) \in \mathcal{D}(\ell)} \mathrm{P}(I_{A,B} I_{C,D} = 1) \qquad (4.12) \\ &= \sum_{\ell=4}^{7} |\mathcal{D}(\ell)| p^{\ell}. \end{aligned}$$

By inserting the upper bounds from Lemma 4.18 into (4.12) we get

$$\Delta = \sum_{\ell=4}^{7} |\mathcal{D}(\ell)| p^{\ell} = O_k(n^{3k+2}p^7 + n^{3k+1}p^6).$$

Suppose for now that $p = o(n^{-(3k+1)/6})$. Then it is straightforward to check that since $k \geq 2$, it holds that

$$n^{3k+1}p^6 + n^{3k+2}p^7 = o(1),$$

which implies $\Delta = o(1)$. If it also holds that $p = \omega(n^{-(2k+1)/4})$, we can now use the lower bounds on $|\mathcal{C}'|$ obtained in Lemma 4.17, such that Equation (4.11) therefore gives

$$\begin{aligned} \mathrm{P}(X = 0) &\leq \mathrm{P}(Y = 0) \\ &\leq (1 - p^4)^{|\mathcal{C}'|} \exp(\Delta) \\ &\leq \exp(-\Omega_k(p^4 n^{(2k+1)}) + o(1)) \\ &= \exp(-\omega(1)). \end{aligned}$$

The proof of the 0-statement can now be completed by noting that the property of being a Sidon system is clearly monotone, and hence if it holds for $\log_n(p)$ between $-(2k+1)/4$ and $-(3k+1)/6$, it will in fact hold for all $p = \omega(n^{-(2k+1)/4})$. ■

The 1-statement of Theorem 4.14

The proof of the 1-statement is very simple, it essentially only uses the upper bound proved in Theorem 4.6 and standard concentration bounds. Specifically, we will make use of the following form of Chernoff's bound (see for example Theorem 1.8 and Corollary 1.9 in [105]).

Lemma 4.19 (Chernoff's inequality). *Assume that $X_1, \ldots, X_m$ are jointly independent random variables where $|X_i - \mathbb{E}(X_i)| \leq 1$ for all i. Set $X = X_1 + \cdots + X_m$ and let $\sigma = \sqrt{\mathbb{V}\mathrm{ar}(X)}$ be the standard deviation of X. Then for any $\lambda > 0$*

$$\mathrm{P}(|X - \mathbb{E}(X)| \geq \lambda\sigma) \leq 2 \max\left(e^{-\lambda^2/4}, e^{-\lambda\sigma/2}\right). \tag{4.13}$$

In particular, if $X = t_1 + \cdots + t_m$ where the t_i are independent boolean random variables, then for any $\epsilon > 0$

$$\mathrm{P}(|X - \mathbb{E}(X)| \geq \epsilon\, \mathbb{E}(X)) \leq 2e^{-\min(\epsilon^2/4, \epsilon/2)\, \mathbb{E}(X)}. \tag{4.14}$$

Now, by Theorem 4.6 it holds that the maximum size $F_k(n)$ of a family of k-element integer subsets of $[n]$ without any nontrivial additive quadruples is at most $M = \binom{n-1}{k-1} + n - k$. Defining $N = \binom{n}{k}$, we see that if $p \geq 2M/(\delta N)$, then the expected size of $\mathcal{A} = \binom{[n]}{k}_p$ is $\mathbb{E}\,|\mathcal{A}| = 2M/\delta$. Clearly $|\mathcal{A}| = \sum \mathbb{I}(A \in \mathcal{A})$ where the sum goes over all sets A in $\binom{[n]}{k}$, hence we apply (4.14) with $\epsilon = 1/3$ and get that with probability at least $P = 1 - 2\exp(-M/18\delta)$, we have that $|\mathcal{A}| \geq 4M/3\delta$. Clearly P tends to 1 as n tends to infinity, so asymptotically almost surely every subsystem $\mathcal{F} \subseteq \mathcal{A}$ with $|\mathcal{F}| \geq \delta|\mathcal{A}|$ has cardinality at least $(4/3)M > M$, and hence it must contain at least one additive quadruple by Theorem 4.6. This proves the 1-statement. ■

The 0-statement of Theorem 4.14

Proving the 0-statement is slightly more involved. Since it will assume Conjecture 4.8 to be true, it will hold unconditionally only for $k = 2, 3$.

To start, by the 1-statement of Theorem 4.10, if $p = o(n^{-(2k+1)/4})$, then with probability $1 - o_k(1)$, the random family $\mathcal{A} = \binom{[n]}{k}_p$ itself contains no nontrivial additive quadruple, and hence the 0-statement in Theorem 4.14 holds for any $0 \le \delta \le 1$. So assume now that $p = \omega(n^{-(2k+1)/4})$ (in particular $p = \omega(1/N)$). Conjecture 4.8 suggests that a necessary and sufficient condition for $\mathcal{A}$ to contain a subsystem of relative density δ without nontrivial additive quadruples is that

$$X = \left|\left\{B \in \binom{[n-1]}{k}_0 : B + z \subset \mathcal{A} \text{ for some } z \in \mathbb{Z}.\right\}\right| \ge \delta|\mathcal{A}|. \tag{4.15}$$

In other words, we need that the sets in $\mathcal{A}$ are distributed somewhat uniformly among all equivalence classes. First note that $\mathbb{E}\,|\mathcal{A}| = Np$ and $\mathbb{V}\mathrm{ar}\,|\mathcal{A}| = N(p - p^2) \le \mathbb{E}\,|\mathcal{A}|$, and by (4.13) with $\lambda = \sigma^{1/2}$, we have that with probability $P = 1 - 2\max\left(e^{-\sigma/4}, e^{-\sigma^{3/4}/2}\right)$ it holds that $|\mathcal{A}| = Np \pm \sigma^{3/4}$. Since $p = \omega(1/N)$ we see that $P = 1 - o(1)$ and $\sigma^{3/4} = o(Np)$, hence with high probability $|\mathcal{A}| \sim Np$. We will now give estimates for $\mathbb{E}(X)$. By the Bonferroni inequalities we see that for any $B \in \binom{[n-1]}{k}_0$

$$\mathrm{P}(B + z \subset \mathcal{A} \text{ for some } z \in \mathbb{Z}) \le (n - \max(B))p \tag{4.16}$$

and

$$\mathrm{P}(B + z \subset \mathcal{A} \text{ for some } z \in \mathbb{Z}) \ge (n - \max(B))p - \binom{n - \max(B)}{2} p^2. \tag{4.17}$$

Note that for any nonnegative integers K, L we also have

$$\sum_{i=K}^{L} \binom{i}{K} = \binom{L+1}{K+1},$$
$$\sum_{i=K}^{L} i\binom{i-1}{K-1} = K\sum_{i=K}^{L} \binom{i}{K} = K\binom{L+1}{K+1},$$

as well as

$$\sum_{i=K+1}^{L} i^2 \binom{i-1}{K} \le \frac{1}{K!}\sum_{i=K+1}^{L} i^{K+2} \le \frac{1}{K!}\int_0^L x^{K+2}dx = \frac{L^{K+3}}{(K+3)K!}.$$

Using all of this we get the upper bound

$$\begin{aligned}
\mathbb{E}(X) &\le \sum_{B \in \binom{[n-1]}{k}_0} (n - \max(B))p \\
&= \sum_{i=k-1}^{n-1} \binom{i-1}{k-2}(n-i)p \\
&= np\sum_{i=k-2}^{n-2} \binom{i}{k-2} - (k-1)p\sum_{i=k-1}^{n-1} \binom{i}{k-1} \\
&= np\binom{n-1}{k-1} - (k-1)p\binom{n}{k} \\
&\sim \frac{pn^k}{k!},
\end{aligned}$$

as well as the lower bound

$$\begin{aligned}
\mathbb{E}(X) &\geq \sum_{B \in \binom{[n-1]}{k}_0} \left((n - \max(B))p - \binom{n - \max(B)}{2} p^2 \right) \\
&= \sum_{i=k-1}^{n-1} \binom{i-1}{k-2} \left((n-i)p - \frac{(n-i)(n-i-1)}{2} p^2 \right) \\
&\geq np\binom{n-1}{k-1} - (k-1)p\binom{n}{k} - \frac{n^2p^2}{2}\binom{n-1}{k-1} + (k-1)np^2\binom{n}{k} \\
&\quad + \frac{np^2}{2}\binom{n-1}{k-1} - \frac{(k-1)p^2}{2}\binom{n}{k} - \frac{n^{k+1}p^2}{2(k+1)(k-2)!} \\
&\sim \frac{pn^k}{k!} - \frac{p^2 n^{k+1}}{(k+1)!}
\end{aligned}$$

The last expression on the right-hand side is larger than $\delta Np \sim \delta pn^k/k!$ for every $p \leq (1-\delta)(k+1)/n$.

Finally, note that X is also the sum of $\binom{n-1}{k-1}$ independent indicator random variables, and hence $\mathbb{V}\mathrm{ar}(X) \leq \mathbb{E}(X) = \Theta(pn^k)$. So we can use the same values as in the case of concentration for $|\mathcal{A}|$, that is, $\lambda = \mathbb{V}\mathrm{ar}(X)^{1/4}$ and by (4.13), we have $X \sim \mathbb{E}(X)$ asymptotically almost surely. Conjecture 4.8 implies that there exists a subfamily $\mathcal{G} \subseteq \binom{[n-1]}{k}_0$ with $|\mathcal{G}| = o(n^{k-1})$ such that $\binom{[n-1]}{k}_0 \setminus \mathcal{G}$ contains no nontrivial additive quadruples. Let us call the sets in $\mathcal{G}$ *bad*, and say that a set B is represented in $\mathcal{A}$ if there is some $z \in \mathbb{Z}$ such that $B + z \subset \mathcal{A}$. We will show now that bad sets are not overrepresented in X. Clearly, the upper bound in (4.16) still holds, that is, for every bad set $B \in \mathcal{G}$, the probability that it is represented in $\mathcal{A}$ is at most $(n - \max(B))p \leq np$, and hence the expected number of represented bad sets is at most $|\mathcal{G}|np = o(pn^k)$. It is easy to verify that we can apply Chernoff's inequality again to see that we also have $o(pn^k)$ bad represented sets with high probability, and hence they do not contribute meaningfully in the random setting, either. Hence, defining $\mathcal{F}$ to be the subset of $\mathcal{A}$ that includes exactly one translation of every represented set that is not bad, we end up with a family of size at least $\delta|\mathcal{A}|$ that contains no nontrivial additive quadruples. ■

4.3 Open problems and partial results

4.3.1 Asymptotically sharp lower bounds for $F_k(N)$

Especially in light of Theorem 4.14's dependence on it, the most begging question left open is whether Conjecture 4.8 is true or not. Note that while the conjecture is nominally stronger than the statement

$$F_k(n) \sim \frac{n^{k-1}}{(k-1)!},$$

for $k \geq 3$ they are essentially identical, since translations cannot generate a significant number of new sets. If we consider the family $\binom{[n-1]}{k}_0 + \binom{[n-1]}{k}_0$, then a randomly chosen element S will asymptotically almost surely have cardinality k^2, so it is reasonable to assume that sumsets of this cardinality are the most important case to consider. It is also not hard to show that this is almost true. Suppose we have four sets $A, B, C, D \in \binom{[n-1]}{k}_0$ that form an additive quadruple and assume that one of the sets, say A, is a Sidon set. We are going to use the trivial fact that $a + b = a' + b'$ if and only if $a - a' = b' - b$. If A is a Sidon set, we know that $r_{A-A}(x) \leq 1$ for every $x \neq 0$, and hence the Cauchy-Schwarz inequality implies

$$\begin{aligned} |A|^2|B|^2 &= \left(\sum_{x \in A+B} r_{A+B}(x) \right)^2 \leq |A+B| \sum_{x \in A+B} r_{A+B}(x)^2 \\ &= |A+B| \sum_{y \in B-B} r_{B-B}(y) r_{A-A}(y) \\ &\leq |A+B|(|A||B| + |B|^2). \end{aligned} \tag{4.18}$$

Using $|A| = |B| = k$ and simplifying results in $|A+B| \geq k^2/2$. Now, any k-set that is not Sidon satisfies at least one nontrivial linear relation between its elements, and hence there can be only $O(n^{k-2})$ of them in $\binom{[n-1]}{k}_0$, so we can remove all of them without affecting the asymptotic density, and hence all remaining nontrivial additive quadruples will consist of four Sidon sets. Note that in some sense, this last statement is stronger than just assuming the sumset $A+B$ to have size k^2, since we can also make it for any other kind of internal linear relation of the summand sets. That is, instead of asking them to be Sidon sets, which is equivalent to saying that $|A+A|$ should be maximal, we can ask that the h-fold sumset $|hA|$ for any fixed h should be maximal. The number of sets in $\binom{[n-1]}{k}_0$ that contain a solution to some fixed linear equation involving h variables is at most $O_h(n^{k-2})$, and hence they can be removed. An example for why this is stronger than asking $A+B$ to be large is the following construction in the case $k = 4$. Let a, b, c, d be intergers such that

$$S = \{0, a\} + \{0, b\} + \{0, c\} + \{0, d\} \quad \text{and} \quad |S| = 16.$$

Then, in general, we have three different representations for S as a sumset of two 4-sets, namely by pairing $\{0, a\}$ with one of the remaining three 2-sets, and pairing the other two. Similar constructions can be done for any k that is composite. In this case, the sets will always contain a solution to the Schur equation $x + y = z$. Still, it might be interesting to see if these two aspects can be combined in some way since it would open up the following two-step approach to establishing Conjecture 4.8 if k is fixed.

i.) Find a set L of size $O_k(1)$ of linear equations with length $O_k(1)$ such that for any sets $A, B \subset [n]$ of size k not containing any solutions to the equations in L, their sumset $|A+B|$ must have size $|A||B| - \binom{|A \cap B|}{2}$.

ii.) Find a second set of linear equations R such that if A, B, C, D is a nontrivial additive quadruple of k-subsets of $[n]$ satisfying $|A+B| = k^2 - \binom{|A \cap B|}{2}$, then one of the sets must contain a solution to an equation in R.

Since L and R will both be small enough, we could then make $\binom{[n-1]}{k}_0$ free of nontrivial additive quadruples by removing all sets containing solutions to either L or R. This general idea actually already solves the $A = B$ case, that is, the study of

$$M_k := \max \left\{ |\mathcal{F}| : \mathcal{F} \subset \binom{[n]}{k}, \forall A, B \in \mathcal{F} : A + A = B + B \iff A = B \right\}.$$

The argument is as follows. Suppose $A = \{a_1 < \cdots < a_k\}$ and $B = \{b_1 < \cdots < b_k\}$ are k-element sets of real numbers such that both are Sidon sets. We will prove by induction on the index that $A = B$. First, we have $2a_1 = 2b_1$ as the unique smallest element in $A + A$, which implies $a_1 = b_1$ as our induction base. So suppose now that there is some $j > 1$ such that $a_i = b_i$ for all $1 \leq i < j$. Consider the element $a_1 + a_j$. By sumset equality, there exist indices $s, t \in [k]$ such that

$$a_1 + a_j = b_s + b_t.$$

We cannot have that $s, t \in [j-1]$, since the elements of A and B agree for those indices, and $a_1 + a_j = a_s + a_t$ with $s, t < j$ would be a violation of A being a Sidon set. Hence without loss of generality $t \geq j$. Since we also clearly have $s \geq 1$, this implies

$$a_1 + a_j = b_s + b_t \geq b_1 + b_j = a_1 + b_j,$$

and hence $a_j \geq b_j$. Equality can then be obtained by just repeating the same argument starting with $b_1 + b_j$.

Since the number of k-sets in $[n]$ that are not Sidon is $O(n^{k-1})$, this implies $M_k \sim \binom{n}{k}$. A different proof by Selfridge and Straus [94] works in the complex numbers, where the ordering argument that was used above is not valid. Their technique also only requires the sums of distinct elements to be unique (so in the language of Chapter 1, the set A is only required to be free of *proper* solutions to the Sidon equation), but they need to exclude the cases where k is a power of 2 and in fact provide constructions for counter-examples of this stronger statement for $k = 2^q$ for any $q \geq 2$.

On the other hand, asking two distinct k-sets A and B to be Sidon is not enough for them to have a uniquely represented sumset among pairs of Sidon sets, as can already be seen in the $k = 3$ case, since

$$\{0,1,3\} + \{0,4,6\} = \{0,1,4\} + \{0,3,5\}$$

and it is easy to check that all four of these sets are Sidon. Finally, we mention that when assuming all four summand sets $A, B, C, D \in \binom{[n]}{k}_0$ to not satisfy any internal relations between three variables, that is, $\bigcup_{i,j,\ell \in \{-1,1\}} iA + jA + \ell A$ is as large as it can be (and similarly for B, C, D), computations for $k = 4$ and $k = 5$ up to n in the 100s have not produced any non-trivial sumset equality $A + B = C + D$.

4.3.2 $B_h[g]$ systems

It is also possible to further generalize the definition of a Sidon system, in the same way that Sidon sets can be generalized to so called $B_h[g]$ sets. For a family $\mathcal{A}$ of integer subsets, a set of integers C, and an integer $h \geq 2$, let $r_{h\mathcal{A}}(C)$ denote the number of different multisets $\{A_1, A_2, \ldots, A_h\}$, $A_i \in \mathcal{A}$ such that $A_1 + A_2 + \cdots + A_h = C$. A $B_h[g]$ *system* is a family $\mathcal{A}$ of integer subsets such that $r_{h\mathcal{A}}(C) \leq g$ for all sets $C \subseteq \mathbb{Z}$. So a Sidon system is a $B_2[1]$ system. We can now try to prove the previous results in this more general setting. Define $F_{k,g,h}(n)$ as the largest cardinality of a $B_h[g]$ system $\mathcal{A} \subseteq \binom{[n]}{k}$. Let us prove Theorems 4.6 and 4.7 for $B_2[g]$ systems, write $F_{k,g,2}(n) = F_{k,g}(n)$. We start by proving an upper bound.

Theorem 4.20. *Let $n > k, g \geq 2$, then*

$$F_{k,g}(n) = O_k(\sqrt{g} n^{k-1/2}).$$

Proof. There are $O_k(n^{2k-1})$ sumsets of the form $A + B$, with $A, B \in \binom{[n]}{k}$. Indeed, since any fixed sumset $A + B$ with $A, B \in \binom{[n]}{k}$ is essentially a translation of a sumset of two sets in $\binom{[n]}{k}_0$, there are at most $O_k(n^{2k-1})$ of them. Now, if $\mathcal{A} \subseteq \binom{[n]}{k}$ is a $B_2[g]$ system, then

$$\binom{|\mathcal{A}|+1}{2} = \sum_{S \subseteq \mathbb{Z}} r_{2\mathcal{A}}(S) = O_k(g n^{2k-1}).$$

Simplifying this gives the upper bound. ∎

Remark. This can be made more precise for specific values of k. For instance, there are exactly $n(n-1)^2/2$ sumsets in the case $k = 2$, which gives a bound $F_{2,g}(n) \leq \sqrt{g} n^{3/2}$.

At first sight this could seem like a rather weak statement, since Theorems 4.6 and 4.7 give $F_{k,1}(n) = \Theta_k(n^{k-1})$. However, we will see that for any $g \geq 2$, $\sqrt{g}n^{k-1/2}$ is indeed the right order for $F_{k,g}(n)$. More specifically, we get the following lower bound.

Theorem 4.21. *Let $n > k, g \geq 2$, then*

$$F_{k,g}(n) = \Omega_k(\sqrt{g}n^{k-1/2}).$$

Proof. Let $A \subseteq \{1, 2, \ldots, n/2\}$ be a $B_2[\lfloor g/2 \rfloor]$ set such that $|A| = \Theta(\sqrt{gn})$, which is well known to exist (c.f. [80]) and let $\mathcal{I} \subseteq \binom{[n/2]}{k}_0$ be a Sidon system. We can use the general construction from the proof of Theorem 4.7 and see that $|\mathcal{I}| = \Omega_k(n^{k-1})$. We will show that the family

$$\mathcal{A} = \{a + I : a \in A, I \in \mathcal{I}\}$$

is a $B_2[g]$ system. Suppose

$$(a + I) + (b + J) = (c + L) + (d + M),$$

then we must have

$$a + b = c + d \quad \text{and} \quad I + J = L + M.$$

Since A is a $B_2[\lfloor g/2 \rfloor]$ set, $a + b$ has at most $\lfloor g/2 \rfloor$ representations, and since $\mathcal{I}$ is a Sidon system, the latter implies $\{I, J\} = \{L, M\}$. Hence there are at most g representations for this sumset. Since we clearly have $|\mathcal{A}| = |A||\mathcal{I}| = \Omega_k(\sqrt{g}n^{k-1/2})$, this completes the proof. ■

Next, we are going to consider $B_h[1]$ systems, that is, we are asking for unique representations, but consider the situation of h-fold sumsets for some fixed $h \geq 2$. We are able to prove the following generalization of the 0-statement of Theorem 4.10 in this setting.

Proposition 4.22. *Let $n > k, h \geq 2$ be integers, and $0 \leq p \leq 1$. Define*

$$p_0(n, k, h) = p_0 = n^{-\frac{hk+1}{h+2}}.$$

If $p = \omega(p_0)$, then $\binom{[n]}{k}_p$ is asymptotically almost surely not a $B_h[1]$ system.

Proof. The proof is an adaptation of the one in Theorem 4.10 for the 0–statement in the case $h = 2$. As in that proof it suffices to exhibit the occurrence of a particular class of h-tuples violating the Sidon property.

We recall that, with the notation for distance sets and minimal elements used there, the equation

$$A_1 + A_2 + \cdots + A_h = B_1 + B_2 + \cdots + B_h$$

violating the Sidon condition is equivalent to

$$A'_1 + \cdots + A'_h = B'_1 + \cdots + B'_h \text{ and } a_1 + \cdots + a_h = b_1 + \cdots + b_h.$$

Let Z be a random system of k–sets. Consider the set

$$F = \{A = (A_1, \ldots, A_h) : A_i \in Z, A_i \preceq A_{i+1}, i = 1, \ldots, h-1\},$$

of ordered h-tuples of sets in Z. Denote by G the family of ordered pairs of distinct h-tuples $(A, B) \in F \times F$ satisfying the following properties:

(i) $A_i = a_i + A'_i$ and $B_i = b_i + A'_i$, $i = 1, \ldots h$.

(ii) $a_i = b_i$ for all but two subscripts in $[h]$, the a_i's are pairwise distinct and the b_i's are pairwise distinct, and

(iii) $\sum_i a_i = \sum_i b_i$.

Thus G consists of pairs violating the Sidon sets with $h+2$ sets in total.

Let Y be the random variable counting the number of pairs in G. The Proposition will be proved if we show that

$$\lim_{N\to\infty} \mathrm{P}(Y=0) = 0.$$

There are $\Omega(n^{h(k-1)})$ choices for the $A'_1, \ldots, A'_h$ and $\Omega(n^{h+1})$ for distinct integers $a_1, a_2, \ldots, a_h$, b_1, b_2 satisfying

$$a_1 + a_2 + a_3 + \cdots + a_h = b_1 + b_2 + a_3 + \cdots + a_h.$$

It follows that

$$|G| = \Omega(n^{hk+1}).$$

By the Janson inequality,

$$\mathrm{P}(Y=0) \le \prod_{(A,B)\in G} \mathrm{P}(I_{A,B}=0) \cdot \exp(\Delta) = (1-p^{h+2})^{|G|} \exp(\Delta), \tag{4.19}$$

where, by denoting by K the subset consisting of distinct pairs $((A,B),(C,D)) \in G \times G$ such that $\{A_1, \ldots, A_h, B_1, \ldots, B_h\} \cap \{C_1, \ldots, C_h, D_1, \ldots, D_h\} \neq \emptyset$ and by $K(m)$ the pairs of K which have m subsets in total,

$$\begin{aligned}\Delta &= \sum_{((A,B),(C,D))\in K} \mathrm{P}(I_{A,B}I_{C,D}=1) \\ &= \sum_{m=h+2}^{2h+3} \sum_{((A,B),(C,D))\in K(m)} \mathrm{P}(I_{A,B}I_{C,D}=1) \\ &= \sum_{m=h+2}^{2h+3} |K(m)| p^m.\end{aligned} \tag{4.20}$$

We will next upper bound the cardinalities of the sets $K(m)$.

Claim. *We have*

$$|K(m)| = \begin{cases} O_{h,k}(n^{\lfloor (m-h)/2 \rfloor (k-1) - k + m - h + hk - 1}) & \text{if } h+4 \le m \le 2h+3, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. By the definition of G we cannot have distinct pairs $(A,B),(C,D) \in G$ having at most $h+3$ sets in total, which shows that $K(h+2) = K(h+3) = \emptyset$.

Suppose that $h+4 \le m \le 2h+3$. For a fixed $(A,B) \in G$, we first note that only $m-h-2$ of the sets defined by a potential (C,D) can still be chosen freely. So the same is true for the minimal elements c_i, d_i. Furthermore, we also have the equation $\sum c_i = \sum d_i$, which results in another non-redundant restriction, and hence there are at most $m-h-3$ choices for the minimal elements of the pair (C,D). We now consider the distance sets. By the definition of G, there are exactly h pairwise distinct distance sets $C'_1, \ldots, C'_h$. Since at most $m-h-2$ of the sets defined by (C,D) are not determined by (A,B), we thus have at most

$$\left\lfloor \frac{m-h-2}{2} \right\rfloor = \left\lfloor \frac{m-h}{2} \right\rfloor - 1$$

undetermined distance sets. There are at most $O(n^{(\lfloor (m-h)/2 \rfloor -1)(k-1)})$ of such sumsets, and so

$$\begin{aligned} K(m) &= |G| \cdot O(n^{(\lfloor (m-h)/2 \rfloor -1)(k-1)}) \cdot O(n^{m-h-3}) \\ &= O(n^{\lfloor (m-h)/2 \rfloor (k-1) - k + m - h + hk - 1}). \end{aligned}$$

■

It follows from the above claim that, if

$$p = o(n^{-\frac{\lfloor (m-h)/2 \rfloor (k-1) - k + m - h + hk - 1}{m}}),$$

then

$$|K(m)|p^m = o(1), \; h+2 \le m \le 2h+3. \tag{4.21}$$

We note that, for $k, h \geq 2$ and $h+4 \leq m \leq 2h+3$, we have

$$\frac{(m-h)(k-1)/2 - k + m - h + hk - 1}{m} < \frac{hk+1}{h+2}.$$

Hence, for any any p such that

$$p = o(n^{-\frac{\lfloor (m-h)/2 \rfloor (k-1) - k + m - h + hk - 1}{m}}) \quad \text{and} \quad p = \omega(n^{-\frac{hk+1}{h+2}}),$$

we see that (4.21) holds, and furthermore when looking at (4.19) we get

$$\begin{aligned} \mathrm{P}(Y=0) &\leq (1-p^{h+2})^{|G|} \exp(\Delta) \\ &\leq \exp(-|G|p^{h+2} + \Delta) \\ &= \exp(-\omega(1) + o(1)), \end{aligned}$$

and so the family is asymptotically almost surely not a $B_h(1)$ system for p in this range. But this property is clearly monotone in p, and hence we get the 0-statement for all $p = \omega(n^{-\frac{hk+1}{h+2}})$. ■

Recall that the threshold in Theorem 4.10 essentially resulted from the fact that the most important case was that of four pairwise distinct sets A, B, C, D such that

$$A + B = C + D.$$

The threshold for $B_h[1]$ sets obtained by Godbole et al. in [45] was similarly defined by the case of $2h$ distinct elements, and hence the initial assumption might be that this corresponds to the case of $2h$ sets for general $B_h[1]$ systems as well, which would lead to a threshold at

$$p_1 = n^{-\frac{h(k+1)-1}{2h}}.$$

Proposition 4.22 in fact shows that the case of $h+2$ pairwise distinct sets is more important, since one can easily check that $p_0 \leq p_1$, with equality only if $h = 2$. Further evidence that the p_0 in Proposition 4.22 might be the correct threshold is given by the fact that it is easy to check that the 1-statement holds in the special case of $k = 2$ for any $h \leq 4$. Note that inserting $k = 1$ into the threshold p_0 gives something strictly weaker than the result by Godbole et al. for all $h > 2$, and so the consistency in the $h = 2$ case seems to be more of a coincidence.

Bibliography

[1] Noga Alon, *Large sets in finite fields are sumsets*, J. Number Theory **126** (2007), no. 1, 110–118 (English).

[2] Noga Alon, József Balogh, Robert Morris, and Wojciech Samotij, *A refinement of the Cameron-Erdős conjecture*, Proc. Lond. Math. Soc. (3) **108** (2014), no. 1, 44–72.

[3] Noga Alon and Joel H. Spencer, *The probabilistic method*, third ed., Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., Hoboken, NJ, 2008, With an appendix on the life and work of Paul Erdős.

[4] Éric Balandraud, *Coloured solutions of equations in finite groups*, J. Comb. Theory, Ser. A **114** (2007), no. 5, 854–866 (English).

[5] Antal Balog and George Shakan, *On the sum of dilations of a set*, Acta Arith. **164** (2014), no. 2, 153–162 (English).

[6] Antal Balog and Endre Szemerédi, *A statistical theorem of set addition*, Combinatorica **14** (1994), no. 3, 263–268.

[7] József Balogh, *A remark on the number of edge colorings of graphs*, Eur. J. Comb. **27** (2006), no. 4, 565–573 (English).

[8] József Balogh, Hong Liu, and Maryam Sharifzadeh, *The number of subsets of integers with no k-term arithmetic progression*, Int. Math. Res. Not. IMRN **2017** (2017), no. 20, 6168–6186.

[9] József Balogh, Robert Morris, and Wojciech Samotij, *Independent sets in hypergraphs*, J. Amer. Math. Soc. **28** (2015), no. 3, 669–709.

[10] ———, *The method of hypergraph containers*, Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. IV. Invited lectures, World Sci. Publ., Hackensack, NJ, 2018, pp. 3059–3092.

[11] Yacine Barhoumi-Andréani, Christoph Koch, and Hong Liu, *Bivariate fluctuations for the number of arithmetic progressions in random sets*, Electron. J. Probab. **24** (2019), 32 (English), Id/No 145.

[12] F. A. Behrend, *On sets of integers which contain no three terms in arithmetical progression*, Proc. Nat. Acad. Sci. U.S.A. **32** (1946), 331–332.

[13] Fabrício S. Benevides, Carlos Hoppen, and Rudini M. Sampaio, *Edge-colorings of graphs avoiding complete graphs with a prescribed coloring*, Discrete Math. **340** (2017), no. 9, 2143–2160 (English).

[14] Ross Berkowitz, Ashwin Sah, and Mehtaab Sawhney, *Number of arithmetic progressions in dense random subsets of* $\mathbf{Z}/n\mathbf{Z}$, Israel J. Math. (2021), 1–32.

[15] R. C. Bose, *An affine analogue of Singer's theorem*, J. Indian Math. Soc., New Ser. **6** (1942), 1–15 (English).

[16] Emmanuel Breuillard, Ben Green, and Terence Tao, *The structure of approximate groups*, Publ. Math. Inst. Hautes Études Sci. **116** (2012), 115–221.

[17] Boris Bukh, *Sums of dilates*, Comb. Probab. Comput. **17** (2008), no. 5, 627–639 (English).

[18] P. J. Cameron and P. Erdős, *On the number of sets of integers with various properties*, Number theory (Banff, AB, 1988), de Gruyter, Berlin, 1990, pp. 61–79.

[19] Marcelo Campos, *On the number of sets with a given doubling constant*, Israel J. Math. **236** (2020), no. 2, 711–726.

[20] Marcelo Campos, Maurício Collares, Robert Morris, Natasha Morrison, and Victor Souza, *The Typical Structure of Sets With Small Sumset*, Int. Math. Res. Not. IMRN (2021), 1–45, rnab021.

[21] Marcelo Campos, Matthew Coulson, Oriol Serra, and Maximilian Wötzel, *The typical approximate structure of sets with bounded sumset*, arXiv e-prints (2021), arXiv:2108.06253.

[22] Augustin-Louis Cauchy, *Recherches sur les nombres*, J. École Polytech. **9** (1813), 99–116.

[23] Javier Cilleruelo, Yahya O. Hamidoune, and Oriol Serra, *On sums of dilates*, Combin. Probab. Comput. **18** (2009), no. 6, 871–880.

[24] ______, *Addition theorems in acyclic semigroups*, Additive number theory. Festschrift in honor of the sixtieth birthday of Melvyn B. Nathanson, New York, NY: Springer, 2010, pp. 99–104 (English).

[25] Javier Cilleruelo, Oriol Serra, and Maximilian Wötzel, *Sidon set systems*, Rev. Mat. Iberoam. **36** (2020), no. 5, 1527–1548.

[26] D. Conlon and W. T. Gowers, *Combinatorial theorems in sparse random sets*, Ann. Math. (2) **184** (2016), no. 2, 367–454 (English).

[27] Harold Davenport, *On the addition of residue classes*, J. Lond. Math. Soc. **10** (1935), 30–32 (English).

[28] Matt DeVos, *The Structure of Critical Product Sets*, arXiv e-prints (2013), arXiv:1301.0096.

[29] Matt DeVos, Luis Goddyn, and Bojan Mohar, *A generalization of Kneser's addition theorem*, Adv. Math. **220** (2009), no. 5, 1531–1548.

[30] George T. Diderrich, *On Kneser's addition theorem in groups*, Proc. Amer. Math. Soc. **38** (1973), 443–451.

[31] Shan-Shan Du, Hui-Qin Cao, and Zhi-Wei Sun, *On a sumset problem for integers*, Electron. J. Comb. **21** (2014), no. 1, research paper p1.13, 25 (English).

[32] Pál Erdős and Alfréd Rényi, *On random graphs. I*, Publ. Math. **6** (1959), 290–297 (English).

[33] Pál Erdős and Pál Turán, *On a problem of Sidon in additive number theory, and on some related problems*, J. Lond. Math. Soc. **16** (1941), 212–215 (English).

[34] Paul Erdős, Daniel J. Kleitman, and B. L. Rothschild, *Asymptotic enumeration of K_n-free graphs*, Colloq. int. Teorie comb., Roma 1973, Tomo II, 19-27 (1976)., 1976.

[35] Paul Erdős, Miklos Simonovits, and Vera T. Sós, *Anti-Ramsey theorems*, Infinite finite Sets, Colloq. Honour Paul Erdős, Keszthely 1973, Colloq. Math. Soc. Janos Bolyai 10, 633-643 (1975)., 1975.

[36] Yushuang Fan and Salvatore Tringali, *Power monoids: a bridge between factorization theory and arithmetic combinatorics*, J. Algebra **512** (2018), 252–294.

[37] Dmitri V. Fomin, *Is the multiset of n integers uniquely determined by the multiset of its s-sums?*, Amer. Math. Monthly **126** (2019), no. 5, 400–417.

[38] Peter Frankl, Ronald L. Graham, and Vojtěch Rödl, *Quantitative theorems for regular systems of equations*, J. Comb. Theory, Ser. A **47** (1988), no. 2, 246–261 (English).

[39] G. A. Freĭman, *On addition of finite sets. I*, Izv. Vyssh. Uchebn. Zaved., Mat. **1959** (1959), no. 6(13), 202–213 (Russian).

[40] ______, *Addition of finite sets*, Sov. Math., Dokl. **5** (1964), 1366–1370 (English).

[41] ______, *Foundations of a structural theory of set addition*, Translations of Mathematical Monographs, Vol 37, American Mathematical Society, Providence, R. I., 1973, Translated from the Russian.

[42] ______, *Groups and the inverse problems of additive number theory*, Number-theoretic studies in the Markov spectrum and in the structural theory of set addition (Russian), Kalinin. Gos. Univ., Moscow, 1973, pp. 175–183.

[43] Shinya Fujita, Colton Magnant, and Kenta Ozeki, *Rainbow generalizations of Ramsey theory: A survey*, Graphs Comb. **26** (2010), no. 1, 1–30 (English).

[44] Zoltán Füredi, *A proof of the stability of extremal graphs, Simonovits' stability from Szemerédi's regularity*, J. Combin. Theory Ser. B **115** (2015), 66–71.

[45] Anant P. Godbole, Svante Janson, Nicholas W. Locantore Jr., and Rebecca Rapoport, *Random Sidon sequences*, J. Number Theory **75** (1999), no. 1, 7–22 (English).

[46] B. Gordon, Aviezri S. Fraenkel, and E.G. Straus, *On the determination of sets by the sets of sums of a certain order*, Pac. J. Math. **12** (1962), 187–196 (English).

[47] W. T. Gowers, *A new proof of Szemerédi's theorem for arithmetic progressions of length four*, Geom. Funct. Anal. **8** (1998), no. 3, 529–551.

[48] Ben Green, *The Cameron-Erdős conjecture*, Bull. London Math. Soc. **36** (2004), no. 6, 769–778.

[49] ______, *Counting sets with small sumset, and the clique number of random Cayley graphs*, Combinatorica **25** (2005), no. 3, 307–326 (English).

[50] ______, *A Szemerédi-type regularity lemma in abelian groups, with applications*, Geom. Funct. Anal. **15** (2005), no. 2, 340–376.

[51] Ben Green and Imre Z. Ruzsa, *Sum-free sets in abelian groups*, Israel J. Math. **147** (2005), 157–188.

[52] ______, *Freiman's theorem in an arbitrary abelian group*, J. Lond. Math. Soc. (2) **75** (2007), no. 1, 163–175.

[53] Ben Green and Terence Tao, *The primes contain arbitrarily long arithmetic progressions*, Ann. Math. (2) **167** (2008), no. 2, 481–547 (English).

[54] Simon Griffiths, Christoph Koch, and Matheus Secco, *Deviation probabilities for arithmetic progressions and irregular discrete structures*, arXiv e-prints (2020), arXiv:2012.09280.

[55] Yahya O. Hamidoune, *The global isoperimetric methodology applied to Kneser's Theorem*, arXiv e-prints (2007), arXiv:0708.2191.

[56] Yahya O. Hamidoune, *Two inverse results*, Combinatorica **33** (2013), no. 2, 217–230.

[57] Yahya O. Hamidoune and Juanjo Rué, *A lower bound for the size of a Minkowski sum of dilates*, Comb. Probab. Comput. **20** (2011), no. 2, 249–256 (English).

[58] Yahya O. Hamidoune and Oriol Serra, *A note on Pollard's Theorem*, arXiv e-prints (2008), arXiv:0804.2593.

[59] Robert Hancock, Katherine Staden, and Andrew Treglown, *Independent sets in hypergraphs and Ramsey properties of graphs and the integers*, SIAM J. Discrete Math. **33** (2019), no. 1, 153–188 (English).

[60] Carlos Hoppen and Hanno Lefmann, *Edge-colorings avoiding a fixed matching with a prescribed color pattern*, Eur. J. Comb. **47** (2015), 75–94 (English).

[61] Veselin Jungić, Jacob Licht, Mohammad Mahdian, Jaroslav Nešetřil, and Radoš Radoičić, *Rainbow arithmetic progressions and anti-Ramsey results*, Comb. Probab. Comput. **12** (2003), no. 5-6, 599–620 (English).

[62] J. H. B. Kemperman, *On small sumsets in an abelian group*, Acta Math. **103** (1960), 63–88.

[63] Daniel J. Kleitman and Kenneth J. Winston, *On the number of graphs without* 4*-cycles*, Discrete Math. **41** (1982), no. 2, 167–172.

[64] Martin Kneser, *Abschätzung der asymptotischen Dichte von Summenmengen*, Math. Z. **58** (1953), 459–484.

[65] Yoshiharu Kohayakawa, Sang June Lee, Vojtěch Rödl, and Wojciech Samotij, *The number of Sidon sets and the maximum size of Sidon sets contained in a sparse random set of integers*, Random Struct. Algorithms **46** (2015), no. 1, 1–25 (English).

[66] Yoshiharu Kohayakawa, Tomasz Łuczak, and Vojtěch Rödl, *Arithmetic progressions of length three in subsets of a random set*, Acta Arith. **75** (1996), no. 2, 133–163 (English).

[67] Tomasz Ł uczak, *On triangle-free random graphs*, Random Structures Algorithms **16** (2000), no. 3, 260–276.

[68] Vsevolod F. Lev, *Structure theorem for multiple addition and the Frobenius problem*, J. Number Theory **58** (1996), no. 1, 79–88.

[69] ———, *Restricted set addition in groups. I. The classical setting*, J. London Math. Soc. (2) **62** (2000), no. 1, 27–40.

[70] ———, *Restricted set addition in groups. III. Integer sumsets with generic restrictions*, Period. Math. Hungar. **42** (2001), no. 1-2, 89–98.

[71] Vsevolod F. Lev and Pavel Y. Smeliansky, *On addition of two distinct sets of integers*, Acta Arith. **70** (1995), no. 1, 85–91.

[72] Bernt Lindström, *An inequality for B_2-sequences*, J. Comb. Theory **6** (1969), 211–212 (English).

[73] Zeljka Ljujić, *A lower bound for the size of a sum of dilates*, J. Comb. Number Theory **5** (2013), no. 1, 31–51 (English).

[74] W. Mantel, *Problem 28 (Solution by H. Gouwentak, W. Mantel, J. Teixeira de Mattes, F. Schuh and W. A. Wythoff)*, Wiskundige Opgaven **no. 10** (1907), 60–61.

[75] Amanda Montejano and Oriol Serra, *Counting patterns in colored orthogonal arrays*, Discrete Math. **317** (2014), 44–52 (English).

[76] J. J. Montellano-Ballesteros and V. Neumann-Lara, *An anti-Ramsey theorem on cycles*, Graphs Comb. **21** (2005), no. 3, 343–354 (English).

[77] Robert Morris, Wojciech Samotij, and David Saxton, *An asymmetric container lemma and the structure of graphs with no induced* 4*-cycle*, arXiv e-prints (2018), arXiv:1806.03706.

[78] Jitsuro Nagura, *On the interval containing at least one prime number*, Proc. Japan Acad. **28** (1952), 177–181 (English).

[79] E. Nazarewicz, M. O'Brien, M. O'Neill, and C. Staples, *Equality in Pollard's theorem on set addition of congruence classes*, Acta Arith. **127** (2007), no. 1, 1–15.

[80] Kevin O'Bryant, *A complete annotated bibliography of work related to Sidon sequences*, Electron. J. Comb. **DS11** (2004), 39 (English).

[81] John E. Olson, *On the symmetric difference of two sets in a group*, European J. Combin. **7** (1986), no. 1, 43–54.

[82] J. M. Pollard, *A generalisation of the theorem of Cauchy and Davenport*, J. London Math. Soc. (2) **8** (1974), 460–462.

[83] R. A. Rankin, *Sets of integers containing not more than a given number of terms in arithmetical progression*, Proc. Roy. Soc. Edinburgh Sect. A **65** (1960/61), 332–344 (1960/61).

[84] Klaus F. Roth, *On certain sets of integers*, J. Lond. Math. Soc. **28** (1953), 104–109 (English).

[85] Andrzej Ruciński, *When are small subgraphs of a random graph normally distributed?*, Probab. Theory Related Fields **78** (1988), no. 1, 1–10.

[86] Juanjo Rué, Christoph Spiegel, and Ana Zumalacárregui, *Threshold functions and Poisson convergence for systems of equations in random sets*, Math. Z. **288** (2018), no. 1-2, 333–360.

[87] Juanjo Rué and Maximilian Wötzel, *Normal limiting distributions for systems of linear equations in random sets*, arXiv e-prints (2021), arXiv:2111.03526.

[88] Imre Z. Ruzsa, *Solving a linear equation in a set of integers. I*, Acta Arith. **65** (1993), no. 3, 259–282.

[89] ______, *Generalized arithmetical progressions and sumsets*, Acta Math. Hungar. **65** (1994), no. 4, 379–388.

[90] Tom Sanders, *The structure theory of set addition revisited*, Bull. Amer. Math. Soc. (N.S.) **50** (2013), no. 1, 93–127.

[91] A. A. Sapozhenko, *The Cameron-Erdős conjecture*, Dokl. Akad. Nauk **393** (2003), no. 6, 749–752.

[92] David Saxton and Andrew Thomason, *Hypergraph containers*, Invent. Math. **201** (2015), no. 3, 925–992.

[93] Mathias Schacht, *Extremal results for random discrete structures*, Ann. Math. (2) **184** (2016), no. 2, 333–365 (English).

[94] J.L. Selfridge and E.G. Straus, *On the determination of numbers by their sums of a fixed order*, Pac. J. Math. **8** (1958), 847–856 (English).

[95] George Shakan, *Sum of many dilates*, Comb. Probab. Comput. **25** (2016), no. 3, 460–469 (English).

[96] Xuancheng Shao, *On an almost all version of the Balog-Szemerédi-Gowers theorem*, Discrete Anal. **2019** (2019), 18 (English), Id/No 12.

[97] Xuancheng Shao and Wenqiang Xu, *A robust version of Freiman's $3k-4$ theorem and applications*, Math. Proc. Cambridge Philos. Soc. **166** (2019), no. 3, 567–581.

[98] James Singer, *A theorem in finite projective geometry and some applications to number theory*, Trans. Am. Math. Soc. **43** (1938), 377–385 (English).

[99] Christoph Spiegel, *A note on sparse supersaturation and extremal results for linear homogeneous systems*, Electron. J. Comb. **24** (2017), no. 3, research paper p3.38, 19 (English).

[100] ______, *Additive structures and randomness in combinatorics*, Ph.D. thesis, Universitat Politècnica de Catalunya, 2020, Available at http://hdl.handle.net/2117/328203.

[101] Yonutz Stanchescu, *On addition of two distinct sets of integers*, Acta Arith. **75** (1996), no. 2, 191–194.

[102] B. Sudakov, E. Szemerédi, and V. H. Vu, *On a question of Erdős and Moser*, Duke Math. J. **129** (2005), no. 1, 129–155.

[103] E. Szemerédi, *On sets of integers containing no k elements in arithmetic progression*, Acta Arith. **27** (1975), 199–245.

[104] Terence Tao, *Noncommutative sets of small doubling*, European J. Combin. **34** (2013), no. 8, 1459–1465.

[105] Terence Tao and Van Vu, *Additive combinatorics*, Cambridge Studies in Advanced Mathematics, vol. 105, Cambridge University Press, Cambridge, 2006.

[106] Matthew C. H. Tointon, *Freiman's theorem in an arbitrary nilpotent group*, Proc. Lond. Math. Soc. (3) **109** (2014), no. 2, 318–352.

[107] P. Varnavides, *On certain sets of positive density*, J. Lond. Math. Soc. **34** (1959), 358–360 (English).

[108] A. G. Vosper, *The critical pairs of subsets of a group of prime order*, J. London Math. Soc. **31** (1956), 200–205.

[109] Lutz Warnke, *Upper tails for arithmetic progressions in random subsets*, Israel J. Math. **221** (2017), no. 1, 317–365.