
**The Removal Lemma:
algebraic versions and applications**

Lluís Vena Cros



Departament de Matemàtica
Aplicada IV

UNIVERSITAT POLITÈCNICA DE CATALUNYA

The Removal Lemma: algebraic versions and applications

Lluís Vena Cros

A Thesis submitted
for the degree of Doctor of Mathematics
in the Universitat Politècnica de Catalunya

Thesis Advisor
Oriol Serra Albó

Doctoral Program
APPLIED MATHEMATICS

Barcelona, May 2012



Facultat de Matemàtiques
i Estadística

UNIVERSITAT POLITÈCNICA DE CATALUNYA

Lluís Vena Cros
Departament de Matemàtica Aplicada IV
Universitat Politècnica de Catalunya
Edifici C3, Jordi Girona 1-3
E-08034, Barcelona
<lvena@ma4.upc.edu>

Acknowledgements

This thesis would not have been possible without the support, encouragement, long talks, and advice of many people.

First and foremost, I would like to thank my supervisor Oriol Serra for all the encouragement, support, advice, talks, comments and remarks he has given to me. They have turned out to be invaluable. His insightful questions and knowledge have guided me through this subject.

This thesis has its roots in the Master thesis ¹ and the work that it fostered. Thus, this work is beneficiary of my participation in the Spring School in Combinatorics held in Vysoká Lípa, Czech Republic, in 2007 and organized, among others, by Daniel Král'. I would like to specially thank Dan for the collaboration we initiated at the Spring School and that has been extended throughout the years. This thesis owes him a lot.

I would also like to specially thank Balázs Szegedy for his encouragement, wise advice and extensive support as well as for his insightful, deep and clear teachings.

I also have to thank Pablo Candela for many fruitful discussions, Javier Cilleruelo and Boris Bukh for their helpful remarks, Vojtěch Rödl and Mathias Schacht for their comments. I would like to thank Anna Lladó and Marc Noy for their advice and encouragement.

Vull agrair a la meua familia i mi familia, molt especialment als meus pares i a la Diana, pel seu increïble suport i paciència en tots aquests anys. Moltíssimes gràcies!

També vull agrair a tots aquells amb qui he compartit les inquietuds d'aquests anys, llargues i interessants xerrades així com molt bones estones: Aida, Àngela, Anna, Arnau, Carla, Cristina, Dieter, Enric, Èric, Guillem, Hèctor, Itziar, Jordi, Juanjo, Marc, Marcel, Maria, Pere Daniel, Teixi, Vincent, Laura, Thomas, Gemma, als Naranjitos, Andreu, David, Joan, Jose, Litus, Ruben, Branka, Fabia, Miguel, Gina ...

Moreover, I also want to thank those whom I have shared my concerns and worries with, as well as many nice moments, long discussions and talks: Anna, Barry, David, David and Donald, Dylan, Eric, Erik, Ida, John, Josh, Leonel, Mary, Nevena, Ruedi, Tina, Yonsue, Dana, Jan, Diana ...

Let me finish this acknowledgements with a little detour. As I was told in the Spring School, “vysoká lípa” means “tall linden” in Czech and Slovak. The place apparently took its name from a huge linden that was rooted by the road, in front of the hotel where the Spring School was held. Sadly, a big snowstorm took the linden down that January and

¹“The Regularity Lemma in Additive Combinatorics” by the author under the supervision of Oriol Serra

where the huge linden once was, only the roots remained. We could say that a seed was planted during the Spring (School) and the fruit is this thesis.

Abstract

This thesis presents some contributions in additive combinatorics and arithmetic Ramsey theory. More specifically, it deals with the interaction between combinatorics, number theory and additive combinatorics. This area saw a great improvement with the Szemerédi Regularity Lemma [75] and some of the results that followed. The Regularity Lemma and its consequences have become a widely used tool in graph theory, combinatorics and number theory. Furthermore, its language and point of view has deeply changed the face of additive number theory, a fact universally acknowledged by the Abel award given to Szemerédi in 2012. One of the main reasons for the prize has been Szemerédi's theorem, a result regarding the existence of arbitrarily long arithmetic progressions in dense sets of the integers [74], the proof of which uses the Regularity Lemma in a key step.

One of the earlier consequences of the Regularity Lemma was the Removal Lemma for graphs [65, 24, 32] that was used by Ruzsa and Szemerédi in [65] to show Roth theorem [64], regarding the existence of 3-term arithmetic progressions in dense sets of the integers, in a combinatorial way. The Removal Lemma states that in any graph K with few copies of a subgraph, say a triangle, we can remove few edges from K so that the result contains no copy of the subgraph. This has become a key tool in the applications of the so-called Regularity Method, which has extensive literature in combinatorics, graph theory, number theory and computer science. In [41] Green introduced a regularity lemma for Abelian groups as well as an algebraic removal lemma. The removal lemma for groups states that, for a given finite Abelian group G , if there are $o(|G|^{m-1})$ solution to $x_1 + \dots + x_m = 0$ with $x_i \in S \subset G$, then we can remove $o(|G|)$ elements from S to make the set S solution-free.

The main contributions of this work corresponds to extensions of the removal lemma for groups to either more general contexts, like non-necessary Abelian finite groups, or to linear systems of equations for finite Abelian groups. The main goal is to give a comprehensive and more general framework for many results in additive number theory like Szemerédi Theorem.

In particular, we show that the removal lemma for groups by Green can be extended to non-necessary Abelian finite groups. Moreover, we prove a removal lemma for linear systems on finite fields: for every $\epsilon > 0$ there exists a $\delta > 0$ such that if A is a $k \times m$ linear system of equations with coefficients in a finite field \mathbb{F}_q and the number of solutions to $Ax = b$, with $x_i \in S_i \subset \mathbb{F}_q$ is less than $\delta |\mathbb{F}_q|^{m-k}$, then by removing less than $\epsilon |\mathbb{F}_q|$ elements in each S_i we can make the resulting sets solution-free, thus solving a conjecture by Green [41] to that respect. Even more, if A is an integer linear system, G is a finite Abelian group, and the determinantal of A and $|G|$ are coprime, then a similar statement holds. Let us mention that the last result allows us to characterize those linear systems where

any set S with size proportional to G has a nontrivial solution in S , provided $|G|$ is large enough. This extends the validity of Szemerédi's theorem to finite Abelian groups [73].

These extensions of the removal lemma have been used in arithmetic Ramsey theory to obtain counting results for the number of monochromatic solutions of linear systems. The main result in Frankl, Graham and Rödl [28] states that the number of monochromatic solutions of regular systems in integer intervals is in fact a positive proportion of the total number of solutions. We give analogous results for solutions in Abelian groups with bounded exponent, for which the main tool in the torsion-free case cannot be applied. Density versions of these counting results are also obtained, in this case with a full characterization.

Contents

1	Introduction	3
1.1	Background	3
1.1.1	Szemerédi Regularity Lemma	3
1.1.2	Extensions	4
1.1.3	Removal Lemma for graphs and hypergraphs	5
1.1.4	Proofs without using the regularity lemma	6
1.1.5	Removal Lemma vs Analytic methods	6
1.2	Framework of this work	6
1.2.1	Previous use of the Removal Lemma in algebraic settings	7
1.2.2	The Removal Lemma for groups	7
1.3	Contributions of this work	7
1.3.1	Extensions of the Removal Lemma for groups	8
1.3.2	Applications	9
1.3.3	Outputs of this work	9
2	Removal Lemma for Groups	11
2.1	Definitions and basic results	11
2.2	The Removal Lemma for groups	12
2.3	Proof of the Removal Lemma for groups	13
2.4	Extension to some systems of equations	14
2.4.1	Systems of equations for Abelian finite groups	15
2.4.2	Systems of equations for nonabelian groups	18
2.4.3	Examples of graph-representable systems	19
3	Removal Lemma for Finite Fields	21
3.1	The hypergraph Removal Lemma and outline of the proof	23
3.2	Reductions of the system	23
3.3	Hypergraph representation and proof of Theorem 3.1	25
3.4	Proof of Lemma 3.6	27
3.4.1	Example of a system and a matrix C	29

4	Removal Lemma for Abelian Groups	31
4.1	Removal Lemma for linear systems over finite Abelian groups	31
4.2	Circular Unimodular Matrices	32
4.3	A reduction lemma	36
4.4	Proof of Theorem 4.1	41
4.5	On the condition over the determinantal	41
4.6	Example of the construction of the extension of the matrix A	42
5	Applications to Arithmetic Ramsey Theory	45
5.1	Introduction	45
5.2	Number of monochromatic solutions inside \mathbb{F}_q^N	46
5.2.1	Introduction	46
5.2.2	Proof of Theorem 5.5	47
5.3	Number of monochromatic solutions inside \mathbb{Z}_n^N	49
5.3.1	Introduction	49
5.3.2	Proof of Theorem 5.7	50
5.3.3	Proof of Lemma 5.8	57
5.4	Extension for bounded torsion groups	64
5.5	Density case	65
5.6	Remarks	66
6	Final remarks and future work	69
6.1	Extension to Polynomials	69
6.2	Orthogonal Arrays	71
6.2.1	General considerations	71
6.2.2	Construction	72
6.2.3	Example	74
6.3	On the dimension of the representation	75
6.3.1	The case of graphs (for systems in Abelian groups)	76
6.3.2	Example of a system with lower dimension	78
6.3.3	Comments on the complexity of the system	79
6.4	Symmetry-preserving lemma	80
6.5	Open problems and future work	80
	Bibliography	83
	Index	89

Introduction

This thesis presents some contributions in additive combinatorics and arithmetic Ramsey theory. More specifically, it deals with the interaction between combinatorics, number theory and additive combinatorics. This area saw a great improvement with the Szemerédi Regularity Lemma [75] and some of the results that followed. The Regularity Lemma and its consequences have become a widely used tool in graph theory, combinatorics and number theory. Szemerédi Theorem, a result regarding the existence of arbitrarily long arithmetic progressions in dense sets of the integers [74], became one of the earliest and more prominent statements where the Regularity Lemma was used in a key step.

The main contributions of this work are generalizations of the Removal Lemma for groups to either more general contexts, like non-necessary Abelian finite groups, or to linear systems of equations for finite Abelian groups. The main purpose behind these extensions is to give a comprehensive and general framework for many results in additive number theory like Szemerédi Theorem, where solutions to linear systems of equations are involved. In this case we have shown that the linear systems behave in a similar way as graphs do: they both fulfill, at least, a Removal Lemma statement.

1.1 Background

1.1.1 Szemerédi Regularity Lemma

The Szemerédi Regularity Lemma [75] has been an important tool in graph theory in particular and for combinatorics in general, with various ramifications to computer science and number theory. It roughly says that any graph can be partitioned into finitely-many clusters of vertices, each of equal size, such that between most of the clusters we find bipartite graphs that have all the properties expected in a random bipartite graph. In other words, between most of the clusters we have a quasi-random graph (see [17]).

In the Regularity Lemma, the number of clusters can be made as large as desired and the quasi-randomness of the pairs, also called regularity, can be made as quasi-random as we need. Since the vast majority of the pairs are regular, if the number of clusters is large, then most of the edges will be between regular pairs. Thus we can approximate any graph by using a finite structure of quasi-random bipartite graphs.

In [55], Lovász and Szegedy present the Szemerédi Regularity Lemma as an analog for graphs of the result in analysis regarding the arbitrarily-precise approximation of a measurable function using an step function with finitely-many steps. Moreover, in the same paper, they use this analogy to complete the set of graphs, with a certain norm, using

symmetric measurable functions from $[0, 1]^2$ to $[0, 1]$ with some additional properties (see [55] for further details.) Another perspective from which Szemerédi's regularity lemma can be seen, due to Tao [76], is as a structure theorem for random variables in a product probability space.

From its earlier formulations, the Regularity Lemma has seen many applications in graph theory as well as in number theory. One of its first uses, thanks to a primitive version of this result valid when K is a bipartite graph, can be found in the original proof of Szemerédi Theorem [74]. Many of these applications in graph theory have been collected in the splendid surveys [48] and [47]. Among many others, we can cite a short proof of Erdős-Stone theorem [21] (the original in [25]), an approximate versions of Lovász-Komlós-Sós conjecture [58] and Lovász conjecture for large enough graphs [82].

Let us mention that the Regularity Lemma can be “weakened” and “strengthened”. The upper bound on the number of sets in the partition of the Regularity Lemma, M , is roughly of order

$$M \approx 2^{\left\{ 2^{2^{2^{\dots}}} \right\} \approx \epsilon^{-5}}.$$

Gowers showed in [35] that a tower-type bound is necessary. However, if one asks for weaker properties to the partition, better bounds can be achieved [31]. On the other side, some applications, for example results in property testing like [6], may require stronger properties to the partition (see [3] also as [4].) The strong regularity lemmas require even worse bounds (Wowzer type) than the ones coming from the Szemerédi Regularity Lemma (tower type) [19].

As it is observed in [55], we can have a continuous range of bounds with a continuous range of properties that can be asked to the partition. Let us note that a strong version can be obtained from weaker versions by iteratively applying the latter.

1.1.2 Extensions

The Regularity Lemma has seen extensions to finitely edge colored graphs [48] as well as to directed graphs [5]. The proofs divide the edge set in different classes, like the different colors, or different directions of the edges in a bipartite graph, and then regularize for each of these classes of edges.

Even more, Szemerédi Regularity Lemma has seen extension to hypergraphs. The first results were due to Frankl and Rödl in [29] and Chung in [16], but the regularity lemma was not strong enough to show the analogous implications for hypergraphs that the Szemerédi Regularity Lemma has for graphs. In [30], Frankl and Rödl proved a version of the hypergraph regularity lemma for 3-uniform hypergraphs that was strong enough to show similar implications as the result for graphs, like the (general) removal lemma, for example.

Recently, several authors have found versions of the regularity lemma for any k -uniform hypergraphs equivalent to the graph case, in the sense that they imply similar results and have similar consequences to the ordinary counting and removal lemmas. Let us mention Rödl and Skokan [62], Gowers [37], Tao [77], Elek and Szegedy [23], Ishigami [46] or Rödl and Schacht [61].

1.1.3 Removal Lemma for graphs and hypergraphs

One of the application of the Szemerédi Regularity Lemma is the Removal Lemma for graphs:

Theorem 1.1 (Removal Lemma for graphs). *Let K be a graph of order n and let H be a graph of order h . For every $\epsilon > 0$, there exists a $\delta = \delta(\epsilon, h) > 0$ such that, if the number of copies of H in K is less than δn^h , then we can remove ϵn^2 edges from K to make it H -free.*

We can see the Removal Lemma as a distance result: if a graph has not many copies of a certain graph H , then it can be made free of copies of such a graph H by deleting not many edges. This is, if a graph has not many copies of H then it is close to an H -free graph. Moreover, the Removal Lemma can be seen as a reciprocal statement of the following obvious one: if a graph has not many edges, then it does not have many copies of a given graph H . The straight converse is not true since there are graphs with not many copies of H but with lots of edges (for example the Turán graphs where H is, for example, a complete graph on 3 vertices and K is a complete bipartite graph.) However, the Removal Lemma states that, if there are not many copies of H in K , then they are supported over few edges, thus it can be seen as a sensible converse.

The first appearance of the Removal Lemma was in Ruzsa and Szemerédi [65] for K being a triangle. The authors used it to give a purely combinatorial proof of Rőth's theorem on the existence of 3-term arithmetic progressions in subsets of integers with positive upper density [64]. However, the bounds on the proportion of integers for which we can ensure a 3-term arithmetic progression given by the Fourier analytic proof by Roth [64] are better than the ones coming from the Removal Lemma.

The version of the Removal Lemma for the complete graphs K_r was first proved by Erdős, Frankl and Rödl [24]. In [32], Füredi showed the general case, which also appears in the surveys [48], [47].

The proof of the Removal Lemma by means of the Regularity Lemma uses what has been called the Szemerédi Regularity Method. The Regularity Lemma is a graph approximation result that, using finitely many information, retrieves, among other information, the approximate number of copies of any bounded-sized subgraphs. In particular, if a graph has not many copies, the regularity partition cannot see them. As the number of edges that the regularity partition does not see is small, if a graph K has not many copies of a fixed sized subgraphs H , then it is close, by removing few edges, to a H -free graph K' .

The Removal Lemma has seen extensions to colored graphs [48], using finitely many colors, and to directed graphs [5].

More recently, the Removal Lemma has seen extensions to the hypergraph case. Rödl and Skokan in [63] have shown it by using the regularity lemma for hypergraphs [62] via Nagel, Rödl and Schacht's counting result in [56] (see also [60]). Gowers, in [37] focused more on a quasi-randomness approach (see also [36]). Tao in [77] uses a more probabilistic and information-theoretic approach to show a Removal Lemma for hypergraphs. Elek and Szegedy in [23] show the result by means of a measure-theoretic approach using non-standard analytic techniques such as ultraproducts.

Let us mention that the Removal Lemma for hypergraphs has seen extensions to directed and colored versions, using finitely many colors. These versions can be obtained in a similar way as the directed and colored versions of the removal lemma for graphs are

obtained from the Szemerédi Regularity Lemma: by classifying the edges into different classes according to directions or colors. More explicit statements of directed and colored versions of the hypergraph removal lemma can be found in Austin and Tao [7] and in Ishigami [46].

The Removal Lemma has seen applications in many areas especially in property testing (see [5, 2, 34, 7] for some examples) and, using the hypergraph version in general, to show Szemerédi Theorem [74], its multidimensional version (see [33] for the first proof of the result and [37] for a proof using the hypergraph removal lemma), or other related results [71].

1.1.4 Proofs without using the regularity lemma

Recently, Fox [27] gave an alternative proof of the Removal Lemma without using the Szemerédi Regularity Lemma for graphs and was able to substantially lower the bounds that relate the number of copies of a graph and the number of edges that needs to be deleted, although the relation is still of tower-type. It is not known whether a similar approach works for hypergraphs. Fox’s result involves revisiting the proof of Szemerédi Regularity Lemma with a less strong notion of “regularity”, but sufficient for the purposes of the removal lemma.

All the known proofs for the removal lemma of hypergraphs come from a regularity lemma in a direct or indirect way: either the proof of the removal lemma uses a regularity lemma (like [56, 63], [36] and [77]), or the strategy and the techniques used to show the result output both at the same time (like in [23]). The Regularity Lemma and the Removal Lemma seem to be intimately tied.

1.1.5 Removal Lemma vs Analytic methods

Since the very beginning, the graph removal lemma has been applied to additive number theory. The most famous example is the proof of Roth’s theorem [64] by Ruzsa and Szemerédi [65] using the triangle removal lemma (to be precise, they used the $(6, 3)$ -theorem). However, as we have mentioned earlier, the removal lemma does not have a good relation between the number of copies of a certain subgraph and the number of edges that support those graphs. This bad relation is translated to the bound for the proportion of integers needed to obtain a 3-term arithmetic progression in $[1, N]$, which can be dramatically improved by using Fourier Analysis techniques. The best currently known bound is due to Bourgain [10].

The removal lemma seems to be a powerful and versatile tool, yet other, more specific, tools work much better in many of its applications.

1.2 Framework of this work

The focus of this work has been to translate the use and philosophy of the so-called Szemerédi Regularity Method from the realm of graphs and hypergraphs, for which it was originally designed, to more algebraic settings, namely to linear equations over finite fields, finite Abelian groups or finite groups in general. In particular, we develop removal lemmas for several algebraic settings and discuss some applications.

1.2.1 Previous use of the Removal Lemma in algebraic settings

As we have mentioned earlier, one of the first examples of this machinery is the proof of the Röth's theorem [64] by Ruzsa and Szemerédi [65] using the graph removal lemma where H is a triangle (indeed, the so-called $(6, 3)$ -theorem), although the bounds of the original proof by Roth are better.

Indeed, one of the first uses of the Regularity Lemma was in the original proof of the Szemerédi's theorem regarding the existence of arbitrarily long arithmetic progressions in dense subsets of the integers [74]. However, in the original proof of Szemerédi's theorem, the Regularity Lemma is one of the many tools and results used, whereas in [65] the Removal Lemma plays the central role. The Removal Lemma for triangles has also been used by Solymosi [69] to show a 2-dimensional version of Szemerédi's result proved by Ajtai and Szemerédi [1] using a different argument: in dense sets ¹ of $[0, N]^2$, there are 3 points of the form $\{(a, b), (a + d, b), (a, b + d)\}$ for some d .

In [30], along with showing regularity and removal lemmas for 3-uniform hypergraphs, Frankl and Rödl highlighted the usefulness of the hypergraph removal lemma for general k -uniform hypergraphs to show Szemerédi's theorem in the case of $(k + 1)$ -term arithmetic progressions. Solymosi [70] used Frankl and Rödl's result [30] to show that any dense set of $[0, N]^3$ has four points in a configuration like $\{(a, b, c), (a + d, b, c), (a, b + d, c), (a + d, b + d, c + d)\}$. Moreover, Solymosi observed that a removal lemma for k -uniform hypergraphs (more precisely, the equivalent of the $(6, 3)$ -theorem for hypergraphs), would imply the multidimensional version of the Szemerédi's theorem, which was proved by Fürstenberg and Katznelson [33] by means of ergodic theory.

1.2.2 The Removal Lemma for groups

In [41] Green introduces a Regularity Lemma for Abelian groups using Fourier Analysis. From the Regularity Lemma, he deduces the Removal Lemma for groups, which is an algebraic version of the Removal Lemma for graphs. The Removal Lemma for groups states that, for any Abelian group G and each m -subsets of G , S_1, \dots, S_m , $S_i \subset G$, if the number of solutions to $x_1 + \dots + x_m = 0$, $x_i \in S_i$, is small, then we can, by removing not many elements, make the new sets solution-free. Due to the use of Fourier Analysis in its proof, the result was restricted to the finite Abelian group setting.

Let us notice that, if the sets S_i are small, then the number of solutions to $x_1 + \dots + x_m = 0$ with $x_i \in S_i$ will be necessarily small. The removal lemma for groups can be interpreted as a reciprocal of this observation: if the number of solutions is small, then the solutions are supported on few elements.

1.3 Contributions of this work

Let us now present the contributions of this work. The first part, corresponding to the contributions in Chapters 2-4, involve extensions of the removal lemma for groups. In particular we show removal lemmas for non-necessarily Abelian groups, a removal lemma

¹By dense we understand that the number of points is, for N large enough, at least a constant times N^2 .

for linear systems on finite fields and a removal lemma for integer linear systems over Abelian groups. The second part, Chapter 5, consists of applications to arithmetic Ramsey theory of the results presented in Chapters 2-4.

1.3.1 Extensions of the Removal Lemma for groups

This work presents several extensions of the result by Green. The first result, Theorem 2.1, extends the removal lemma for groups to non-necessarily Abelian groups. The proof uses the removal lemma for directed graphs [5]. The main idea of the proof is to represent the equation using a directed graph in which the edges have identifications with elements of the group and each copy of a particular subgraph is identified with a solution to the equation $x_1 + \dots + x_m = 0$. Then we apply the removal lemma for graphs to this construction and, using a pigeonhole principle, extract the information of which elements should be removed in order to delete all the solutions. This construction is similar to all the other applications of the removal lemma for graphs in the algebraic settings presented in this thesis.

This idea of representing the elements of the algebraic structure by edges and certain subgraphs as solutions is used in the other extensions of the result. In particular, we use a $(k+1)$ -uniform hypergraph to represent the given linear system with k equations in a finite field and we deduce a removal lemma for linear systems over finite fields, Theorem 3.1, thus solving a conjecture by Green [41]. This result was independently obtained by Shapira [68] although his representation uses s -uniform hypergraphs with $s \geq k+1$.

The uniformity, or size of the edges, of the hypergraphs can also be denoted as its dimension and it seems to be related with the complexity of the linear systems; more specifically, with the different interactions of each variable with various equations of the system. Although some systems can be represented by graphs (see Theorem 2.4 and Theorem 2.5 and the discussion in Chapter 6), higher dimensional hypergraphs appear to be a better framework for those results. An indication for the suitability of the higher dimension hypergraphs to represent the systems of equations can be indicated by the use, among the combinatorial removal lemmas, of the hypergraph version of the removal lemma to prove Szemerédi Theorem ([30] or [70] and [37]); the dimension of the hypergraphs increase with the length of the progression. Furthermore, there are systems that cannot be represented using graphs (not representable using graphs by the techniques presented in this work, see Chapter 6 for a further discussion.)

Allowing an even higher dimension on the hypergraph, in comparison to the dimension used to represent a similar system in the finite fields case, we have been able to show a removal lemma for linear systems of equations with integer coefficients in finite Abelian groups Theorem 4.1. The construction of the hypergraph used to represent the linear system and its solution set is similar to the one used for Theorem 3.1, but we increase the dimension of the edges to accommodate the fact that the product by an integer is not always a bijective application. Furthermore, there are some restrictions in the statement of Theorem 4.1: the order of the group has to be coprime with the k -th determinantal divisor of the matrix (see Chapter 4 for details.) Although the result seems to be true if the condition is removed, more refined ideas and a deeper understanding of the construction and the problem seems to be needed. We refer the reader to Section 6.5 for further discussion.

These extensions of the removal lemma for groups open the way to continuous versions of the algebraic removal lemma to compact Abelian groups, a project already started

by Candela and Sissak [15] for the circle and with a continuation to higher dimensional compact Abelian groups by Szegedy and the author.

1.3.2 Applications

As it has been mentioned, the removal lemma can be used to provide counting versions of certain results. For example, a well known application shows that the number of 3-term arithmetic progressions in a subset of $[1, N]$ with density ϵ grows as cN^2 . This counting statement was due to Varnavides [78]. Most of these counting applications follow the same scheme: they use the contrapositive of the removal lemma when unavoidable and fairly popular structures are present. In the previous example the unavoidable structures are the trivial 3-term arithmetic progressions formed with three copies of the same element.

Using this reasoning, we can characterize the $k \times m$ integer linear systems A for which any set of positive density in a finite Abelian group G will have, at least, $c|G|^{m-k}$ solutions to $Ax = 0$ (Theorem 5.27 and Theorem 5.28). This result is similar to the one found in [28] by Frankl, Graham and Rödl for the integers.

Moreover, using these algebraic versions of the removal lemma, we can find a Roth theorem for finite groups (Corollary 2.2 or its counting version), as well as versions of Szemerédi Theorem for finite fields, (Corollary 3.3) or for Abelian groups in general (showed by Szegedy using a similar framework in [73].)

Combined with the appropriate Ramsey result (Theorem 5.4 or the specifically developed for this application, Lemma 5.8), the suitable extension of the removal lemma (Theorem 3.1 or Theorem 4.1 respectively) allows us to find some asymptotic counting results for monochromatic structures. In particular, Theorem 5.5 shows that: for any r -coloring of $\mathbb{F}_q \setminus \{0\}$, the number of monochromatic solutions to the linear system $Ax = 0$ is proportional to $|\mathbb{F}_q|^{m-k}$, provided that A is a $k \times m$ matrix over \mathbb{F}_q and fulfills analogous conditions to Rado's characterization of partition regular systems (see [59] or Section 5.1).

Furthermore, Theorem 5.26 shows that a similar counting result can be achieved for finite Abelian groups where the maximal order of the elements is bounded. In this case, the matrix of the system has integer coefficients and its column space has certain properties that ensure monochromatic solutions where all the elements have maximal order (Lemma 5.8). Finally we give a characterization of the systems for which one can state a density version. Namely, every set with positive density has a positive proportion of the total number of solutions. These applications are presented in more detail in Chapter 5.

We discuss in Chapter 6 some further directions prompted by our results, including the consideration of polynomial versions, extensions to orthogonal arrays (where the removal lemma does not hold in general), the relation with complexity, the conditions on the determinantal of the matrix in the case of Abelian groups or the extension to arbitrary abelian groups of the counting results.

1.3.3 Outputs of this work

The results presented in this thesis have been published in several journals. The extension of the removal lemma for groups to non-necessarily Abelian groups, Theorem 2.1, as well as Theorem 2.4 and Theorem 2.5 can be found in a joint paper with Král' and Serra [50].

Also with Král' and Serra, the removal lemma for finite fields, Theorem 3.1 can be found in [52].

The removal lemma for Abelian groups, Theorem 4.1, has been published in the arXiv as a preprint with Král and Serra [51]; the work has been submitted and is currently under review.

The algebraic removal lemmas Theorem 3.1 and Theorem 4.1 have seen applications in Arithmetic Ramsey theory. For instance, Theorem 5.5 is an application of the removal lemma for finite fields to qualitatively count monochromatic solutions for linear systems in finite fields. Similarly, Theorem 5.7 or Theorem 5.26 are two counting statement regarding monochromatic solutions for integer systems in finite Abelian groups. The density version of these counting results, Theorem 5.28, is also addressed. All these applications can be found in a preprint with Serra posted on arXiv [66]; this work has been submitted for publishing.

Let us mention that the removal lemma for finite fields, which was independently proved by Shapira [68], has been used by Candela and Sissak to show a removal lemma for configurations on the circle [15].

Removal Lemma for Groups

We begin this chapter by introducing some definitions and basic results that will be useful throughout the thesis. Afterwards we proceed to present the Removal Lemma for non-necessary Abelian groups, which is a natural extension of the Removal Lemma for groups by Green [41]. The original proof uses Fourier Analysis therefore is restricted to the Abelian group setting. Here we present a graph-theoretic argument using the Removal Lemma for directed graphs that allows us to extend the Removal Lemma for groups from finite Abelian groups to finite groups in general.

2.1 Definitions and basic results

Let us begin by recalling some definitions that are used throughout the work. A graph H is a pair (V, E) , where V is the set of vertices, E is the set of edges, $|V|$ is called the order of the graph and $|E|$ its size. If E is a collection of two-sets of V , $E \subset \binom{V}{2}$ then H is said to be an *undirected graph* or a graph for short. If E is a collection of ordered pairs, $E \subset V^2$, then H is said to be a *directed graph*. Moreover, if $E \subset \binom{V}{i}$, with $i \geq 3$, then H is said to be an *i -uniform hypergraph*.

If a graph has t distinguished sets of vertices $V_1 \cup \dots \cup V_t = V$ and all the edges have, for every i , at most one vertex in every V_i , we say that the graph/hypergraph is *t -partite*.

A *circuit* c in H is a set of edges $\{e\}_{e \in c}$ such that there exists an ordering of the edges that allows us to write $e_i = \{v_i, v_{i+1}\}$ with $e_{|c|} = \{v_{|c|}, v_1\}$. A *cycle* is a circuit in which no edge or vertex is repeated.

For a given positive integer r and a graph or a hypergraph H , the function $\chi : E \rightarrow [r]$ is said to be an *r -coloring*, and $\chi(e)$, $e \in E$, is said to be the color of the edge e .

A map ϕ is called a *graph homomorphism* between two graphs $H' = (V', E')$ and $H = (V, E)$ if $\phi : V' \rightarrow V$ and, for every $\{v_1, v_2\} \in E'$, then $\{\phi(v_1), \phi(v_2)\} \in E$; this is, ϕ induces a well defined map between the edge sets of H' and H . A graph $H' = (V', E')$ is a *subgraph* of H if there exists an injective graph homomorphism between H' and H . A copy of H' in H is the pair (H', ϕ) , where ϕ is the homomorphism that exhibits H' as a subgraph of H . We say that two copies (H', ϕ_1) and (H', ϕ_2) of H' in H are *edge-disjoint* if $\{\phi_1(e)\}_{e \in E'} \cap \{\phi_2(e)\}_{e \in E'} = \emptyset$. Notice that these notions can be naturally extended to hypergraphs.

Let A be an integer matrix with k rows and m columns. Assume that $k \leq m$. We define the *i -th determinantal* as the greatest common divisor of all the $i \times i$ square submatrices of A . This notion can be found in Newman's book [57] as determinantal divisor.

Recall also the fundamental theorem of finite Abelian groups: if G is a finite Abelian group of order n then there are some positive integers n_1, \dots, n_s with $n_1 | n_2 | \dots | n_s$ and $n_1 n_2 \dots n_s = n$ such that $G \cong \prod_{i=1}^s \mathbb{Z}_{n_i}$, where \mathbb{Z}_{n_i} denotes the cyclic group with n_i elements.

We denote by $[1, N]$ the set of integers from 1 to N .

2.2 The Removal Lemma for groups

In the following sections we show the Removal Lemma for groups:

Theorem 2.1 (Kráľ, Serra, V. [50]). *Let G be a finite group of order N . Let A_1, \dots, A_m , $m \geq 2$, be sets of elements of G and let g be an arbitrary element of G . If the equation $x_1 x_2 \dots x_m = g$ has $o(N^{m-1})$ solutions with $x_i \in A_i$, then there are subsets $A'_i \subseteq A_i$ with $|A_i \setminus A'_i| = o(N)$ such that there is no solution of the equation $x_1 x_2 \dots x_m = g$ with $x_i \in A'_i$.*

Rigorously speaking, Theorem 2.1 asserts that for every $\delta > 0$ there exists $\delta'(\delta, m)$ such that, if the equation $a_1 \cdot \dots \cdot a_m = 0$ has less than δN^{m-1} solutions with $a_i \in A_i$, then there are subsets $A'_i \subseteq A_i$, $|A_i \setminus A'_i| \leq \delta' N$ such that the equation $a_1 \cdot \dots \cdot a_m = 0$ has no solution with $a_i \in A'_i$, and the value of δ' tends to 0 as $\delta \rightarrow 0$. Let us emphasize that the value of δ' does not depend on the order N of the group G (nor its structure).

The idea of our proof of Theorem 2.1 comes from the original proof of Ruzsa and Szemerédi [65] of Roth's theorem [64] on the existence of 3-term arithmetic progressions in sets of integer in $[1, N]$ with positive density, this consequence being one of the main motivations of their Triangle Removal Lemma. Theorem 2.1 can be similarly used to prove an analogous version of Roth's theorem for groups.

Corollary 2.2. *Let G be a finite group of odd order N and A a subset of its elements. If the number of solutions of the equation $xz = y^2$ with $x, y, z \in A$ is $o(N^2)$, then the size of A is $o(N)$.*

Proof. Apply Theorem 2.1 to the equation $x_1 x_2 x_3^{-1} = 1$ and $A_1 = A_2 = A$, $A_3 = A^2$. We need that the map $x \rightarrow x^2$ is a bijection: indeed, Lagrange's theorem yields that $x^N = 1$ for every $x \in G$. If it holds that $x^2 = y^2$ for two elements of G , then $x^N x^{-2\lfloor N/2 \rfloor} = y^N y^{-2\lfloor N/2 \rfloor}$ and consequently $x = y$. \square

In the proof of Theorem 2.1, we use a directed cycle to represent an equation. Using similar ideas, we can use a directed graph to represent some systems of equations and obtain a similar result as Theorem 2.1. Section 2.4 is devoted to this results.

The use of directed graphs to represent systems of equations limits us in the typology of systems of equation we can deal with (see Chapter 6 for a further discussion). On the other side, allows us to prove the result using the removal lemma for graphs, instead of the removal lemma for hypergraphs, and for the more general framework of general finite groups whereas the results in Chapter 3 and Chapter 4 are restricted to finite fields or finite Abelian finite groups. Following the result of Fox [27], the use of the removal lemma for graphs ensures better relations between the bound on the number of solutions and how many elements we should remove.

2.3 Proof of the Removal Lemma for groups

In our arguments the following removal lemma, consequence of a variant of Szemerédi Regularity Lemma for directed graphs [5], becomes useful:

Lemma 2.3 (Alon and Shapira [5] Lemma 4.1). *Let H be a fixed directed graph of order h . If K contains less than $o(n^h)$ copies of H , there exists a set E of at most $o(n^2)$ arcs of K such that the graph obtained from K by removing the arcs of E is H -free.*

The proof of Theorem 2.1 consists in constructing a blow-up graph of a small graph H such that any solution of the equation gives rise to N edge-disjoint copies of H and every copy of H comes, in fact, from a solution of the equation; this construction is similar to that in [65] or in [70]. We then apply the removal lemma for graphs and, by a pigeonhole principle, reduce the $o(N^2)$ arcs from Lemma 2.3 to the $o(N)$ elements stated in Theorem 2.1.

Proof. [Proof of Theorems 2.1] Fix $\delta_0 > 0$ and $m \geq 2$. Let G be a finite group of order N , let g be an element of G and let A_1, \dots, A_m be sets of elements of G .

We define an auxiliary directed graph K whose vertex set is the set $G \times \{1, \dots, m\}$, i.e., they are pairs formed by an element of the group G and an integer between 1 and m . There is an arc in K from a vertex (x, i) , $1 \leq i \leq m-1$, to a vertex $(y, i+1)$ if there exists an element $a_i \in A_i$ such that $xa_i = y$. This arc is labeled by the pair $[a_i, i]$. The digraph K also contains an arc from a vertex (x, m) to a vertex $(y, 1)$ if there exists an element $a_m \in A_m$ such that $xa_m g^{-1} = y$. This arc is labeled by the pair $[a_m, m]$. Let $N_0 = mN$ denote the order of K . Note that, for each element $a_i \in A_i$, K contains exactly N arcs labelled with $[a_i, i]$.

Observe that any directed cycle of K with length m gives a solution of the equation: if $[a_1, 1], [a_2, 2], \dots, [a_m, m]$ are the labels of the arcs in the cycle and it contains the vertex $(z, 1)$, then $za_1 a_2 \dots a_m g^{-1} = z$ by the definition of K . In the opposite way, each solution a_1, \dots, a_m of (2.2) corresponds to N directed cycles of length m in K :

$$(z, 1), (za_1, 2), (za_1 a_2, 3), \dots, (za_1 \dots a_{m-1}, m), (za_1 \dots a_m g^{-1}, 1) = (z, 1) \quad (2.1)$$

one for each of the N distinct possible choices of $z \in G$. These N directed cycles are vertex disjoint (and thus edge disjoint) since it is possible to recover the value of z from any vertex contained in the cycle.

Suppose, using the hypothesis, that there are less than $\delta_0 N^{m-1}$ solutions of the equation

$$x_1 x_2 \dots x_m = g \text{ with } x_i \in A_i. \quad (2.2)$$

By the correspondence of the cycles of K and the solutions of (2.2), the directed graph K contains no more than $\delta_0 N^m$ distinct directed cycles of length m .

Apply Lemma 2.3 to K , the directed cycle of length m as H and with $\delta = \delta_0/m^m$: since K has less than $\delta_0 N^m = \delta N_0^m$ copies of the directed cycle of length m , there is a set E of at most $\delta' N_0^2$ arcs such that $K - E$ contains no directed cycle of length m with some δ' depending only on δ and m .

Let B_i be the set of those elements $a \in A_i$ such that E contains at least N/m arcs labeled with $[a, i]$. Since $|E| \leq \delta' N_0^2$, the size of each B_i is at most $m|E|/N \leq \delta' m^3 N$. Set $A'_i = A_i \setminus B_i$. Since the size of B_i is bounded by $\delta' m^3 N$, δ' depends on δ and m only, and

$\delta' \rightarrow 0$ as $\delta_0 \rightarrow 0$, the theorem will be proven after we show that there is no solution of the equation (2.2) with $a_i \in A'_i$.

Assume that there is a solution with $a_i \in A'_i$ of the equation (2.2). Consider the N edge disjoint directed cycles of length m corresponding to a_1, \dots, a_m which are given by (2.1). Each of these N cycles contains at least one of the arcs of E and the arcs of these N edge disjoint cycles are labelled only with the pairs $[a_1, 1], [a_2, 2], \dots, [a_m, m]$. Since these N directed cycles are disjoint and the set E contains at least one arc of each of them, the set E contains at least N/m arcs labelled $[a_i, i]$ for some $1 \leq i \leq m$. Consequently, $a_i \in B_i$ and thus $a_i \notin A'_i$. We conclude that there is no solution of (2.2) with $a_i \in A'_i$. \square

Let us notice that we have chosen H to be a directed cycle and not an undirected cycle to break some internal graph homomorphisms from H to itself that a cycle have. Those would have created more copies of H in the constructed blowup graph than the ones coming from solutions of the equation. Even though we could have used directed hypergraphs to break those internal homomorphisms in the blow-up hypergraph coming from H , in the next chapters we have used colored hyperedges instead. The coloring breaks the homomorphisms in clearer way.

2.4 Extension to some systems of equations

In this section we consider an extension of Theorem 2.1 which applies to a particular class of systems of equations, which however holds in general finite groups and its proof just requires the removal lemma for graphs.

We start with the Abelian version of the extension. Let G be an Abelian group (with additive notation) and consider an equation system of the following type:

$$\left. \begin{array}{r} \epsilon_{11}x_1 + \cdots + \epsilon_{1m}x_m = 0 \\ \vdots \\ \epsilon_{k1}x_1 + \cdots + \epsilon_{km}x_m = 0 \end{array} \right\} \quad (2.3)$$

where $\epsilon_{ij} \in \{-1, 0, 1\}$, $k \geq 1$ and $m \geq 2$. The vector $(\epsilon_{i1}, \dots, \epsilon_{im})$ is referred to as the *characteristic vector* of the i -th equation. We say that the system (2.3) is *graph representable* by a directed graph H with m arcs (one for each variable in the system) if the characteristic vectors of cycles in H are precisely integer linear combinations of the characteristic vectors of the equations, see Section 2.4.1 for details. With this notation, we state the following result:

Theorem 2.4. *Let G be a finite Abelian group of order N . Let A_1, \dots, A_m , $m \geq 2$, be sets of elements of G . If the equation system (2.3) is graph-representable and has $o(N^{m-k})$ solutions with $x_i \in A_i$, then there are subsets $A'_i \subseteq A_i$ with $|A_i \setminus A'_i| = o(N)$ such that there is no solution of the system (2.3) with $x_i \in A'_i$.*

Theorem 2.4 can be also extended to non-abelian groups at the expense of strengthening the notion of graph representability. With this strong version which is explained in Section 2.4.2 the same technique allows us to prove:

Theorem 2.5. *Let G be a finite group of order N written multiplicatively. Let A_1, \dots, A_m ,*

$m \geq 2$, be sets of elements of G . Consider the equation system

$$\left. \begin{array}{cccc} x_{\sigma_1(1)}^{\epsilon_{1\sigma_1(1)}} & \cdots & x_{\sigma_1(m)}^{\epsilon_{1\sigma_1(m)}} & = 1 \\ \vdots & & \vdots & \vdots \\ x_{\sigma_k(1)}^{\epsilon_{k\sigma_k(1)}} & \cdots & x_{\sigma_k(m)}^{\epsilon_{k\sigma_k(m)}} & = 1 \end{array} \right\} \quad (2.4)$$

where $\sigma_1, \dots, \sigma_k$ are permutations of $[1, m]$, $\epsilon_{ij} \in \{-1, 0, 1\}$, $k \geq 1$, $m \geq 2$. If the system is strongly graph-representable and has $o(N^{m-k})$ solutions with $x_i \in A_i$, then there are subsets $A'_i \subseteq A_i$ with $|A_i \setminus A'_i| = o(N)$ such that there is no solution of the system (2.3) with $x_i \in A'_i$.

Being graph-representable implies that there exists an equivalent linear system which is, essentially, the incidence matrix of a graph, so, in particular, is totally unimodular. However, not all the totally unimodular matrices are incidence matrices of directed graphs (see [11, Theorem 2.3.7]). In particular, this shows that not all the matrices with coefficients in \mathbb{Z}_2 are graph-representable.

2.4.1 Systems of equations for Abelian finite groups

Let us now recall the notion of cycle spaces of directed graphs. If H is a directed graph with m arcs, then the *cycle space* of H is the vector space over \mathbb{Q} spanned by the characteristic vectors of cycles of H where the *characteristic vector* of a cycle C of H is the m -dimensional vector v with each coordinate associated with one of the arcs such that the i -th coordinate of v is $+1$ if the i -th arc is traversed by C in its direction, it is -1 if it is traversed by C in the reverse direction, and it is 0 if the arc is not traversed by C .

A set of integer vectors contained in the cycle space is said to *integrally generate* the cycle space of H if they are independent and every vector of the cycle space can be expressed as a linear combination of these vectors with integer coefficients. It is known [53] that the vectors integrally generate the cycle space if and only if every maximum square submatrix of the matrix formed by these vectors has determinant 0 , $+1$ or -1 . This turns out to be equivalent to the fact that a determinant of one such non-singular submatrix is $+1$ or -1 , more precisely Liebchen and Peeters [54] established the following, also see [53]:

Proposition 2.6. *Let H be a connected directed graph with h vertices and m edges. Let A be the $k \times m$ matrix of the characteristic vectors of a set of $k = m - h + 1$ cycles of H . This set of cycles integrally generates the cycle space of H if and only if A contains a $k \times k$ square submatrix with determinant in $\{-1, 1\}$.*

Let us now give some examples. If T is a spanning tree of H , this is, a tree containing all vertices in H . The addition of an edge in $E(H) \setminus E(T)$ creates a single cycle, which is called a fundamental cycle with respect to T . The characteristic vectors of the fundamental cycles with respect to T always integrally generate the cycle space of H [53]. On the other hand, an example of a set of characteristic vectors that generate but not integrally generate the cycle space of a graph is given in Figure 2.1.

Consider now the equation system (2.3). The vector $(\epsilon_{i1}, \dots, \epsilon_{im})$ is referred to as the *characteristic vector* of the i -th equation. The system is said to be *graph-representable* if there exists a directed graph H with m arcs, each associated with one of the variables x_1, \dots, x_m , such that the characteristic vectors of the equations integrally generate the




Figure 2.1 An example of a set of cycles generating but not integrally generating the cycle space of a directed graph: the cycles $v_0v_1v_2v_0v_3v_4v_0$ and $v_0v_1v_2v_0v_4v_3v_0$ generate the cycle space of the depicted directed graph but they do not integrally generate it: the cycle $v_0v_1v_2v_0$ can only be written as a rational (not integral) linear combination of the two cycles in the generating set.

cycle space of H . Such a directed graph H is called a *graph representation* of the equation system (2.3). Note that the condition that the characteristic vectors of the equations integrally generate the cycle space can be efficiently tested since it is equivalent to computing the value of the determinant of a matrix as explained in the previous paragraph.

The proof of Theorem 2.4 follows the lines of the one for Theorem 2.1. In this case we use the following colored version of Lemma 2.3.

Lemma 2.7 (Removal Lemma for arc-colored directed graphs). *Let m be a fixed integer and H a directed graph with its arcs colored with m colors. If a directed graph G with edges colored with m colors contains less than $o(n^h)$ copies of H (the colors of edges in the copy and H must be the same), there exists a set E of at most $o(n^2)$ arcs such that the graph obtained from G by removing the arcs contained in E is H -free.*

Lemma 2.7 can be proved by combining the proof of Lemma 2.3 with the edge-colored version of the Regularity Lemma stated for instance in [48, Lemma 1.18].

Proof. [Proof of Theorem 2.4] Let H be a graph representation of the equation system (2.3). We can assume without loss of generality that H is connected. We view the arc corresponding to the variable x_i as colored with the color i . In this way, the arcs of H are colored with numbers from 1 to m .

Since the dimension of the cycle space of H is k (as the characteristic vectors of the equations from (2.3) are assumed to be independent) and H is comprised of m arcs, the number of the vertices of H is $h = m - k + 1$ (recall that the cycle space of H has dimension $|E(H)| - V(H) + 1$, Biggs [9, Theorem 4.5]).

Next, we construct an auxiliary directed graph K . The vertex set of K is $G \times V(H)$. For every arc (u, v) of H associated with x_i , the directed graph K contains $N|A_i|$ arcs from (g, u) to (ga, v) , one for each $g \in G$ and each $a \in A_i$. The arc from (g, u) to (ga, v) is colored i and labeled by the pair $[a, i]$. The order of H is $N_0 = hN$, its size is $N(|A_1| + \dots + |A_m|)$ and its arcs are colored with numbers $1, \dots, m$. We call K the *blow-up graph* of H by A_1, \dots, A_m .

Let H' be a subgraph of K isomorphic to H (preserving the colors). The arc of H' colored with i is an arc from a vertex (g, u) to a vertex (ga_i, v) for some $a_i \in A_i$. Setting $x_i = a_i$ yields a solution of the system (2.3): indeed, if C is a cycle corresponding to the j -th equation, then the cycle C is also present in H' as a cycle $(g_1, u_1)(g_2, u_2) \dots (g_l, u_l)$. If γ_t is the color of the arc $((g_t, u_t), (g_{t+1}, u_{t+1}))$ (indices taken modulo l), then $a_{\gamma_t} = g_{t+1} - g_t$,

if the arc is traversed in its direction, and $a_{\gamma_t} = g_t - g_{t+1}$ otherwise, and thus

$$0 = \sum_{i=1}^l (g_{i+1} - g_i) = \sum_{i=1}^l \epsilon_{j\gamma_i} a_{\gamma_i} = \epsilon_{j1} a_1 + \dots + \epsilon_{jm} a_m .$$

Note that we can freely rearrange the summands in the above equation as the group G is Abelian.

We have seen that every subgraph of K isomorphic to H corresponds to a solution of the system (2.3). Let us now show that every solution of (2.3) corresponds to N edge disjoint copies of H . Fix a vertex u_0 of H , an element z of G and a solution of the system $a_1 \in A_1, \dots, a_m \in A_m$. Define $\varphi : V(H) \rightarrow G$ such that $\varphi(u_0) = z$ and $\varphi(u') - \varphi(u) = a_i$ for an arc (u, u') of H corresponding to the variable x_i . By the graph representability of the system, the function φ is well defined: if there are two paths from u_0 to a vertex u in H they close a cycle C which can be expressed as an integral linear combination of the cycles in the system. Since the a_i 's form a solution of the system, the sum of the labels on the edges along each of the cycles arising from the system is zero, and therefore this is also the case for C . Since H is connected, the set of vertices $\{(u, \varphi(u)), u \in V(H)\}$ induce a copy of H in K . Since there are N choices for z , and two different choices yield edge-disjoint copies of H , every solution of the system with $a_i \in A_i$ gives rise to N edge-disjoint copies of H .

The proof now proceeds as in Theorem 2.1 except that instead of a cycle of length m we aim to consider copies of the graph H . Fix $\delta_0 > 0$ and apply Lemma 2.7 for $\delta = \delta_0/h^h$ which yields $\delta' > 0$. If there are less than $\delta_0 N^{m-k} = \delta_0 N^{h-1}$ solutions of the system (2.3), the directed graph K contains at most $\delta_0 N^h = \delta N_0^h$ distinct copies of H . By the choice of δ , there is a set E of at most $\delta' N_0^2$ arcs such that $K \setminus E$ has no copy of H .

Let B_i be the set of those elements $a \in A_i$ such that E contains at least N/m arcs $((g, u), (ga, v))$ colored with i . Since $|E| \leq \delta' N_0^2$, the size of each B_i is at most $m|E|/N \leq \delta' m N_0^2/N = \delta' m h^2 N \leq \delta' m^3 N$. Set $A'_i = A_i \setminus B_i$. Since the size of B_i is bounded by $\delta' m^3 N$, and $\delta' \rightarrow 0$ as $\delta_0 \rightarrow 0$, the theorem will be proven after we show that there is no solution of the system (2.3) with $a'_i \in A'_i$.

Assume that there is a solution a'_1, \dots, a'_m of the equation system (2.3) such that $a'_i \in A'_i$ and consider the N disjoint copies of H corresponding to this solution. For every i , the N copies of H contain together N arcs colored with i that are of the form $((g, u), (ga'_i, v))$. Hence, there exists an i_0 such that E contains at least N/m arcs that are colored with i_0 and are of the form $((g, u), (ga'_{i_0}, v))$. Consequently, $a'_{i_0} \in B_{i_0}$ and thus $a'_{i_0} \notin A'_{i_0}$ which violates the choice of the solution. \square

As a final remark, we briefly discuss the condition of graph representability. The key point in the proof of Theorem 2.4 is the correspondence between copies of H and solutions of the system: every copy of H in the constructed graph K yields a solution of the system and every solution gives rise to N edge-disjoint copies of H . This correspondence can be broken if the system is not graph representable in the sense we have defined. For instance, in the example from Figure 2.1, it is possible to express only $2C$ as an integer combination of the base and thus the stated correspondence does not need to exist for groups with elements of order two.

2.4.2 Systems of equations for nonabelian groups

We have already mentioned that, if the characteristic vectors of the equations from a system correspond to fundamental cycles of a graph H with respect to one of its spanning trees, then the equation system is graph-representable. If there exists a representation of this special type, then we say the system is strongly graph representable. More precisely, the system (2.4) is strongly graph representable if there is a directed graph H with m arcs colored by $1, \dots, m$ and a spanning tree T of H such that the fundamental cycles of H with respect to T are cycles C_i defined as follows: C_i is the cycle traversing the arcs of H in the order $e_{i1} \dots e_{im}$ where some of e_{ij} are “empty”, i.e., they do not define an arc of H . If $\epsilon_{i\sigma_i(j)} = +1$, then e_{ij} is the arc colored with $\sigma_i(j)$ traversed in its direction, if $\epsilon_{i\sigma_i(j)} = -1$, then e_{ij} is the arc colored with $\sigma_i(j)$ traversed in the opposite direction, and if $\epsilon_{i\sigma_i(j)} = 0$, then e_{ij} is empty. Note that the condition on the equation system being strongly representable implies that every equation contains a variable that is not in any of the other equations. An example can be found in Figure 2.2 where the graph strongly represents the equation system arising from Corollary 2.8.

This stronger condition suffices to extend Theorem 2.4 to the non-abelian case.

Proof. [Proof of Theorem 2.5] The proof is analogous to the one of Theorem 2.4. Let H be a strong representation of the system. In particular, H has m edges and $h = m - k + 1$ vertices. Let K be the graph with vertex set $G \times V(H)$ that contains an arc $((g, u), (ga, v))$ for each arc (u, v) in H that has color i and each $a \in A_i$. Such an arc has also color i in K .

Let H' be a subgraph of K that is isomorphic to H (preserving the colors of the edges). If $((g, u), (g', v))$ is an arc of H' colored with i , set $x_i = g^{-1}g'$. Observe that $x_i \in A_i$. We claim that x_1, \dots, x_m is a solution of the equation system. Consider the i -th equation and the cycle $(g_0, u_0)(g_1, u_1) \dots (g_t, u_t) = (g_0, u_0)$ of H' corresponding to this equation (note that t can be smaller than m as some arcs can be “empty”); note that this cycle does not need to be consistently directed. We infer from the choice of x_i the following:

$$\prod_{j=1}^m x_{\sigma_i(j)}^{\epsilon_{i\sigma_i(j)}} = \prod_{j=1}^m g_{j-1}^{-1} g_j = g_0^{-1} g_1 g_1^{-1} \dots g_{m-1} g_{m-1}^{-1} g_m = g_0^{-1} g_m = g_0^{-1} g_0 = 1.$$

Hence, x_i 's are indeed a solution of the equation system.

On the other hand, each solution $x_i \in A_i$ gives rise to N edge-disjoint copies of H in K . Indeed, let T be a spanning tree of H such that the cycles C_1, \dots, C_m corresponding to the equations of the system (2.4) are fundamental cycles with respect to T . Root T at an arbitrarily chosen vertex v_0 . Set g_{v_0} to an arbitrary element of G and define the values g_v for other vertices of the graph H as follows: if v' is the parent of v in T and the arc vv' has color i and is oriented from v to v' , then $g_v = g_{v'} x_i^{-1}$; if the arc is oriented from v' to v , then $g_v = g_{v'} x_i$.

Let H' be the subgraph of K with the vertices (g_v, v) that contains the arc from (g_v, v) to $(g_{v'}, v')$ with color i for every arc vv' of H with color i . In order to be sure that H' is properly defined, we have to verify that an arc from (g_v, v) to $(g_{v'}, v')$ with the color i is present in K . If vv' is an arc of T , then K contains the arc from (g_v, v) to $(g_{v'}, v')$ by the definition of g_v . If vv' is not contained in T , there is exactly one equation in the system that contains the variable x_i . We infer by a simple manipulation from the definition of g_v that $g_{v'} = g_v x_i$. Since $x_i \in A_i$, the arc from (g_v, v) to $(g_{v'}, v')$ is contained in H and its color is i . Since the choice of g_{v_0} was arbitrary, K contains N edge-disjoint copies of H .




Figure 2.2 A directed graph representing the equation system (2.5).

The rest of the proof is the same as the last three paragraphs of the proof of Theorem 2.4. \square

2.4.3 Examples of graph-representable systems

Let us present two corollaries of Theorems 2.1–2.5 which illustrates possible applications of our results.

Let G be a finite group and $A, B \subseteq G$. The representation function $r_{A,B} : G \rightarrow \mathbb{N}$, defined as $r_{A,B}(g) = |\{(a, b) \in A \times B : ab = g\}|$, counts the number of representations of an element $g \in G$ as a product of an element in A and one in B . We write r_A for $r_{A,A}$.

Corollary 2.8. *Let G be a finite group of order N and let $A, B, C, D, E \subseteq G$. If*

$$\frac{1}{N} \sum_{g \in E} r_{A,B}(g)r_{C,D}(g) = o(N^2),$$

then it is possible to eliminate $o(N)$ elements in each of the sets to obtain sets A', B', C', D', E' such that

$$\sum_{g \in E'} r_{A',B'}(g)r_{C',D'}(g) = 0.$$

In particular,

1. *If $\frac{1}{N} \sum_{g \in E} r_A^2(g) = o(N^2)$, then $(A')^2 \cap E' = \emptyset$ (A' is E' -product-free).*
2. *If $\frac{1}{N} \sum_{g \in G \setminus A} r_A^2(g) = o(N^2)$, then $|(A')^2| = |A'| + o(N)$ (A' has small doubling).*
3. *If $\frac{1}{N} \sum_{g \in G} r_{A,B}(g)r_{B,A}(g) = o(N^2)$, then $|A'B' \cap B'A'| = o(N)$ (almost all pairs do not commute).*

Proof. Consider the following equation system:

$$\left. \begin{aligned} x_1 x_2 x_4^{-1} x_3^{-1} &= 1 \\ x_1 x_2 x_5^{-1} &= 1 \end{aligned} \right\} \quad (2.5)$$

The system (2.5) is strongly representable by the graph H depicted in Figure 2.2.

The number of solutions of (2.5) with $x_1 \in A$, $x_2 \in B$, $x_3 \in C$, $x_4 \in D$ and $x_5 \in E$ is $\sum_{g \in E} r_{A,B}(g)r_{C,D}(g)$. Hence, if it holds that

$$\frac{1}{N} \sum_{g \in E} r_{A,B}(g)r_{C,D}(g) = o(N^2),$$

then there are $o(N^3)$ solutions of the system (2.5). By Theorem 2.5 applied with $m = 5$, $k = 2$, $A_1 = A, A_2 = B, A_3 = C, A_4 = D$ and $A_5 = E$, it is possible to remove $o(N)$ elements from each of the sets A, \dots, E obtaining sets A', \dots, E' such that the system (2.5) has no solution with $x_1 \in A', \dots, x_5 \in E'$.

Applying the above argument with $A = B = C = D$, we obtain that $\sum_{g \in E'} r_{A'}^2(g) = 0$, which is equivalent to $(A')^2 \cap E' = \emptyset$. This proves 1. Setting $E = G \setminus A$, we get $\sum_{g \in E'} r_{A'}^2(g) = 0$. Since $(A')^2 \subseteq A \cup (E \setminus E')$, $|A \setminus A'| = o(N)$, $|E \setminus E'| = o(N)$, we obtain 2. Similarly, 3 is derived by applying the Corollary for $A = C$ and $B = D$. \square

Removal Lemma for Finite Fields

In this chapter we show the so-called removal lemma for finite fields. This result has been independently proved by Shapira in [68] and by Král', Serra and the author [52].

Theorem 3.1 (Removal Lemma for systems of equations over finite fields; Shapira [68]; Král', Serra, V. [52]). *For all positive integers k and m , $k \leq m$, and every $\varepsilon > 0$, there exists $\delta > 0$ such that the following holds: Let $F = \mathbb{F}_q$ be the finite field of order q and X_1, \dots, X_m be subsets of F , let A be a $(k \times m)$ matrix with coefficients in F and let b be a vector in F^k .*

If there are at most δq^{m-k} solutions of the system $Ax = b$, $x = (x_1, \dots, x_m)$, with $x_i \in X_i$, then there exist sets X'_1, \dots, X'_m with $X'_i \subseteq X_i$ and $|X_i \setminus X'_i| \leq \varepsilon q$ such that there is no solution of the system $Ax = b$ with $x_i \in X'_i$.

Using the little o-notation, Theorem 3.1 asserts that if there are $o(q^{m-k})$ solutions of the system $Ax = b$ with $x_i \in X_i$, then there exist sets $X'_i \subseteq X_i$ such that $|X_i \setminus X'_i| = o(q)$ and there is no solution of the system $Ax = b$ with $x_i \in X'_i$. We will use more precise formulations without the little o-notation, but we occasionally use this notation if no confusion can arise.

By a standard argument Theorem 3.1 implies an analogous result in the integers. In particular it provides a proof of the following result conjectured by Green [41, Conjecture 9.4]:

Theorem 3.2. *Let k and m be integers with $k \leq m$ and let A be an integer $k \times m$ matrix of rank k . For every $\varepsilon > 0$, there exists $\delta > 0$ with the following property. Let $X \subseteq [N]$, and suppose that there are at most δN^{m-k} vectors x in X^m for which $Ax = 0$. Then $X = B \cup C$, such that there are no solutions of the system $Ax = 0$ with $x \in B^m$ and $|C| \leq \varepsilon N$.*

Proof. Let $c(A)$ be twice the sum of the absolute values of the coefficients in A plus 1. Let p be a prime such that $c(A) \cdot N \leq p \leq 2c(A) \cdot N$.

By the choice of p , there is a natural bijective correspondence between the solutions of the linear system $Ax = 0$ in \mathbb{F}_p with $x \in X^m$ and the ones in the integers.

We apply Theorem 3.1 with $F = \mathbb{F}_p$ and $X_i = X$ for all i to obtain the result. \square

A natural application of Theorem 3.2 is the proof of the celebrated theorem of Szemerédi on the existence of k -term arithmetic progressions in sets of integers with positive density [74]. Actually Theorem 3.2 proves the strengthening by Varnavides [78] that a set of integers in $[1, n]$ with positive density contains $\Omega(n^2)$ arithmetic progressions of length

k . This is so because the linear system which defines a k -term arithmetic progression in a set X has $|X|$ trivial solutions (corresponding to constant k -term progressions) which can only be removed by deleting all elements in X . Theorem 3.1 provides the analogous statement in the finite field context.

Corollary 3.3. *For every positive integer k and every $\varepsilon > 0$, there exists $\delta > 0$ such that if a subset X of the elements of the q -element field \mathbb{F}_q contains at most δq^2 arithmetic progressions of length k , then the set X has at most εq elements.*

Corollary 3.3 above can also be proved by using the construction from Frankl and Rödl [30] and the hypergraph removal lemma (see [63] for an explicit construction). Indeed, the fact that Szemerédi Theorem can be shown using the hypergraph removal lemma, was one of the motivation to extend the removal lemma for graphs to hypergraphs.

Our proof of Theorem 3.1 follows the main idea used to show the removal lemma for groups, Theorem 2.1. When the system is reduced to one equation, our construction coincides with the one in the proof of Theorem 2.1, thus it can be viewed as its natural generalization.

Although the techniques used in the proof of the removal lemma for groups to represent a linear equation using a graph, namely a cycle, can be pushed forward to represent some systems of equations, see Theorem 2.4 and Theorem 2.5, the use of graphs seems to impose serious limitations to extend the result for a general system of equations. Instead, the extensions to hypergraphs of the removal lemma, which have been recently proved by Nagle, Rödl, Schacht and Skokan [56, 63], Gowers [37], Tao [77] or Elek and Szegedy [22] seem to be the natural tool to achieve this goal. In particular, and for the same technical conditions explained at the end of Section 2.3, we use the edge-colored version of the hypergraph Removal Lemma, see Theorem 3.4 in Section 3.1. Theorem 3.4 follows from a more general result of Austin and Tao [7, Theorem 1.5].

Independently of us, Conjecture 9.4 from [41] was proved by Shapira [68] (see also [67]) whose method also yields a different proof of Theorem 3.1. Shapira's proof also reduces the problem to finding an appropriate representation of the system by a hypergraph in which one can identify certain subgraphs with solutions, and uses the colored version of the hypergraph Removal Lemma (Theorem 3.4) as our proof does. However, his proof involves $O(m^2)$ -uniform hypergraphs where our proof involves $(k+1)$ -uniform hypergraphs. The two proofs follow a common approach but they differ in the particular ideas used to represent systems by hypergraphs. For a further discussion on the dimension of the edges of the hypergraph used in our construction see Chapter 6.

We note that Theorem 3.1 might also be derived from the main result in Szegedy [73]. There the author proves a Symmetry-preserving Removal Lemma and describes a framework to apply it to Cayley Hypergraphs. Theorem 3.1 would follow from the Symmetry-preserving Removal Lemma once the conditions of validity within this setting are properly verified.

Let us also mention that the conclusion of Theorem 3.1 can be proven in a substantially easier way if we assume that every k columns of the matrix are linearly independent; Král', Serra and the author have reported on this result in [49]. Candela [13] has also proved this result independently of us.

3.1 The hypergraph Removal Lemma and outline of the proof

Our main tool for the proof of Theorem 3.1 is the following version of the hypergraph Removal Lemma which follows from a more general result of Austin and Tao [7, Theorem 1.5].

Theorem 3.4 (Austin and Tao [7]). *Let H be an edge-colored $(k+1)$ -uniform hypergraph with m vertices. For every $\varepsilon > 0$ there exists $\delta > 0$ with the following property.*

Let K be an edge colored $(k+1)$ -uniform hypergraph with M vertices. If the number of copies of H in K (preserving the colors of the edges) is at most δM^m , then there is a set $E' \subseteq E(K)$ of size at most εM^{k+1} such that the hypergraph K' with edge set $E(K) \setminus E'$ is H -free.

The general idea of the proof of Theorem 3.1 is to associate to the linear system $Ax = b$, where A has size $k \times m$, a pair of edge-colored $(k+1)$ -uniform hypergraphs H and K . The hypergraph H has m edges and m vertices, and K is an m -partite hypergraph with m vertices. The edges of H and of K are defined in such a way that there is a correspondence between copies of H in K and solutions of the linear system in $X_1 \times \cdots \times X_m$. More precisely, each solution gives rise to exactly q^k edge disjoint copies of H in K .

The bound on the number of solutions of our linear system translates to the fact that K contains $o(q^m)$ copies of H . At this point we apply the Removal Lemma for hypergraphs, Theorem 3.4, to find a set E' of edges with size $o(q^{k+1})$, such that, by removing E' from K we delete all copies of H .

Since the q^k copies of H corresponding to the same solution are edge-disjoint, a pigeonhole argument allows us to find $o(q)$ elements from each set X_i whose removal eliminates all the solutions of the system of equations.

3.2 Reductions of the system

The key point in our argument is the construction of the auxiliary hypergraphs H and K . Before we explain the details of this construction, we show that we can assume some properties of the given linear system $Ax = b$. In what follows, M^i denotes the i -th column of a matrix M and M_j denotes its j -th row.

Lemma 3.5. *Theorem 3.1 holds if it can be proved under the following assumptions.*

- (i) *The matrix A has the form $A = (I_k | B)$ where I_k is the identity matrix.*
- (ii) *$b = 0$.*
- (iii) *$m \geq k + 2$.*
- (iv) *Every two rows of B are linearly independent.*
- (v) *Each row of B has at least two non-zero entries.*
- (vi) *No column of A is the zero vector.*

Proof. We will establish these properties sequentially and assume the previous ones at each step.

- (i) Observe that, by the nature of the statement of Theorem 3.1, there is no loss of generality in assuming that the matrix A has full rank k . Indeed, choose δ to be the minimum $\delta_{k'}$, $k' = 1, \dots, k$, where $\delta_{k'}$ is the constant for full rank $k' \times m$ matrices. Consider a $k \times m$ matrix A . If the rank k' of the matrix A is smaller than k but the rank of the matrix $(A|b)$ is $k' + 1$, then there is no solution of the system $Ax = b$ at all and there is nothing to prove. Otherwise, let A' be a full-rank $k' \times m$ submatrix of A and b' the subvector b with entries corresponding to the rows of A' . Observe that if the system $Ax = b$ has at most δq^{m-k} solutions, then the system $A'x = b'$ has at most $\delta_{k'} q^{m-k'}$ solutions and the statement follows.

By an appropriate choice of basis, the matrix A can be assumed to be of the form $A = (I_k|B)$, where I_k denotes the $k \times k$ identity matrix.

- (ii) If A is written in the form $(I_k|B)$, then the general statement of Theorem 3.1 follows by applying it to the system $Ax = 0$ once we replace the given first k sets X_1, \dots, X_k by $X_1 - b_1, \dots, X_k - b_k$, where $b = (b_1, \dots, b_k)$ (and leave the remaining sets X_{k+1}, \dots, X_m unchanged.)
- (iii) Note that if $m = k + 1$ then Theorem 3.1 trivially holds with $\delta = \varepsilon$. Indeed, for each element $a \in X_{k+1}$ there is at most one solution to the system $Ax = 0$ with last coordinate a ; since the number of solutions is at most δq , there must be at most $\varepsilon q = \delta q$ elements in X_{k+1} which belong to a solution of $Ax = 0$ with $x \in X_1 \times \dots \times X_{k+1}$; by deleting these elements from X_{k+1} , we delete all the solutions. Thus, we can assume that $m \geq k + 2$.
- (iv) Suppose on the contrary that rows B_i and B_j of B are not linearly independent, say $B_i = \lambda B_j$. This implies that every solution of the system $Ax = 0$ satisfies $x_i = \lambda x_j$. Therefore we can replace X_i by $X_i \cap (\lambda \cdot X_j)$, delete the j -th equation together with the j -th variable, and apply our theorem in the resulting setting: the obtained system contains one less equation and one less variable.
- (v) We may assume that any row B_i of B has at least two non-zero entries. Otherwise the i -th equation would read $x_i + b_{i,j}x_j = 0$ for some $j \in [k + 1, \dots, m]$. As in the preceding paragraph, we can replace the set X_j by $X_j \cap (-b_{i,j}^{-1} \cdot X_i)$ and consider the system obtained by eliminating the i -th equation and the i -th variable.
- (vi) Suppose that A has a zero column, say $A^m = 0$. Set δ to be $\varepsilon \delta'$ where δ' obtained for $k' = k$ and $m' = m - 1$. If the set X_m contains at most εq elements, we can delete all elements of the set X_m and no solution of the system is left. Otherwise, if the system $Ax = b$ has at most δq^{m-k} solutions, then the system $A'x' = b$, where A' is the matrix obtained from A by deleting the m -th column, has at most $\delta q^{m-k} / (\varepsilon q) = \delta' q^{m-1-k}$ solutions and we can apply the statement for $m' = m - 1$ and $k' = k$.

□

3.3 Hypergraph representation and proof of Theorem 3.1

Let $Ax = 0$ be a linear system, where A is a $k \times m$ matrix with entries in F satisfying the properties (i)–(vi) of Lemma 3.5. For the hypergraph representation of the system $Ax = 0$ we shall use an auxiliary matrix associated to the matrix A which is described in Lemma 3.6 below. The support of a vector $x \in F^n$, denoted by $s(x)$, is the set of coordinates with a nonzero entry.

Lemma 3.6. *Let $A = (I_k|B)$ be a $(k \times m)$ -matrix with coefficients in \mathbb{F}_q satisfying the properties (i)–(vi) of Lemma 3.5. There are an $(m \times m)$ matrix C and m pairwise distinct $(k + 1)$ -subsets $S_1, \dots, S_m \subseteq [1, m]$ with the following properties:*

1. $AC = 0$
2. $\text{rank}(C) = m - k$ (maximal under the first condition).

Moreover, there is an ordering of the columns of B such that

3. For every i , $s(C_i) \subseteq S_i$ and $i \in s(C_i)$.
4. For every i , there exists a subset $S'_i \subseteq S_i$ with $|S'_i| = k$ and $S_i \setminus S'_i \subseteq s(C_i)$ such that the set of columns $\{C^j, j \in [1, \dots, m] \setminus S'_i\}$ has rank $m - k$.

The proof of Lemma 3.6 is postponed to Section 3.4. We now proceed to define a suitable hypergraph representation of the linear system which leads to a proof of Theorem 3.1.

Let C be the matrix associated to A and S_1, \dots, S_m be the $(k + 1)$ -subsets of $[1, m]$ satisfying the properties stated in Lemma 3.6.

The hypergraph H is the $(k + 1)$ -uniform edge-colored hypergraph with vertex set $\{1, 2, \dots, m\}$ and with edges S_1, S_2, \dots, S_m , where the edge S_i is colored i .

The hypergraph K is the $(k + 1)$ -uniform m -partite edge-colored hypergraph with vertex set $\mathbb{F}_q \times [1, m]$ and with the following edge set. For every $u \in X_i$, K contains an edge $\{(a_j, j), a_j \in \mathbb{F}_q, j \in S_i\}$ if and only if

$$\sum_{j \in S_i} C_{ij} a_j = u,$$

and this edge is colored by i and labeled by u . Since the support $s(C_i)$ is nonempty and $|S_i| = k + 1$, K contains precisely q^k edges colored by i and labeled by x for each $x \in X_i$.

We next show that the hypergraphs K and H have the needed properties for the proof.

Claim 3.7. *If H' is a copy of H in K , then $x = (x_1, \dots, x_m)$ is a solution of the system, where x_i is the label of the edge colored by i in H' .*

Proof. Since H' is a copy of H , it has m vertices and an edge of each color. By Lemma 3.6 (3) we have $i \in S_i$ for each i which implies $\cup_{i=1}^m S_i = [1, m]$. Hence the vertex set of H' is of the form $\{(a_1, 1), (a_2, 2), \dots, (a_m, m)\}$. By the construction of K , it holds that $Ca = x$ where $a = (a_1, a_2, \dots, a_m)$. Hence, $0 = ACa = Ax$ and x is a solution of the system. \square

Claim 3.8. *For any solution $x = (x_1, \dots, x_m)$ of the system $Ax = 0$ with $x_i \in X_i$, there are precisely q^k edge-disjoint copies of the edge-colored hypergraph H in the hypergraph K .*

Proof. Fix a solution $x = (x_1, \dots, x_m)$ of $Ax = 0$ with $x_i \in X_i$, $1 \leq i \leq m$. First, we will show that there is a copy of H in K in which the edge colored i has label x_i , $1 \leq i \leq m$.

Since the matrix C has rank $m-k$ and satisfies $AC = 0$, the columns in C span the solution space in F^m and thus there is a vector $u = (u_1, \dots, u_m)$ with $x = Cu$. In particular,

$$x_i = C_i \cdot u = \sum_{j=1}^m C_{ij}u_j = \sum_{j \in S_i} C_{ij}u_j,$$

where the second equality follows from Lemma 3.6 (3). Therefore, for every i , the set $\{(u_j, j), j \in S_i\}$ is an edge of K colored i and labeled x_i . It follows that the edges $\{(u_j, j), j \in S_i\}$, $i = 1, \dots, m$, span a copy of H in K .

Since the kernel of C is k -dimensional, there are q^k vectors u satisfying $x = Cu$, and each of them corresponds to a copy of H in K . We next verify that these q^k copies are edge-disjoint.

Let $e = \{(a_j, j), j \in S_i\}$ be an edge of K colored by i and labeled $x_i \in X_i$. We show that all the q^k copies of H in K contain different edges colored by i and labeled x_i for each i . By Lemma 3.6 (4), there is a subset $S'_i \subseteq S_i$ of size k such that $\{C^j, j \notin S'_i\}$ is a set of $m-k$ linearly independent solutions of the system $Ax = 0$. Hence, we may find a vector $u = (u_1, \dots, u_m)$ with $x = Cu$ such that $u_j = a_j$ for each $j \in S'_i$. Moreover, as the element $j \in S_i \setminus S'_i$ is such that $C_{ij} \neq 0$, we must also have $u_j = a_j$ for each $j \in S_i$ and the copy of H associated to this u contains the edge e . Thus, for each edge colored i and labeled x_i there is a copy of H associated to x in K which contains this edge.

Since there are q^k such edges and there is the same number of copies of H associated to the solution x , no two copies can share the same edge colored i and labeled x_i . By applying the same argument to each of the colors $1, \dots, m$, we conclude that the q^k copies of H associated to the solution x are edge-disjoint. \square

We now proceed with the proof of Theorem 3.1.

Proof. [Proof of Theorem 3.1] Let \mathcal{H} be the family of $(k+1)$ -uniform edge colored hypergraphs with m vertices and m edges. Note that \mathcal{H} has a finite number of members. Set $\epsilon' = \epsilon/m$ and, for each $H \in \mathcal{H}$ let δ_H be the quantity obtained from Theorem 3.4 applied to H . Choose δ to be the smallest such δ_H .

Assume that the matrix A and the vector b have the form described in Lemma 3.5, and that the number of solutions of the system $Ax = b$ is at most δq^{m-k} . Let H and K be the hypergraphs constructed in this section. By Claims 3.7 and 3.8, K contains at most $\delta q^m \leq \delta_H q^m$ copies of H . By the Removal Lemma for colored hypergraphs (Theorem 3.4), there is a set E' of edges of K , $|E'| \leq \epsilon' q^{k+1}$ such that, by deleting the edges in E' from K , the resulting hypergraph is H -free.

The sets X'_i are constructed as follows: if E' contains at least q^k/m edges colored with i and labelled with x_i , remove x_i from X_i . In this way, the total number of elements removed from all the sets X_i together is at most $m \cdot |E'|/q^k \leq \epsilon q$. Hence, $|X_i \setminus X'_i| \leq \epsilon q$ as desired. Assume that there is still a solution $x = (x_1, x_2, \dots, x_m)$ with $x_i \in X'_i$. Consider the q^k edge-disjoint copies of H in K corresponding to x . Since each of these q^k copies contains at least one edge from the set E' and the copies are edge-disjoint, E' contains at

least q^k/m edges with the same color i and the same label x_i for some i . However, such x_i should have been removed from X_i . \square

3.4 Proof of Lemma 3.6

In this section, we prove Lemma 3.6 by constructing explicitly a matrix C with the required properties.

We first define a family of auxiliary subsets T_1, \dots, T_m . For each i let T_i be the maximum k -subset of $[i-m+1, i]$ in the lexicographic order such that the set of columns $\{A^j, j \in T_i\}$ (indices taken modulo m) has rank k .

Lemma 3.9. *With indices taken modulo m , the following conditions hold:*

- (i) *For each $i \in [1, m]$ we have $i \notin T_{i-1}$.*
- (ii) *For each $i \in [2, m]$ we have $i \notin T_{i-2}$.*
- (iii) *For each i , the set T_i is obtained by adding i to T_{i-1} and deleting some element in T_{i-1} .*

Proof. Note that the set of columns $\{A^j, j \in [1, m] \setminus \{i\}\}$ span the column space of A . This is clearly so for $k+1 \leq i \leq m$ since A^1, \dots, A^k is the canonical base. On the other hand, for $1 \leq i \leq k$, it follows from Lemma 3.5(v) as every row of B has (at least two) nonzero entries. The maximality of T_{i-1} implies (i).

Similarly, it follows from Lemma 3.5(iv) applied to rows $i-1$ and i with $2 \leq i \leq k$ that the set of columns $\{A^j, j \in [1, m] \setminus \{i-1, i\}\}$ also span the column space of A . The same conclusion follows from Lemma 3.5(v) when $i = k+1$, and it is obvious when $k+2 \leq i \leq m$ since the first k columns of A form the identity matrix. This proves (ii).

By Lemma 3.5(vi) no column of A is the zero vector, so that $i \in T_i$ for each i . It follows from (i) and the maximality of T_i that the symmetric difference $T_i \Delta T_{i-1}$ has cardinality two. \square

We now define the function $g : [1, m] \rightarrow [1, m]$ as $g(i) = T_{i-1} \setminus T_i$ (indices taken modulo m). It follows from Lemma 3.9(iii) that the function g is well defined. Moreover the following holds:

Lemma 3.10. *We have:*

- (i) *The function g is bijective.*
- (ii) *There is an ordering of the columns of B such that g is increasing in $[k+1, m]$.*

Proof. If $g(r) = g(s) = i$ for some distinct r and s then i has been deleted twice in the circular process described in Lemma 3.9 (iii) but inserted only once, a contradiction. This proves (i).

We have $T_k = [1, k]$ for every ordering of the columns A^{k+1}, \dots, A^m . For each $i = k+1, \dots, m$, we may choose A^i to be a column for which the first nonzero coefficient when expressed as a linear combination of the columns in the base corresponding to T_{i-1} occurs more to the left. This choice minimizes the value of $g(i)$ and makes the function g increasing in $[k+1, m]$. \square

We will assume that the last $m - k$ columns of A are ordered in such a way that g is increasing in $[k + 1, m]$, a choice which is possible by Lemma 3.10 (ii).

We can now define the matrix C . The j -th column of C has its support in $T_{j-1} \cup \{j\}$. For $i \in T_{j-1}$, the entry C_{ij} is the coefficient of A^i in the expression of A^j in the base $\{A^i, i \in T_{j-1}\}$:

$$A^j = \sum_{i \in T_{j-1}} C_{ij} A^i,$$

and $C_{jj} = -1$ (recall that, by Lemma 3.9 (i), we have $j \notin T_{j-1}$.)

Clearly, each column of C belongs to the space of solutions of the system $Ax = 0$, so that Lemma 3.6 (1) holds.

Since all the elements of T_i , $i \in [k, m-1]$, are in $[1, i]$, the submatrix of C formed by the last $m - k$ columns and the last $m - k$ rows is an upper triangular matrix with nonzero entries on the diagonal which implies that the rank of C is $m - k$. This proves Lemma 3.6 (2).

By the definition of C the support of column C^j is included in T_j . For $j = k$ we have $T_k = [1, k]$. Since g is increasing in $[k + 1, m]$ and, by Lemma 3.9 (iii), each T_j is obtained from T_{j-1} by adding j and deleting $g(j)$, the support of C^j is included in $[g(j), j]$ if $j \in [k + 1, m]$. For $j \in [1, k]$, Lemma 3.9 (iii) and the maximality of the T_i 's imply that the support of C^j is included in $[1, j] \cup [g(j), m]$.

Let $R \subseteq [1, m] \times [1, m]$ be the area defined by the T_i 's, i.e., $(i, j) \in R$ if and only if either $j \in [1, k]$ and $i \in [1, j] \cup [g(j), m]$ or $j \in [k + 1, m]$ and $i \in [g(j), j]$ (see Figure 3.4 for a typical portrait of R .)




Figure 3.1 An example of the area R in matrix C which corresponds to the permutation $g(1, 2, 3, 4, 5, 6, 7, 8) = (3, 4, 6, 7, 8, 1, 2, 5)$.

We define the family $\{S_1, \dots, S_m\}$ of $(k + 1)$ -subsets of $[1, m]$: S_i is the set of indices j such that $C_{i,j} \in R$. In other words, the sets S_i are obtained by reading off the area R by rows:

$$S_i = \begin{cases} g^{-1}([1, i]) \cup [i, k], & i \in [1, k] \\ g^{-1}(T_i) \cup \{i\}, & i \in [k + 1, m]. \end{cases}$$

By the definition of g , the support of the row C_i is contained in S_i for every $i \in [1, m]$ and none of the rows is zero (the entry in the main diagonal is -1).

Let us show that $|S_i| = k + 1$.

It follows from the definition of g that $g(i) \notin T_i$. Since g is a bijection, S_i has indeed cardinality $k + 1$ for $i \in [k + 1, m]$. On the other hand, we can not have $1 \leq g(j) \leq i \leq k$

for $j \in [i, k]$ since this would imply $T_k \neq [1, k]$, a contradiction. Thus $g^{-1}([1, i])$ and $[i, k]$ are disjoint and S_i has also cardinality $k + 1$ for $i \in [1, k]$.

Let us now show that the sets S_i are pairwise distinct.

Recall that the region R contains in a column $j \in [1, k]$ the rows $[1, j] \cup [g(j), m]$. It follows from Lemma 3.9 (ii) that $j \notin T_{j-2}$ for $j = 2, \dots, k+1$, which implies $g(j-1) > j$. Hence S_j does not contain $j-1$ but it does contain j . On the other hand, the column $j \in [k+1, m]$ contains in the region R the rows $[g(j), j]$, so again S_j contains j but does not contain $j-1$.

Let $j < j'$. If $j' \leq k$ then $\{j' - 1, j'\} \subseteq [j, k] \subseteq S_j$, which implies $S_j \neq S_{j'}$. If $j' > k$ then, either $j' \notin S_j$ or, as g is increasing in $[k+1, m]$, $\{j' - 1, j'\} \subseteq S_j$, which again implies $S_j \neq S_{j'}$.

In order to prove the last part of Lemma 3.6, we show that the columns $\{C^j, j \notin S_i\}$ form a set of $m - k - 1$ linearly independent vectors. Together with Lemma 3.6 (2) and (3), this fact implies Lemma 3.6 (4) and completes the proof of the Lemma.

Let $C' = \{C^j : j \notin S_i\}$ be the submatrix of C formed by the columns with indices not in S_i . We divide this matrix into four parts: the upper left $UL = \{C_{rs} : r < i, s \in [1, i] \setminus S_i\}$ formed by the first $i - 1$ rows of C and the columns with index at most i , the upper right $UR = \{C_{rs} : r < i, s \in [i+1, m] \setminus S_i\}$ formed by the same rows and the remaining columns, the lower right $LR = \{C_{rs} : r \geq i, s \in [1, i] \setminus S_i\}$ formed by the last $m - i + 1$ rows and the columns with index at most i and the lower left $LR = \{C_{rs} : r \geq i, s \in [i+1, m] \setminus S_i\}$ with the remaining entries.

By our construction of the matrix C , UR is an all-zero matrix, while, as discussed in the proof of Lemma 3.6 (2), the columns C^j with $j \in [i+1, m] \setminus S_i$ are linearly independent because the columns $C^j, j \in [k+1, m]$, are linearly independent. On the other hand, again by the construction of C , UL is an upper triangular matrix (maybe with the steps higher than one). It follows that the columns of C' are linearly independent. This completes the proof of Lemma 3.6.

3.4.1 Example of a system and a matrix C

To illustrate the construction, let us present an example. Let $F = \mathbb{F}_7$, $k = 6$, $m = 13$ and let $A = (A_1, \dots, A_{13})$ be the matrix

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 1 & 2 & 0 & 1 & 2 & 4 & 2 \\ 0 & 1 & 0 & 0 & 0 & 0 & 3 & 4 & 0 & 1 & 4 & 6 & 3 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 2 & 2 & 5 & 2 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 2 & 5 & 1 & 5 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 2 & 1 & 4 & 1 & 2 & 4 & 2 \\ 0 & 0 & 0 & 0 & 0 & 1 & 3 & 5 & 5 & 4 & 1 & 2 & 1 \end{pmatrix}.$$

We can observe that A already fulfills Lemma 3.5. Moreover, the order in the columns of the part B of the matrix A is such that maximizes the sets T_i . The matrix C associated with A is

$$\begin{pmatrix} -1 & \mathbf{2} & \mathbf{6} & \mathbf{0} & \mathbf{3} & \mathbf{5} & \mathbf{1} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & \mathbf{0} & \mathbf{4} & \mathbf{1} & \mathbf{4} & \mathbf{3} & -2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & \mathbf{2} & \mathbf{3} & \mathbf{6} & \mathbf{0} & \mathbf{1} & \mathbf{2} & 0 & 0 & 0 & 0 \\ \mathbf{6} & 0 & 0 & -1 & \mathbf{2} & \mathbf{0} & \mathbf{0} & \mathbf{2} & \mathbf{5} & 0 & 0 & 0 & 0 \\ \mathbf{2} & \mathbf{4} & 0 & 0 & -1 & \mathbf{5} & \mathbf{2} & -3 & \mathbf{4} & 0 & 0 & 0 & 0 \\ \mathbf{0} & \mathbf{0} & 0 & 0 & 0 & -1 & \mathbf{3} & -1 & \mathbf{5} & \mathbf{3} & 0 & 0 & 0 \\ \mathbf{4} & \mathbf{0} & \mathbf{2} & 0 & 0 & 0 & -1 & \mathbf{2} & \mathbf{0} & \mathbf{6} & 0 & 0 & 0 \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & 0 & 0 & 0 & 0 & -1 & \mathbf{0} & \mathbf{1} & \mathbf{1} & 0 & 0 \\ \mathbf{0} & \mathbf{3} & \mathbf{1} & 0 & \mathbf{1} & 0 & 0 & 0 & -1 & \mathbf{4} & \mathbf{2} & 0 & 0 \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & 0 & \mathbf{0} & 0 & 0 & 0 & 0 & -1 & \mathbf{0} & \mathbf{2} & 0 \\ \mathbf{3} & \mathbf{4} & \mathbf{6} & \mathbf{3} & \mathbf{0} & 0 & 0 & 0 & 0 & 0 & -1 & \mathbf{1} & 0 \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & 0 & 0 & 0 & 0 & 0 & 0 & -1 & \mathbf{4} \\ \mathbf{6} & \mathbf{2} & \mathbf{4} & \mathbf{4} & \mathbf{2} & \mathbf{1} & 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix}$$

The darker part of the matrix C (the elements in bold face) corresponds to the region R .

The sets T_i and the values of g are the following:

$$\begin{aligned} T_6 &= [1, 2, 3, 4, 5, 6], g(7) = 1 & T_1 &= [5, 7, 9, 11, 13, 1], g(2) = 5 \\ T_7 &= [2, 3, 4, 5, 6, 7], g(8) = 2 & T_2 &= [7, 9, 11, 13, 1, 2], g(3) = 7 \\ T_8 &= [3, 4, 5, 6, 7, 8], g(9) = 3 & T_3 &= [9, 11, 13, 1, 2, 3], g(4) = 11 \\ T_9 &= [4, 5, 6, 7, 8, 9], g(10) = 6 & T_4 &= [9, 13, 1, 2, 3, 4], g(5) = 9 \\ T_{10} &= [4, 5, 7, 8, 9, 10], g(11) = 8 & T_5 &= [13, 1, 2, 3, 4, 5], g(6) = 13 \\ T_{11} &= [4, 5, 7, 9, 10, 11], g(12) = 10 \\ T_{12} &= [4, 5, 7, 9, 11, 12], g(13) = 12 \\ T_{13} &= [4, 5, 7, 9, 11, 13], g(1) = 4 \end{aligned}$$

The sets S_i for the matrix C are:

$$\begin{aligned} S_1 &= [1, 2, 3, 4, 5, 6, 7] & S_8 &= [1, 2, 3, 8, 9, 10, 11] \\ S_2 &= [2, 3, 4, 5, 6, 7, 8] & S_9 &= [1, 2, 3, 5, 9, 10, 11] \\ S_3 &= [3, 4, 5, 6, 7, 8, 9] & S_{10} &= [1, 2, 3, 5, 10, 11, 12] \\ S_4 &= [1, 4, 5, 6, 7, 8, 9] & S_{11} &= [1, 2, 3, 4, 5, 11, 12] \\ S_5 &= [1, 2, 5, 6, 7, 8, 9] & S_{12} &= [1, 2, 3, 4, 5, 12, 13] \\ S_6 &= [1, 2, 6, 7, 8, 9, 10] & S_{13} &= [1, 2, 3, 4, 5, 6, 13] \\ S_7 &= [1, 2, 3, 7, 8, 9, 10] \end{aligned}$$

We can check that all the properties from Lemma 3.6 are fulfilled.

Removal Lemma for Abelian Groups

In Chapter 2, we showed the removal lemma for groups, valid for one equation and any finite group. In Chapter 3, we proved a removal lemma for systems of linear equations in finite fields. In this chapter, we prove a removal lemma for integer linear systems over finite Abelian groups, Theorem 4.1. This result attempts to answer a question by Shapira [68] regarding an extension of the removal lemma for finite fields to Abelian groups.

4.1 Removal Lemma for linear systems over finite Abelian groups

Before proceeding to the result, recall that the k -th determinantal divisor $d_k(A)$ of an integer matrix A is the greatest common divisor of the determinants of all the $k \times k$ submatrices of A [57]. For simplicity, we use the shorter term *k -th determinantal* instead of *k -th determinantal divisor*. Our main result is the following:

Theorem 4.1 (Removal lemma for integer linear systems over Abelian groups; Král', Serra, V. [51]). *Let A be an integer $(k \times m)$ matrix, $m \geq k$. For any $\epsilon > 0$ there exists $\delta(\epsilon, A) > 0$ such that the following holds.*

For every Abelian group G of order n coprime with $d_k(A)$, for every family of subsets X_1, \dots, X_m of G and for every vector $b \in G^k$, if the linear system $Ax = b$ has at most δn^{m-k} solutions with $x_1 \in X_1, \dots, x_m \in X_m$ then there are sets $X'_1 \subset X_1, \dots, X'_m \subset X_m$ with $|X'_i| \leq \epsilon n$, for all i , such that there is no solution of the system with $x_1 \in X_1 \setminus X'_1, \dots, x_m \in X_m \setminus X'_m$.

In the little ‘o’ notation, Theorem 4.1 states that, if an integer linear system over an Abelian group of order n (with the condition that the determinantal of the matrix is coprime with the order of the group), has $o(n^{m-k})$ solutions, then we can destroy all the solutions by removing $o(n)$ elements in each set.

Let us remark that the condition over the determinantal $d_k(A)$ in the statement of Theorem 4.1 indicates that the system has, in total, $|G|^{m-k}$ solutions. It also tells us that the system behaves in a similar way to the finite field case. For a further discussion on the structure of the system when the coprimality condition is removed see Section 4.5. In Section 6.5 there can be found several cases for which the condition can be removed, as well as some thoughts about a general result.

A general framework for the study of this type of results is discussed by Szegedy [73]. The author proves a Symmetry-preserving removal lemma and applies it to give a diagonal version of the Szemerédi Theorem on arithmetic progressions in Abelian groups. The arguments presented to show Theorem 4.1 follow the ones from Chapter 2 and 3 and the result provides a general answer for linear systems $Ax = b$ with $d_k(A) = 1$, which includes the case of arithmetic progressions [73, Theorem 3]. Let us notice that the relation between δ and ϵ will be worse using our result than the straight construction found in [73], the main reason being the use of higher-dimensional hypergraphs.

As in Theorem 3.1, the proof of Theorem 4.1 uses the Removal Lemma for colored hypergraphs. This result can be deduced from Austin and Tao's [7, Theorem 1.5] and is stated also in Ishigami's [46]. This result appears in Chapter 3 as Theorem 3.4.

Theorem 4.2. *For every positive integers $m \geq k \geq 2$ and every $\epsilon > 0$ there is a $\delta > 0$ depending on m, k and ϵ such that the following holds.*

Let H and K be colored k -uniform hypergraphs with $m = |V(H)|$ and $M = |V(K)|$ vertices respectively. If the number of copies of H in K (preserving the colors of the edges) is at most δM^m , then there is a set $E' \subseteq E(K)$ of size at most ϵM^k such that the hypergraph K' with edge set $E(K) \setminus E'$ is H -free.

4.2 Circular Unimodular Matrices

In this section we prove Theorem 4.1 in the particular case of homogeneous linear systems with what we call standard circular unimodular matrices, which enjoy some useful particular properties. We show in Section 4.3 how the statement extends to the general case.

Throughout this chapter, and similar as in the other chapters, A_i denotes the i -th row of a matrix A and A^j its j -th column. Recall that a square integer matrix is unimodular if it has determinant ± 1 .

We say that a $(k \times m)$ integer matrix is *standard circular unimodular* if the following properties hold:

- (U1) $A = (I_k | B)$, where I_k denotes the identity matrix of order k .
- (U2) For each $j = 1, \dots, m$, the determinant formed by k consecutive columns in the circular order, $\{A^{j+1}, A^{j+2}, \dots, A^{j+k}\}$ is ± 1 , where the superscripts are taken modulo m .

We simply call matrices satisfying property U2 *circular unimodular*. Note that property U1 can always be imposed to a circular unimodular matrix by using elementary matrix transformations. The next key lemma proves Theorem 4.1 for circular unimodular matrices by constructing an hypergraph associated to a given linear system. The approach is similar to the one by Candela [13] and by Král', Serra and the author [52].

Lemma 4.3. *Let A be a $(k \times m)$ circular unimodular matrix with $m \geq k + 2$. For each $\epsilon > 0$ there is a $\delta(\epsilon, A) > 0$ such that the following holds.*

For every Abelian group G of order n and every collection of subsets $X_1, \dots, X_m \subset G$, if the number of solutions of the system $Ax = 0$ with $x \in \prod_{i=1}^m X_i$ is at most δn^{m-k} , then

there are subsets $X'_i \subset X_i$ with $|X'_i| < \epsilon n$ for all i such that there is no solution of the system $Ax = 0$ with $x \in \prod_{i=1}^m (X_i \setminus X'_i)$.

Moreover, if we have $X_j = G$, for $j \in I$, where $I \subset \{1, \dots, m\}$ has cardinality $|I| \leq k$, then we can choose the sets X'_i in such a way that $X'_j = \emptyset$ for each $j \in I$.

Proof. We start by defining an integer $(m \times m)$ matrix C from which we will construct a pair of colored hypergraphs H and K . The purpose of this construction is to establish a correspondence between solutions of the system $Ax = 0$ with copies of H in K .

By property U2, the j -th column of A can be written, for every j , as an integer linear combination of the preceding k columns in the circular ordering:

$$A^j = \sum_{i=j-k}^{j-1} C_{i,j} A^i,$$

where the superscript i is taken modulo m .

For $j = 1, 2, \dots, m$ we let $C_{j,j} = -1$ and, if i does not belong to the circular interval $[j-k, j]$, then we set $C_{i,j} = 0$. Thus,

$$\sum_i C_{i,j} A^i = 0, \quad j = 1, 2, \dots, m. \quad (4.1)$$

Notice that, since all the determinants of k consecutive columns of A in the circular ordering are ± 1 , the coefficients of C are integers (apply the Cramer's rule to solve the corresponding linear systems). By the same reason, we have

$$C_{j-k,j} = \pm 1,$$

since the determinants of the matrices formed by the columns A^{j-k+1}, \dots, A^j and by the columns A^{j-k}, \dots, A^{j-1} are both ± 1 .

The integer $(m \times m)$ matrix $C = (C_{i,j})$ will be used to define our hypergraph model for the given linear system.

Let H be a $(k+1)$ -uniform colored hypergraph with m vertices labelled $\{1, 2, \dots, m\}$. The edges of H are the m "cyclic" $(k+1)$ -subsets

$$\{1, \dots, k+1\}, \{2, \dots, k+2\}, \dots, \{m, 1, \dots, k\},$$

(entries taken modulo m). The i -th edge $\{i, i+1, \dots, i+k\}$ is colored with color i . Since $m \geq k+2$, H contains m different edges of mutually different colors.

Let K be a $(k+1)$ -uniform colored hypergraph with vertex set $G \times [1, m]$. For each element $a_i \in X_i$, the $(k+1)$ -subset $\{(g_i, i), \dots, (g_{i+k}, i+k)\}$ form an edge labelled a_i and colored with color i if

$$a_i = \sum_{j=i}^{i+k} C_{i,j} g_j. \quad (4.2)$$

Thus the edges of K bear both, a color and a label. Note that, for each fixed $a_i \in X_i$, the system (4.2) has n^k solutions. Indeed, since $C_{i,i} = -1$, we can fix arbitrary values g_{i+1}, \dots, g_{i+k} and get a value for g_i satisfying the equation. Therefore each element $a_i \in X_i$ gives rise to n^k edges colored i and labeled a_i .

We next show that each solution to $Ax = 0$ creates n^k edge-disjoint copies of the hypergraph H inside K and, also, that each copy of H inside K comes from a solution of the system $Ax = 0$.

Claim 4.4. *If H' is a copy of H in K , then $x = (x_1, \dots, x_m)$ is a solution of the system, where x_i is the label of the edge colored by i in H' .*

Proof. The copy H' has an edge of each color and is supported over m vertices. Indeed, since the edge colored i contains a vertex in $G \times \{i\}$, then the copy H' has one vertex on each $G \times \{i\}$, $1 \leq i \leq m$. Hence the vertex set of H' is of the form $\{(g_1, 1), (g_2, 2), \dots, (g_m, m)\}$ for some $g_1, \dots, g_m \in G$. If the edge $((g_i, i), \dots, (g_{i+k}, i+k))$ colored i in H' has label x_i then, by the construction of K , we have $x_i = \sum_s C_{i,s} g_s$. Therefore, it holds that $Cg = x$ where $g = (g_1, g_2, \dots, g_m)$. Hence, as all the columns in C are in the kernel of A , we have $0 = ACg = Ax$ and x is a solution of the system. \square

Claim 4.5. *For any solution $\alpha = (\alpha_1, \dots, \alpha_m)$ of the system $Ax = 0$ with $\alpha_i \in X_i$, there are precisely n^k edge-disjoint copies of the edge-colored hypergraph H in the hypergraph K with edges labelled with $\alpha_1, \dots, \alpha_m$.*

Proof. Fix a solution $\alpha = (\alpha_1, \dots, \alpha_m)$ of $Ax = 0$ with $\alpha_i \in X_i$, $1 \leq i \leq m$.

Observe that, by property U2, α is uniquely determined by any of its subsequences $(\alpha_i, \alpha_{i+1}, \dots, \alpha_{i+m-k-1})$ of $m-k$ consecutive coordinates in the circular ordering.

By the construction of the matrix C , its i -th row C_i has an entry -1 in the i -th column and has its support contained in columns $C^i, C^{i+1}, \dots, C^{i+k}$ (where the superscripts are taken modulo m .) Therefore, the $m-k$ columns of C with indices in $[1, m] \setminus [i+1, \dots, i+k]$ contain a unique nonzero entry in the i -th row, is located in the main diagonal and have a value of -1 .

With the previous remark in mind, we observe that, for every choice of a vector $(g_{i+1}, \dots, g_{i+k}) \in G^k$ (subscripts modulo m), there is a unique vector $(g_{i+k+1}, \dots, g_{i-1}, g_i) \in G^{m-k}$ which satisfies the system $Cg = \alpha$, where $\alpha = (\alpha_1, \dots, \alpha_m)$ with $\alpha_i \in X_i$ is the solution of the system $Ax = 0$ we have fixed from the beginning and $g = (g_1, g_2, \dots, g_m)$. Indeed, for each t , once the values $(g_{i+1-t}, g_{i+2-t}, \dots, g_{i+k-t})$ have been found, we can determine g_{i-t} from the equation

$$\alpha_{i-t} = \sum_{s=i-t}^{i+k-t} C_{i-t,s} g_s, \quad (4.3)$$

since $C_{i-t,i-t} = -1$. In this way, starting with the vector $(g_{i+1}, \dots, g_{i+k-1}, g_{i+k}) \in G^k$ and $m-k$ consecutive elements of α , $\{\alpha_{i+k+1}, \dots, \alpha_{i-1}, \alpha_i\}$, we find a unique m -dimensional vector $g = (g_1, \dots, g_m)$. Observe that, if we let $\beta = Cg \in G^m$, then β satisfies $A\beta = A(Cg) = (AC)g = 0g = 0$. Therefore β is a solution of the system $Ax = 0$ which shares $m-k$ consecutive values with the given solution α , hence $\beta = \alpha$. It follows that the equations (4.3) hold for all t . Since these are the defining equations (4.2) for the $(k+1)$ -tuple $(g_i, i), \dots, (g_{i+k}, i+k)$ to be an edge of K colored i and labeled x_i , we conclude that each vector $(g_{i+1}, \dots, g_{i+k}) \in G^k$ defines uniquely a copy of H in K . Hence the solution α induces n^k copies of H in K .

Recall that each entry $\alpha_i \in X_i$ of α gives rise to n^k edges labeled α_i in the hypergraph K . On the other hand, each of these edges belong to a unique copy of H inside K related to the solution α . Since this holds for each of the edges and for each α_i , $1 \leq i \leq m$, we conclude that the n^k copies of H with edges labelled with $\alpha_1, \dots, \alpha_m$ are edge-disjoint. \square

Claims 4.4 and 4.5 show that there is a correspondence between the solutions of the system $Ax = 0$, with $x_i \in X_i$ for each i , and the copies of H inside K . More precisely, each solution appears in the ordered labels of n^k copies of H in K , and the labels of each copy of H in K form a solution. Recall that the copies should respect the coloring.

We now proceed with the proof of Lemma 4.3. Given $\epsilon > 0$ let $\delta > 0$ be the value given by the Removal Lemma of colored hypergraphs (Theorem 4.2) for the positive integers $m, k+1$ and $\epsilon' = \epsilon/m > 0$. If the number of solutions of the system $Ax = 0$ is at most δn^{m-k} , it follows from Claims 4.4 and 4.5, that K contains δn^m copies of H . By Theorem 4.2, there is a set E' of edges of K with size $\epsilon' n^{k+1}$ such that, by deleting the edges in E' from K , the resulting hypergraph is H -free.

The subsets $X'_i \subset X_i$ of removed elements are constructed as follows: if E' contains at least n^k/m edges colored with i and labeled with x_i , we remove x_i from X_i (that is, $x_i \in X'_i$.) In this way, the total number of elements removed from all the sets X_i together is at most $m\epsilon'n = \epsilon n$. Hence, $|X'_i| \leq \epsilon n$ as desired. Suppose that there is still a solution $x = (x_1, x_2, \dots, x_m)$ with $x_i \in X_i \setminus X'_i$. Consider the n^k edge-disjoint copies of H in K corresponding to x . Since each of these n^k copies contains at least one edge from the set E' and the copies are edge-disjoint, E' contains at least n^k/m edges with the same color i and the same label x_i for some i . However, such x_i should have been removed from X_i , a contradiction.

It remains to show the last part of Lemma 4.3. Let I be a subset of $[1, m]$ with $|I| \leq k$, and suppose that $X_j = G$ for each $j \in I$. Let L be the subgraph of H formed by all the edges in H except the ones colored with $i \in I$. Note that H contains a single copy of L . Since every vertex of H belongs to $(k+1)$ edges, the subgraph L has no isolated vertices. It follows that a copy L' of L in K has precisely one vertex in $G \times \{i\}$ for each $i = 1, 2, \dots, m$. By the construction of K , there is at most one copy H' of H in K containing L' , namely the one whose labels are given by equation (4.2) given the g_i 's. Since $X_j = G$ for each $j \in I$, then the label of each missing edge in L' , given by this equation, belongs to the corresponding set X_j , thus such an edge is indeed present in K . Hence, every copy of L in K can be uniquely extended to a copy of H . Thus, K contains as many copies of H as of L . We can apply Theorem 4.2 to L in the above argument to remove all copies of L by removing only elements from sets X_i with $i \in \{1, \dots, m\} \setminus I$. This completes the proof. \square

The condition $m \geq k+2$ in the hypothesis of Lemma 4.3 has been used in the proof for the construction of the hypergraphs associated to the linear system. However, this condition is not restrictive for the proof of Theorem 4.1; in the remaining cases (when m is k or $k+1$), we apply the following lemma:

Lemma 4.6. *Let $A = (I_k|B)$ be a $(k \times m)$ integer matrix. If $m = \{k, k+1\}$ then the statement of Theorem 4.1 holds for A .*

Proof. For $m = k$ the system has a unique solution and there is nothing to prove. Suppose that $m = k+1$. Then, for each element $\alpha \in X_{k+1}$ there is at most one solution to the system $Ax = b$ with last coordinate $x_{k+1} = \alpha$. Let X'_{k+1} be the set of elements $\alpha \in X_{k+1}$ such that $x_{k+1} = \alpha$ is the last coordinate of some solution x . Since there are at most δn solutions we have $|X'_{k+1}| \leq \delta n$ and we are done by removing the set X'_{k+1} . Thus the statement of Theorem 4.1 holds with $\delta = \epsilon$. \square

4.3 A reduction lemma

In this section we prove some technical lemmas that will allow us to derive Theorem 4.1 from Lemma 4.3 via a series of transformations to the given linear system. We have devoted Section 4.6 to an example for clarification of the construction presented in this section.

Recall that the adjugate matrix of L , denoted by $\text{adj}(L)$, is the matrix C with $C_{i,j} = (-1)^{i+j} M_{j,i}(L)$, where $M_{j,i}(L)$ is the determinant of the matrix L with the row j and the column i deleted.

Throughout the section G denotes an Abelian finite group of order n . For an integer a coprime with the order n of G the map $g \mapsto ag$ is an automorphism of the group. We will also denote by a this automorphism and by a^{-1} its inverse. Observe that if an $(r \times r)$ integer matrix L has determinant $a = \det L$ coprime with n then the action $x \mapsto Lx$ of L on G^r is invertible with $L^{-1}x = a^{-1}(\text{adj}(L)x)$. Thus the linear system $Lx = b$ has the unique solution $x = L^{-1}b$. By abuse of notation, in what follows we write $L^{-1}b$ and, for a matrix M with appropriate dimensions, $L^{-1}M$, in the sense that division by a means the action of the automorphism a^{-1} .

Definition 4.7 (Restricted system). *A restricted system is a triple $\{A, b, \mathcal{X}\}$ where*

- $\mathcal{X} = X_1 \times X_2 \times \cdots \times X_m$ is an m -tuple of subsets of G .
- A denote a $(k \times m)$ integer matrix such that its k -th determinantal $d_k(A)$ satisfies $\gcd(d_k(A), |G|) = 1$.
- b is an element of G^k , and we usually refer to it as the independent vector.

A solution of the restricted system $\{A, b, \mathcal{X}\}$ is a vector $x = (x_1, \dots, x_m) \in G^m$ such that $Ax = b$ and $x_i \in X_i$, $i = 1, 2, \dots, m$.

Definition 4.8 (Extension of a restricted system). *A restricted system $\{A', b', \mathcal{Y}\}$ is an extension of the restricted system $\{A, b, \mathcal{X}\}$ if the following two conditions hold:*

- E1: $k' \geq k$, $m' \geq m$, $m' - k' = m - k$, and
- E2: *There is a subset $I_0 \subset [1, m']$ with cardinality $|I_0| = m$ a bijection $\sigma : I_0 \rightarrow [1, m]$ and maps $\phi_i : Y_i \rightarrow X_{\sigma(i)}$ such that the map $\phi : \mathcal{Y} \rightarrow \mathcal{X}$ with $(\phi(y))_i = \phi_{\sigma^{-1}(i)}(y_{\sigma^{-1}(i)})$ induces a bijection between the set of solutions of $\{A', b', \mathcal{Y}\}$ and the set of solutions of $\{A, b, \mathcal{X}\}$. Moreover, for each $i \in [1, m'] \setminus I_0$, we have $Y_i = G$.*

Thus, an extension $\{A', b', \mathcal{Y}\}$ of $\{A, b, \mathcal{X}\}$ has the same number of solutions and one can define a map ϕ with the following property. Let $\mathcal{Y} \setminus \mathcal{Y}'$ stand for $\prod_{i=1}^{m'} Y_i \setminus Y'_i$ and assume $Y'_j = \emptyset$ when $j \notin I_0$. If $\{A', b', \mathcal{Y} \setminus \mathcal{Y}'\}$ has no solutions, then $\{A, b, \mathcal{X} \setminus \phi(\mathcal{Y}')\}$ has no solutions either ($\mathcal{X} \setminus \phi(\mathcal{Y}')$ refers to $\prod_{i=1}^m X_i \setminus \phi_{\sigma^{-1}(i)}(Y'_{\sigma^{-1}(i)})$).

When $\{A', b', \mathcal{Y}\}$ is an extension of $\{A, b, \mathcal{X}\}$ with $k = k'$, any bijection for σ , and the ϕ_i 's are bijective for each i , we say that the two systems are equivalent.

The purpose of this section is to show that any restricted system which fulfills the hypothesis of Theorem 4.1 can be extended to an homogeneous one with a circular unimodular matrix. This will lead to a proof of Theorem 4.1 from Lemma 4.3.

We first show that the matrix A can be enlarged to an integer square matrix M of order m such that $\det(M) = d_k(A)$. The following Lemma uses the ideas of Zhan [81] and Fang [26] to extend partial integral matrices to unimodular ones. We include the proof of the simpler version we need for our purposes.

Lemma 4.9 (Matrix extension). *Let M be an $r \times s$ integer matrix, $s \geq r$. Let d_M denote the greatest common divisor of the determinants of the $\binom{s}{r}$ square $(r \times r)$ submatrices of M .*

There is an $s \times s$ integer matrix \overline{M} such that

- (i) \overline{M} contains M in its r first rows, and
- (ii) $\det(\overline{M}) = d_M$.

Proof. Let $S = U^{-1}MV^{-1}$ be the Smith Normal Form of M , where U and V are unimodular matrices. We have $S = (D|0)$, where D is an $(r \times r)$ diagonal integer matrix with $|\det(D)| = |d_M|$ and 0 is an all-zero $(r \times (s - r))$ matrix.

Recall that U and V are the row and column operations respectively which transform M into S . Observe that the row operations do not modify the value of the determinant of any $(r \times r)$ square submatrix of M . The column operations may modify individual determinants but do not change the value of d_M .

Let \overline{S} be the matrix:

$$\overline{S} = \begin{pmatrix} D & 0 \\ 0 & I_{s-r} \end{pmatrix},$$

where I_k denotes the identity matrix of order k . We have $\det(\overline{S}) = \det(D) = d_M$.

Then, if we let $\overline{V} = V$ and

$$\overline{U} = \begin{pmatrix} U & 0 \\ 0 & I_{s-r} \end{pmatrix},$$

we obtain the matrix

$$\overline{M} = \overline{U} \overline{S} \overline{V}$$

which clearly fulfills (i) as it contains M as a submatrix in its first r rows, and it also satisfies (ii) since $\det(\overline{M}) = \det(\overline{S}) = d_M$, as \overline{U} and \overline{V} are still unimodular. \square

We say that the restricted system $\{A, b, \mathcal{X}\}$ is *thin* if the set of solutions is a subset of $X_1 \times \cdots \times X_{j-1} \times \{\gamma_j\} \times X_{j+1} \times \cdots \times X_m$, for some j and $\gamma_j \in X_j$. Note that the statement of Theorem 4.1 is obvious if the system is thin since it suffices to delete the element γ_j to remove all solutions. Thus there is no loss of generality in assuming that our restricted system is not thin.

Lemma 4.10. *The restricted system $\{A, b, \mathcal{X}\}$ is either thin or it has an extension $\{A', b', \mathcal{Y}\}$ such that*

- (i) $k' = m$ and $m' = 2m - k$;
- (ii) the matrix A' has the form $A' = (I_{k'}|B)$;
- (iii) $b' = 0$;
- (iv) $\gcd(B_i) = 1$, where B_i denotes the i -row of the submatrix B and
- (v) $\max_{i,j} \{|A'_{i,j}|\}$ depends on the entries of A but not on the group G .
- (vi) for every $k' < j \leq m'$, the restricting set Y_j is the whole group G .

Proof. By using Lemma 4.9 we extend the matrix A into an $m \times m$ square matrix

$$M = \begin{pmatrix} A \\ E \end{pmatrix}$$

with determinant $\det(M) = d_k(A)$. We complete the square matrix M to the $m \times (2m - k)$ matrix

$$M' = \begin{pmatrix} A & 0 \\ E & I_{m-k} \end{pmatrix} = (M|B').$$

We now consider the restricted system $\{M', b', \mathcal{X}'\}$ where $b' = (b, 0)$ is obtained from b by adding zeros in the last $m - k$ coordinates and

$$X'_i = \begin{cases} X_i, & 1 \leq i \leq m; \\ G, & m + 1 \leq i \leq 2m - k. \end{cases}$$

By letting $I_0 = [1, m]$ and σ and ϕ_i be the identity maps we see that ϕ induce a bijection between the solutions to $\{M', b', \mathcal{X}'\}$ and the solutions of $\{A, b, \mathcal{X}\}$. In particular, if y is a solution of $\{M', b', \mathcal{X}'\}$ then $x = \phi(y)$ is a solution of $\{A, b, \mathcal{X}\}$. Moreover, for any solution x of $\{A, b, \mathcal{X}\}$, there exists a unique $y' \in G^{2m-k}$ solution of $\{M', b', \mathcal{X}'\}$ such that $x = \phi(y')$. Therefore $\{M', b', \mathcal{X}'\}$ is an extension of the original system.

Let $U = \text{adj}(M)$ denote the adjugate of M . Since $a = d_k(A)$ is relatively prime with n , we get an equivalent restricted system $\{M'', b'', \mathcal{X}''\}$ by setting

$$M'' = (UM|UB') = (a \cdot I_m|UB'), \quad b'' = Ub'$$

and, by replacing each X'_i , for $i \in [1, m]$, by $\bar{X}_i'' = a^{-1}X'_i$ and $\bar{X}_i'' = X'_i$, for $i \in [m + 1, 2m - k]$, we get a an equivalent system of the form $\{(I_m|B''), b'', \bar{\mathcal{X}}''\}$ where $B'' = UB'$. The system is equivalent since the matrix U is invertible in G .

At this point we can replace b'' with the zero vector by letting $X_i'' = \bar{X}_i'' - b_i''$ for $i = 1, \dots, m$ and leaving the other sets untouched. The solutions of the homogeneous system $(I_m|B'')x = 0$ with $x_i \in X_i''$ are in bijective correspondence with the solutions of $M''x = b''$ with $x_i \in \bar{X}_i''$. So $\{(I_m|B''), 0, \mathcal{X}''\}$ is a system equivalent to $\{(I_m|B''), b'', \bar{\mathcal{X}}''\}$, which fulfills conditions (i)-(iii) of the Lemma.

We observe that, if $B_j'' = 0$ for some j , then the j -th equation implies $x_j = 0$. Thus, the solution set of $\{(I_m|B''), 0, \mathcal{X}''\}$ is inside $X_1'' \times \dots \times X_{j-1}'' \times \{0\} \times X_{j+1}'' \times \dots \times X_m''$, which implies that the solution set for the original system is inside $X_1 \times \dots \times X_{j-1} \times \{\gamma_{j'}\} \times X_{j'+1} \times \dots \times X_m$, for some $\gamma_{j'} \in X_{j'}$. Thus, if $B_j'' = 0$, then the system is thin. Therefore we can assume that all the rows in B'' are non-zero.

Suppose that $\text{gcd}(B_i'') = s > 1$, where B_i'' denotes the i -th row of B'' . Then the i -th coordinate y_i , $i \in [1, m]$, of a solution of $(I_m|B'')y = 0$ belongs to the subgroup $s \cdot G$ of G . Thus we may assume that $X_i'' \subset s \cdot G$. Let $Y_i = s^{-1}(X_i'')$, where now s^{-1} denotes the preimage of the canonical projection $s : G \rightarrow s \cdot G$ defined by $s(g) = sg$, and divide the entries of the i -row B_i'' by s . In this way we obtain an extension of $\{(I_m|B''), 0, \mathcal{X}''\}$ where the map $\phi_i : Y_i \rightarrow X_i''$, $i \in [1, m]$, is the multiplication by s . Notice that, for each value of $y_i \in X_i'' \subset s \cdot G$, we can distinguish different solutions according to the value of $(0, -B_i''/s) \cdot y$. In particular, each solution has a unique value of $(0, -B_i''/s) \cdot y$. Even though the transformation is not bijective between X_i'' and Y_i , this transformation allows us to extend the system by distinguish the solutions according to the different values of $(0, -B_i''/s) \cdot y$. Moreover, it allows the values of the variable y_i to be in the whole G and

not only in $s \cdot G$, thus making the variable y_i more useful in order to parametrize the solutions.

By repeating the same procedure with each row of B'' we eventually obtain an extension $\{A', 0, \mathcal{Y}\}$ satisfying the conditions (i)-(iv) of the Lemma. Moreover, since all operations performed on A to obtain A' depend only on the entries of A and not on G , the condition (v) also holds. The condition (vi) is satisfied as we have added the last variables corresponding to the columns in B and they run over the full group G . This completes the proof. \square

Our final step is to show that, if the restricted system $\{A, 0, \mathcal{X}\}$, where A satisfies the conclusions of Lemma 4.10, is non-thin, then it admits an extension with a circular unimodular matrix.

Lemma 4.11. *Let $\{A, 0, \mathcal{X}\}$ be a non-thin restricted system where $A = (I_k | B)$ and $\gcd(B_i) = 1$ for every row B_i . There is an extension $\{A', 0, \mathcal{X}'\}$ with $k' = k'(A)$ depending only on the entries of A such that all matrices formed by k' consecutive columns of A' in the circular ordering are unimodular. Moreover, up to a reordering on the indices j , $\mathcal{X}' = \mathcal{X} \times \prod_{j=m+1}^{k'+m-k} G$.*

Proof. The stated extension is based on the following construction. Let M be a unimodular matrix of order $m - k$. By adding to M a row at the bottom of the form $M_1 + \sum_{i=2} \lambda_i M_i$, where M_i denotes the i -th row of M and $\lambda_i \in \mathbb{Z}$, the last $(m - k)$ rows of the resulting matrix form a unimodular matrix. By choosing appropriate row operations at each step we may transform M into the identity matrix. By putting each such transformation as a new row at the bottom of M we obtain a matrix of the form

$$M' = \begin{pmatrix} M \\ T \\ I_{m-k} \end{pmatrix}$$

such that every $(m - k) \times (m - k)$ submatrix of M' formed by consecutive rows is unimodular. The same procedure can be repeated by adding rows to the top of M to obtain a matrix of the form

$$M'' = \begin{pmatrix} I_{m-k} \\ S \\ M \\ T \\ I_{m-k} \end{pmatrix}$$

and again every $(m - k) \times (m - k)$ submatrix of M'' formed by consecutive rows is unimodular. Note that the dimensions of S and T depend on the number of row operations needed to transform M into the identity matrix. The operations to build T and S involve performing an Euclidian algorithm on the rows of M .

We apply the above procedure to the matrix B in the following manner. As each row B_i of the submatrix B is such that $\gcd(B_i) = 1$, we can apply Lemma 4.9 to the row B_i , by using $M = B_i$, $r = 1$ with $s = m - k$, and obtain a $(m - k) \times (m - k)$ square matrix \overline{B}_i with determinant ± 1 . Thus, by applying the above procedure to each of the resulting

matrices $\overline{B_1}, \dots, \overline{B_k}$ we may construct the following $k' \times (m - k)$ rectangular matrix:

$$B' = \begin{pmatrix} I_{m-k} \\ \overline{S_1} \\ \overline{B_1} \\ T_1 \\ I_{m-k} \\ \overline{S_2} \\ \overline{B_2} \\ T_2 \\ I_{m-k} \\ \dots \\ I_{m-k} \\ \overline{S_k} \\ \overline{B_k} \\ T_k \\ I_{m-k} \end{pmatrix},$$

for some k' depending on B . Let

$$A' = (I_{k'} | B').$$

Observe that every set of k' consecutive columns in the circular order in A' form a unimodular matrix. To check this, let $M(i)$ be the square submatrix formed by k' consecutive columns of A' in the circular order starting with the i -th column.

Since the matrix A' has the form

$$A' = \left(I_{k'} \mid \begin{array}{c} I_{m-k} \\ X \end{array} \right)$$

for some matrix X , then each matrix $M(i)$ for $i = 1, \dots, m - k$ is a circular permutation of a lower triangular matrix with all ones in the diagonal. Hence $M(i)$ is unimodular for these values of i . Moreover, if $i = k' - (m - k) + 1, \dots, k'$ then $M(i)$ is an upper triangular matrix with all ones in the diagonal.

For the remaining values of i , $\det M(i)$ equals, up to a sign, the determinant of a submatrix of B' formed by $m - k$ consecutive rows which, by construction, is unimodular. More precisely, $\det [M((m - k) + t)]$ equals, up to a sign, the determinant of the matrix formed by the rows $B'_{t+1}, B'_{t+2}, \dots, B'_{t+(m-k)}$.

In order to complete the proof of the Lemma we must construct the family \mathcal{X}' of $m' = k' + m - k$ sets. Let $I_0^1 \subset [1, k']$ be the set of subscripts i for which the i -row of B' corresponds to a row $\sigma(i)$ of the original matrix B and let $I_0^2 = [k' + 1, m']$. Let $I_0 = I_0^1 \cup I_0^2 \subset [1, m']$. By setting $X'_i = X_{\sigma(i)}$ for $i \in I_0^1$, $X'_i = X_{i-m'+m}$ for $i \in I_0^2$, and $X'_i = G$ otherwise, we get an extension $(A', 0, \mathcal{X}')$ of the given restricted system with

$$\phi : \prod_{i=1}^k X'_{\sigma^{-1}(i)} \times \prod_{i=k+1}^m X'_{i+m'-m} \rightarrow \prod_{i=1}^k X_i \times \prod_{i=k+1}^m X_i$$

the identity map. This completes the proof. \square

Observe that Lemma 4.10 and Lemma 4.11 can be concatenated to obtain a single, coherent, extension. The variables added in Lemma 4.10, that run over the whole group G , will also be moving over G after the second extension provided by Lemma 4.11. We summarize the results of this section in the following Proposition.

Proposition 4.12. *Let G be an Abelian group of order n . Let $\{A, b, \mathcal{X}\}$, where A is an integer $(k \times m)$ matrix, be a non-thin restricted system with $\gcd(d_k(A), n) = 1$. There is an extension $\{A', b', \mathcal{X}'\}$ of $\{A, b, \mathcal{X}\}$ with $k' = k'(A)$ such that A' is of the form $A' = (I_{k'} | B)$, $b' = 0$ and every k' consecutive columns of A' form a unimodular matrix.*

4.4 Proof of Theorem 4.1

We complete here the proof of Theorem 4.1. We assume that the system is not thin, otherwise, the result holds by deleting just one element of one set.

By Lemma 4.6 we may assume that $m' - k' \geq 2$. Let $\epsilon > 0$ and an integer $(k \times m)$ matrix A be given. Let G be an Abelian group of order n coprime with $d_k(A)$, and let $\{A, b, \mathcal{X}\}$ be a restricted system in G . It follows from Proposition 4.12 that there is an extension $\{A', 0, \mathcal{X}'\}$ of $\{A, b, \mathcal{X}\}$ such that A' is a circular unimodular matrix of dimension $(k' \times m')$ with $m' - k' = m - k$ and $k' = k'(A)$. Moreover there is a subset $I_0 \subset [1, m']$ with cardinality m , a bijection $\sigma : I_0 \rightarrow [1, m]$ and maps $\phi_i : X'_i \rightarrow X_{\sigma(i)}$, $1 \leq i \leq m$ such that the map $\phi : \mathcal{X}' \rightarrow \mathcal{X}$ with $(\phi(x'))_i = \phi_{\sigma^{-1}(i)}(x'_{\sigma^{-1}(i)})$ induces a bijection between the set of solutions of $\{A', 0, \mathcal{X}'\}$ and the set of solutions of $\{A, b, \mathcal{X}\}$. In addition, $I = [1, m'] \setminus I_0$ has cardinality less than k' and $X'_i = G$ for each $i \in I$.

We apply Lemma 4.3 to the extension $\{A', 0, \mathcal{X}'\}$ to obtain a set $\bar{\mathcal{X}}'$ with $|\bar{X}'_i| < \epsilon n$ for all $i \in [1, m']$ such that $\{A', 0, \mathcal{X}' \setminus \bar{\mathcal{X}}'\}$ has no solution. We use the last part of Lemma 4.3 to ensure that $\bar{\mathcal{X}}'$ can be chosen in such a way that $\bar{X}'_i = \emptyset$ for each $i \in I = [1, m'] \setminus I_0$. This shows that $\{A, b, \mathcal{X} \setminus \phi(\bar{\mathcal{X}}')\}$ is solution free and $|(\phi(\bar{\mathcal{X}}'))_i| < \epsilon n$ for $i \in [1, m]$. This completes the proof of Theorem 4.1.

4.5 On the condition over the determinantal

As has been already shown in Theorem 4.1, the removal lemma for linear systems over Abelian groups holds under the condition of the determinantal condition,. In this section, we expose a characterization of the systems whose determinantal is not coprime, and we see how this condition makes the system “well defined”. Let b an integer and let $b \cdot G$ denote the subgroup of G obtained with the images $\{b \cdot g\}_{g \in G}$.

Proposition 4.13 (Union of systems). *Let G be an Abelian group of order n , let A be a $k \times m$, $m \geq k + 2$, integer linear matrix, let b be a vector in G^k . Let S be the Smith Normal Form of A and let $\bar{d}_1, \dots, \bar{d}_k$ be the elements in the diagonal of S . Let $d_k(A) = \prod_{i=1}^k \bar{d}_i$ denote the k -th determinantal of A . Assume $\gcd(d_k(A), n) = d > 1$, then the system $Ax = b$, $x \in G^m$, is either incompatible, with no solutions at all, or the solution set is the disjoint union of the solution sets of $\prod_{i=1}^k |G|/|\bar{d}_i \cdot G|$ systems of the form $A'x = b'_i$ for $\prod_{i=1}^k |G|/|\bar{d}_i \cdot G|$ different independent vectors b'_i and where A' is such that $\gcd(d_k(A'), n) = 1$.*

Proof. [Proof of Proposition 4.13] Let $S = U^{-1}AV^{-1}$ be the Smith Normal Form of A , where U and V are unimodular matrices. We have $S = (D|0)$, where D is an $(k \times k)$ diagonal integer matrix with $|\det(D)| = |d_k(A)|$ and 0 is an all-zero $(k \times (m - k))$ matrix. Denote by \bar{d}_i the i -th element in the main diagonal of D .

Recall that U and V are the row and column operations respectively which transform A into S . Let $\overline{A}' = SV = U^{-1}A$ and let $\overline{b} = U^{-1}b$. Notice that the system $\overline{A}'x = \overline{b}$ is equivalent to $Ax = b$.

As \overline{A}' has been obtained from S by column operations, all the i -th row are integer multiples of \overline{d}_i . Since $\gcd(d_k(A), n) = d > 1$, there will be a row j whose \overline{d}_j is such that $\gcd(\overline{d}_j, n) = d_j > 1$. If we read of the j -th equation of the system $\overline{A}'x = \overline{b}$, we obtain:

$$d_j (A'_{j,1}x_1 + A'_{j,2}x_2 + \cdots + A'_{j,m}x_m) = \overline{b}_j. \quad (4.4)$$

Notice that, $A'_{j,1}x_1 + \cdots + A'_{j,m}x_m$ is an element of G , and $d_j (A'_{j,1}x_1 + \cdots + A'_{j,m}x_m)$ is an element of the subgroup $d_j \cdot G \subsetneq G$. Therefore, if \overline{b}_j is not in the subgroup $d_j \cdot G$, the system is incompatible.

On the other hand, if \overline{b}_j is in $d_j \cdot G$, then the solutions to the equation (4.4) are the union of the solutions to the equations

$$A'_{j,1}x_1 + A'_{j,2}x_2 + \cdots + A'_{j,m}x_m = b_j^i, \quad (4.5)$$

for all the possible preimages of \overline{b}_j through the application “multiply by d_j ”: $b_j^i \in d_j^{-1}\overline{b}_j$. The set of solutions to the system of equations is the union of the solution sets of the systems where all the other equations remain untouched and the j -th equation is replaced by all the possible equations (4.5).

By doing this procedure in all the equations for which $\gcd(\overline{d}_i, n) = d_i > 1$ and considering that, each time we have to process a new equation, we should consider the equation for each of the possible independent vectors previously found, we obtain a new matrix A' , whose row j is the j -th row from \overline{A}' divided by d_j , and a set of independent vectors \mathcal{B} with cardinality $s = \prod_{i=1}^k |G|/|\overline{d}_i \cdot G|$, $\mathcal{B} = \{b_i\}_{i \in [1,s]}$.

The whole union of systems, considering A' and \mathcal{B} , has the same solution set as $Ax = b$, and fulfills the desired properties since $\gcd(d_k(A'), n) = 1$. \square

As we have seen in the proposition, if the determinantal is not coprime with the order of the group, then we do not have a single linear system but a union of systems of equations where the matrix have coprime determinantal and we change the independent vector b . In the finite field context, this collapses to the case where the matrix A does not have full-rank. However, in the framework of integer systems for Abelian groups, we are dealing with a richer variety of possibilities. See Section 6.5 for further comments.

4.6 Example of the construction of the extension of the matrix A

This final section expose a little example of the construction of the extension of the matrix A using the procedures in this chapter. This extension allows the matrix C to be an integer matrix with -1 in the diagonal.

Let A be the matrix:

$$\begin{pmatrix} 2 & 3 & 15 & 14 & 8 \\ 4 & 9 & 10 & 7 & 6 \end{pmatrix},$$

for which $d_k(A) = 1$.

In this case, the first extension of the matrix using Lemma 4.9 leads us to:

$$\begin{pmatrix} 2 & 3 & 15 & 14 & 8 & 0 & 0 & 0 \\ 4 & 9 & 10 & 7 & 6 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ -3 & -6 & -13 & -10 & -7 & 0 & 1 & 0 \\ -4 & -8 & -17 & -14 & -9 & 0 & 0 & 1 \end{pmatrix},$$

whose equivalent matrix of the form $(I_5|B)$ is

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & -44 & -35 & -27 \\ 0 & 1 & 0 & 0 & 0 & 17 & 14 & 10 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 2 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 3 \end{pmatrix}.$$

These are some possible extensions of the different rows B_i (the boldface marks the original rows in B):

B_1			B_2			B_3			B_4			B_5		
			1	0	0									
			0	1	0									
1	0	0	0	0	1									
0	1	0	1	1	0									
0	0	1	0	1	0									
1	0	0	0	1	1									
0	1	0	1	0	0									
0	1	1	0	3	2				1	0	0			
1	0	0	0	17	11				0	1	0	1	0	0
0	4	3	1	3	2				0	0	1	0	1	0
0	9	7	17	14	10				1	0	1	0	0	1
1	0	0	-8	-7	-5	1	0	0	1	1	0	0	0	1
0	13	10	3	3	2	0	1	0	1	2	0	0	1	0
-44	-35	-27	1	4	2	0	0	1	0	0	1	1	0	3
1	0	0	0	7	3				1	1	0	0	1	0
16	13	10	0	-5	-2				1	0	0	1	0	0
0	9	7	1	0	0				0	1	0	0	1	0
1	0	0	0	2	1				0	0	1	0	0	1
0	4	3	0	1	0									
0	1	1	1	0	0									
1	0	0	0	1	1									
0	1	0	0	0	1									
0	0	1	1	0	0									
			0	1	0									
			0	0	1									

and we just have to sew together those extensions to obtain our final matrix B that, along with an identity of the appropriate dimension in front like $(I_{k'}|B)$, will fulfill the properties of Lemma 4.11. This extension of A makes the construction of the matrix C easier.

Applications to Arithmetic Ramsey Theory

In this chapter we discuss some applications of the algebraic removal lemmas in arithmetic Ramsey theory. In particular we show that, for any r -coloring, the number of monochromatic solutions is, under some natural conditions, as large as we might expect for finite fields Theorem 5.5 and for Abelian group with bounded torsion Theorem 5.26. Moreover, we show the density case of these Ramsey results.

The chapter is organized as follows. Section 5.1 gives an introduction to the subject and presents some known results. Section 5.2 is devoted to show the result for finite fields, Theorem 5.5; it also serves as an example of how to proceed in the proof of the counting results for finite Abelian groups, Theorem 5.26 and Theorem 5.7. In Section 5.3 we show the counting result for monochromatic solutions in groups of the type \mathbb{Z}_n^M for fixed n and growing M ; this is the main result and the statements for bounded torsion, Theorem 5.26, follows from it. Section 5.5 is devoted to show Theorem 5.28, the characterization of the systems such that, in any dense subset, there is a large number of solutions. Finally, Section 5.6 discusses some remarks on the results of the chapter; for a further discussion see Chapter 6.

5.1 Introduction

In 1933 Rado [59] characterized the homogeneous integer linear systems for which, for any coloring (using finitely many colors) of the integers except the zero, there would exist a monochromatic solution. The notion he used to characterize those systems is the so-called *columns condition*.

Let A be a $k \times m$ integer matrix. We say that A fulfills the *columns condition* if we can order the column vectors A^1, \dots, A^m and find $1 \leq k_1 < k_2 \dots < k_t = m$ (with $k_0 = 0$) such that, if we set

$$S^i = \sum_{j=k_{i-1}+1}^{k_i} A^j,$$

we have that

- (i) $S^1 = 0$ in \mathbb{Z}^k .
- (ii) for $1 < i \leq t$, S^i can be expressed as a linear combination of $A^1, \dots, A^{k_{i-1}}$ using coefficients in \mathbb{Q} .

We say that A is *partition regular* if, for any partition of $\mathbb{Z} \setminus \{0\}$ using finite many colors, there is always a monochromatic solution to $Ax = 0$.

Theorem 5.1 (Rado [59]). *A is partition regular if and only if A fulfills the columns condition.*

In 1988 Frankl, Graham and Rödl [28] showed that if A , a $k \times m$ integer matrix, satisfies the columns condition, then there is not just one, but many monochromatic solutions to the homogeneous system $Ax = 0$.

Theorem 5.2 (Frankl, Graham, Rödl [28]). *Let r be a positive integer. Assume A fulfills the columns condition. Then, there exists a constant $c = c(r, A) > 0$ such that for every r -coloring of $[1, N]$, there are at least $c(N^{m-k})$ monochromatic solutions to $Ax = 0$.*

In this chapter, we show that, under certain conditions similar to the columns condition, the number of monochromatic solutions has the order we might expect in other contexts such as the finite fields or finite Abelian groups. In particular, we obtain that the number of monochromatic solutions has to be a constant times the size of the whole algebraic structure raised to the power of the degrees of freedom of the system; the constant depends heavily on the number of colors.

5.2 Number of monochromatic solutions inside \mathbb{F}_q^N

5.2.1 Introduction

In [8], Bergelson, Deuber and Hindman proved Theorem 5.4, a Rado-like theorem for finite fields. They characterize the partition regular systems for finite fields; this is, the systems that have monochromatic solutions for any finite coloring of the \aleph_0 -dimensional vector space over a finite field F . The notion they used, following the works by Deuber [20], is the F -columns condition.

Let F be a finite field of order q and let N be a positive integer. For this section, we let A be a $k \times m$ matrix with coefficients in F . Let $\chi : F^N \rightarrow [r]$ be a coloring with r colors. We are interested in solutions of the system $Ax = 0$, with $x = (x_1, \dots, x_m)$, $x_i \in F^N \setminus \{0\}$ and $\chi(x_1) = \chi(x_2) = \dots = \chi(x_m)$ (monochromatic). This is, $x = (x_1, \dots, x_m)$ is a solution of $Ax = 0$ if, given $(a_{i,1}, \dots, a_{i,m})$ the i -th row of A , then $\sum_{j=1}^m a_{i,j}x_j^k = 0$ for all rows $(a_{i,1}, \dots, a_{i,m})$ and for each of the $k \in [1, N]$ coordinates of $x = (x_1, \dots, x_m)$ simultaneously. Moreover, to have a monochromatic solution, each of the x_i has to be painted with the same color.

Let us denote with A^i the i -th column and A_j the j -th row.

Definition 5.3 (F -columns condition). *We say that A fulfills the F -columns condition if we can order the column vectors A^1, \dots, A^m and find $1 \leq k_1 < k_2 < \dots < k_t = m$ (with $k_0 = 0$) such that, if we set*

$$S^i = \sum_{j=k_{i-1}+1}^{k_i} A^j,$$

we have that

(i) $S^1 = 0$ in F^k .

(ii) for $1 < i \leq t$, S^i can be expressed as a linear combination of A^1, \dots, A^{k_i-1} using coefficients in F .

Using the notion of F -columns condition, the main result from [8] can be stated.

Theorem 5.4 (Bergelson, Deuber, Hindman [8]). *Let F be a finite field, let k, m be positive integers, and let A be a $k \times m$ matrix with coefficients in F . The following two statements are equivalent.*

1. *For each $r \in \mathbb{N}$, there is some $M \in \mathbb{N}$, $M = M(r, m, |F|)$, so that whenever $n \geq M$ and V is an n -dimensional vector space over F , and $V \setminus \{0\}$ is r -colored, there exist monochromatic $x_1, \dots, x_m \in V \setminus \{0\}$ with $Ax = 0$.*
2. *A satisfies the F -columns condition.*

Following the relation of [28, Theorem 1] (stated here as Theorem 5.2) with respect to Rado's Theorem, asking how many monochromatic solutions are there, for a fixed number of colors, appears to be a natural question. Theorem 5.5 tries to address this issue, at least asymptotically, in the context of finite fields.

Theorem 5.5 (Number of monochromatic solutions in Finite Fields). *Let F be a finite field with $q = p^l$ elements, let k, m, N, r be positive integers, $m \geq k$, and let A be a $k \times m$ matrix with coefficients in F . Assume that A satisfies the F -columns condition. Then, for any coloring of the elements of F^N with r colors, there exists a $\delta = \delta(r, q, m) > 0$ such that the system $Ax = 0$ with $x = (x_1, \dots, x_m)$ and $x_i \in F^N$ has, at least, $\lfloor \delta(q^N)^{m-k} \rfloor$ monochromatic solutions.*

Theorem 5.5 says that if the set of monochromatic solutions is non-empty for each coloration, then it grows as expected. The proof uses the Removal Lemma for finite fields, Theorem 3.1, proved independently by Shapira [68] and Král', Serra and the author [52].

Following the proof in [8], one can check that the M on Theorem 5.4 depends on r, q , and on the minimal number of parts in the partition used to determine that A fulfills the F -columns partition (t in Definition 5.3). However, since we use the Removal Lemma to find the constant in Theorem 5.5, δ depends on m .

5.2.2 Proof of Theorem 5.5

The ingredients of the proof are the following.

1. A Ramsey result that finds monochromatic solutions in a given substructure. This result is Theorem 5.4.
2. A Counting Statement: the number of substructures where we are able to find monochromatic solutions is as large as expected.
3. A Removal Lemma: if there are not many solutions, then, by destroying not many elements, we are capable of destroying all the solutions. This is Theorem 3.1.
4. Argue that, if we remove not many elements, a substructure where we find a monochromatic solution has to survive.

The steps above are the same we use to show Theorem 5.7 and Theorem 5.26.

Let us notice that the operations in the system of equations involve the group structure and the product by the elements of the finite field. Therefore, we use Theorem 3.1 restricting the coefficients of A to the field F and we exchange the base field where the subsets X_i are located by an N -dimensional vector space over F , F^N .

Proof. [Proof of Theorem 5.5] Let $F = \mathbb{F}_q$ be the finite field over q elements. Let F^N be an N -dimensional space over F . Let r be the number of colors and let m be the number of columns of the matrix A . Let Y_i be the set colored with i , $i \in [1, r]$. Let $M = M(r, m, q)$ be the constant obtained from Theorem 5.4 such that F^M contains a monochromatic solution.

Recall that the number of M -dimensional subspaces in F^N is given by the gaussian coefficient:

$$\binom{N}{M}_q = \frac{(q^N - 1)(q^{N-1} - 1) \cdots (q^{N-M+1} - 1)}{(q - 1)(q^2 - 1) \cdots (q^M - 1)}.$$

The number of M -dimensional subspaces that go through a point $a \in F^N$, $a \neq 0$ is

$$\binom{N}{M}_q \frac{q^M - 1}{q^N - 1} = \binom{N-1}{M-1}_q.$$

Indeed, a generates, using the multiples of a by elements of F , a linear variety inside F^N ; name this variety $\langle a \rangle$. The number of M -dimensional subspaces of F^N that go through a is the same as the number of $(M-1)$ -dimensional subspaces inside $F^N / \langle a \rangle \cong F^{N-1}$, this is: $\binom{N-1}{M-1}_q$.

Since M is constant, if N is large enough, there exist constants $c_1 = c_1(q, M)$, $c_2 = c_2(q, M)$, $c_3 = c_3(q, M)$ and $c_4 = c_4(q, M)$ such that:

$$c_3 q^{MN} \geq \binom{N}{M}_q \geq c_1 q^{MN}$$

and

$$c_4 q^{(M-1)(N-1)} \geq \binom{N-1}{M-1}_q \geq c_2 q^{(M-1)(N-1)}.$$

At this point, we apply Theorem 3.1 r times, one for each color, with $\epsilon = \frac{c_1 q^M}{m^2 r c_4 q}$ and $X_1 = X_2 = \cdots = X_m = Y_i$ for $i \in [1, r]$. If the number of monochromatic solutions is at most $\delta_{\text{Theorem 3.1}} (q^N)^{m-k}$, we obtain sets $Y'_i, Y''_i \subset Y_i$ with $|Y'_i| \leq \epsilon m q^N$ such that $S = F^N \setminus \bigcup_{i \in [1, r]} Y'_i$ has no monochromatic solution.

As each element $x \in \bigcup_{i \in [1, r]} Y'_i$ belongs to, at most, $c_4 q^{(M-1)(N-1)}$ M -dimensional subspaces and we have removed, at most, $r \epsilon m q^N$ elements, we have destroyed, at most:

$$r \epsilon m q^N c_4 q^{(M-1)(N-1)} = \frac{c_1}{2} q^{MN}$$

M -dimensional subspaces of F^N .

However, since the number of M -dimensional subspaces is, at least, $c_1 q^{MN}$, we should have, at least, $\frac{c_1}{2} q^{MN}$ M -dimensional subspaces that do not contain any removed element. Therefore, using Theorem 5.4, there should exist a monochromatic solution in S , contradicting the assumptions of the Removal Lemma. Thus, we should have more than

$\delta_{\text{Theorem 3.1}}(q^N)^{m-k}$ monochromatic solutions. Let N_0 be such that $\frac{c_1}{2}q^{MN_0} \geq 1$, then by picking

$$\delta = \min \left(\frac{1}{2(q^{N_0})^{m-k}}, \delta_{\text{Theorem 3.1}} \right)$$

the result is proved. \square

5.3 Number of monochromatic solutions inside \mathbb{Z}_n^N

5.3.1 Introduction

In this section, we show that the number of monochromatic solutions in a group \mathbb{Z}_n^N is asymptotically as large as one might expect. Let us introduce the notion of n -columns condition.

Definition 5.6 (n -columns condition). *Let A be a $(k \times m)$ matrix with integer entries. We say that A fulfills the n -columns condition if we can order the column vectors A^1, \dots, A^m and find $1 \leq k_1 < k_2 < \dots < k_t = m$ (with $k_0 = 0$) such that, if we set*

$$S^i = \sum_{j=k_{i-1}+1}^{k_i} A^j,$$

we have that

- (i) $S^1 = 0$ in \mathbb{Z}^k .
- (ii) for $1 < i \leq t$, S^i can be expressed as a linear combination of $A^1, \dots, A^{k_{i-1}}$ using coefficients in \mathbb{Z}_n ; this is, $S^i = \sum_{j=1}^{k_{i-1}} \lambda_j A^j$ with $\lambda_j \in \mathbb{Z}_n$.

The main result of this section is the following.

Theorem 5.7 (Number of monochromatic solutions in \mathbb{Z}_n^N). *Let r be a positive integer and let A be a $(k \times m)$ matrix with integer entries. Assume that A satisfies the n -columns condition. There is a constant $c = c(r, n, m) > 0$ such that every r -coloring of $\mathbb{Z}_n^N \setminus \{0\}$ has at least $\lfloor c|\mathbb{Z}_n^N|^{m-k} \rfloor$ monochromatic solutions of the equation $Ax = 0$ in $\mathbb{Z}_n^N \setminus \{0\}$.*

The scheme of the proof is the same as in Section 5.2. Unfortunately, a result like Theorem 5.4 or, more precisely, Theorem 5.5 does not give us enough solutions, as the number of subspaces isomorphic to \mathbb{Z}_p^N inside \mathbb{Z}_n^N , for any prime p dividing n , is not large enough. For this purpose, we prove Lemma 5.8.

Lemma 5.8 (Solutions outside the finite fields). *Let n be a composite number (not prime). Let A be a $k \times m$ integer matrix fulfilling the n -columns condition. There exists an $M = M(r, n, m)$ such that, for any coloring of the elements of $\mathbb{Z}_n^M \setminus \{0\}$, with r colors, there exist a monochromatic solution to the system $Ax = 0$, with $x_i \in \mathbb{Z}_n^M \setminus \bigcup_{p|n} \mathbb{Z}_p^M$. Moreover, the order of all the x_i is n .*

Section 5.3.2 is devoted to show Theorem 5.7 using Lemma 5.8 following the scheme detailed in Section 5.2, while Section 5.3.3 is devoted to the proof of Lemma 5.8.

5.3.2 Proof of Theorem 5.7

As in Section 5.2, we use a counting result for subgroups isomorphic to \mathbb{Z}_n^M in \mathbb{Z}_n^N . Recall that if $n = p_1^{\alpha_1} \cdots p_l^{\alpha_l}$ is the decomposition of n in its prime factors, then $\mathbb{Z}_n^M \cong \prod_{i=1}^l \mathbb{Z}_{p_i}^{M\alpha_i}$. Thus, the number of subgroups isomorphic to \mathbb{Z}_n^M is the product of the number of subgroups isomorphic to $\mathbb{Z}_{p_i}^{M\alpha_i}$, for each i :

$$|\mathbb{Z}_n^M \subset \mathbb{Z}_n^N| = \prod_{i=1}^l |\mathbb{Z}_{p_i}^{M\alpha_i} \subset \mathbb{Z}_n^N| = \prod_{i=1}^l |\mathbb{Z}_{p_i}^{M\alpha_i} \subset \mathbb{Z}_{p_i}^{N\alpha_i}|.$$

Let us recall a result regarding the number of subgroups of a p -group, p prime. A p -group H is of type $(k_1, k_2, \dots, k_\eta)$ if $H \cong \prod_{i=1}^\eta \mathbb{Z}_{p^{k_i}}$.

Theorem 5.9 (Number of subgroups of a p -group, [80]). *Let G be a prime power Abelian group of order $p^{k_1+k_2+\dots+k_\eta}$, type $(k_1, k_2, \dots, k_\eta)$, where k 's are arranged in ascending order of magnitude. Let*

$$\begin{aligned} h_1 = h_2 = \dots = h_{m_1} &> h_{m_1+1} = \dots = h_{m_1+m_2} > \dots \\ &> h_{m_1+m_2+\dots+m_{r-1}+1} = \dots = h_{m_1+m_2+\dots+m_r}, \end{aligned} \quad (5.1)$$

where $m_1 + m_2 + \dots + m_r = m \leq \eta$, be m positive integers not greater than k_η , and let ν_i be such that $k_{\nu_i} < h_i \leq k_{\nu_i+1}$ ($i = 1, 2, \dots, m$; $k_0 = 0$). Then the number of subgroups of type (5.1) is given by

$$p^H \prod_{i=1}^m (p^{\eta-\nu_i-i+1} - 1) \Big/ \prod_{\mu=1}^r \prod_{\nu=1}^{m_\mu} (p^\nu - 1)$$

where

$$\begin{aligned} H &= \sum_{i=1}^m (\eta - \nu_i + 1 - 2i)(h_i - 1) \\ &+ \frac{1}{2}(m_1^2 + m_2^2 + \dots + m_r^2 - m^2) + \sum_{i=1}^m \sum_{\mu=0}^{\nu_i} k_\mu. \end{aligned}$$

Let us show a small proposition that will be used in the proof of Proposition 5.11.

Proposition 5.10. *Let $G = \mathbb{Z}_{n_1} \times \dots \times \mathbb{Z}_{n_s}$ be an Abelian group with $n_1|n_2|\dots|n_s$. If H is a subgroup of G isomorphic to \mathbb{Z}_{n_s} , then $G/H \cong \prod_{i=1}^{s-1} \mathbb{Z}_{n_i}$.*

Let us notice that this fact is not true if H is isomorphic to a smaller cyclic group.

Proof. [Proof of Proposition 5.10] Let $n_i = p_1^{\alpha_1^i} \cdots p_t^{\alpha_t^i}$ be the decomposition of n_i into its prime factors. So, we have $\alpha_i^j \geq 0$ and $\alpha_i^{j_1} \geq \alpha_i^{j_2}$ if $j_1 \geq j_2$. Then, we can rewrite G as:

$$G_0 = \prod_{j=1}^s \prod_{i=1}^t \mathbb{Z}_{p_i^{\alpha_i^j}}.$$

Let a be an element of order n_s generating H , $\langle a \rangle = H$. If we write a in the canonical coordinates of G_0 then, for any $i \in [1, t]$, there exists a coordinate of a in a component isomorphic to $\mathbb{Z}_{p_i^{\alpha_i^s}}$ whose value is coprime with p_i ; otherwise we would have that a has order smaller than n_s . Let I , $|I| = t$, be the set of indices such that:

- if $j \in I$, then the j -th component of G_0 is isomorphic to $\mathbb{Z}_{p_i^{\alpha_i^s}}$, for some p_i . Moreover, for any p_i , there exists a unique $j \in I$ with $G_0^j \cong \mathbb{Z}_{p_i^{\alpha_i^s}}$.
- if $j \in I$, the j -th component of a , a_j , is coprime with the order of the corresponding component:

$$\langle (0, \dots, 0, \pi_j(a), 0, \dots, 0) \rangle = \langle (0, \dots, 0, \overbrace{1}^j, 0, \dots, 0) \rangle,$$

where π_j represents the projection onto the j -th coordinate of G_0 .

Thus, we can choose as our new base, or generators of the group:

$$\{(0, \dots, 0, \overbrace{1}^i, 0, \dots, 0)\}_{i \in [1, ts] \setminus I} \cup \{a\} = \{e_i\}_{i \in [1, ts] \setminus I} \cup \{a\},$$

as the elements generated by $\{e_i\}_{i \in I}$ can be obtained by combining a and the $\{e_i\}_{i \in [1, ts] \setminus I}$. Therefore, we obtain that $G_0/\langle a \rangle = G/\langle a \rangle \cong \prod_{i=1}^{s-1} \mathbb{Z}_{n_i}$. \square

Let us show a proposition that is used in this section, as well as in Section 5.4 and Section 5.5.

Proposition 5.11. *Let G be an Abelian group of order n . Let n_s be the maximum among the orders of the elements of G . Assume that $\mathbb{Z}_{n_s}^M$ is a subgroup of G . Then, there exist two constants $c_1 = c_1(M, n_s) > 0$ and $c_2 = c_2(M, n_s) > 0$ such that*

$$|\mathbb{Z}_{n_s}^M \subset G| \geq c_1 |G|^M$$

and

$$|\mathbb{Z}_{n_s}^{M-1} \subset G/\mathbb{Z}_{n_s}| \leq c_2 |G|^{M-1}.$$

Throughout this work, n_s and M will be constants, and the order of the group will be thought to be large, tending to infinity. Consequently, we can treat c_1 and c_2 as a constant for asymptotic purposes.

Proof. [Proof of Proposition 5.11] For each $p|n$, we use Theorem 5.9 to each of the components $G_p = \prod_{i=1}^l \mathbb{Z}_{p^{\alpha_i}}$ of G , where $\alpha_i \geq \alpha_j$ whenever $i > j$, to find the number of subgroups of $\mathbb{Z}_{p^{\alpha_l}}^M$ in G_p . Let l_{α_l} be the number of copies of $\mathbb{Z}_{p^{\alpha_l}}$ in G_p . We use Theorem 5.9 with:

$$\begin{aligned} \eta &= l \\ m &= m_1 = M \\ h_1 &= \dots = h_{m_1} = \alpha_l \\ \nu_i &= \nu_0 = (l - l_{\alpha_l}), i \in [1, m_1]. \end{aligned}$$

Therefore the number of copies of $\mathbb{Z}_{p^{\alpha_l}}^M$ in G_p is

$$\begin{aligned} \frac{1}{\prod_{\nu=1}^M (p^\nu - 1)} p^H \prod_{i=1}^M (p^{l - (l - l_{\alpha_l}) - i + 1} - 1) &> c \frac{1}{2^M} p^H p^{(l_{\alpha_l} - M + 1)M} \\ &= c' p^H p^{l_{\alpha_l} M} \end{aligned}$$

for some $c = c(M, p^{\alpha_l}) > 0$ and $c' = c'(M, p^{\alpha_l}) > 0$. On the other side:

$$\begin{aligned}
H &= \sum_{i=1}^M (l - (l - l_{\alpha_i}) + 1 - 2i)(\alpha_i - 1) + 0 + \sum_{i=1}^M \sum_{\mu=0}^{\nu_i} k_{\mu} \\
&= \sum_{i=1}^M (l_{\alpha_i} + 1 - 2i)(\alpha_i - 1) + M \sum_{\mu=0}^{\nu_0} k_{\mu} = \sum_{i=1}^M (l_{\alpha_i} + 1 - 2i)(\alpha_i - 1) + M \binom{l - l_{\alpha_l}}{\sum_{i=1}^{l-l_{\alpha_l}} \alpha_i} \\
&\geq M(l_{\alpha_l} + 1 - 2M)(\alpha_l - 1) + M \binom{l - l_{\alpha_l}}{\sum_{i=1}^{l-l_{\alpha_l}} \alpha_i} = d' + Ml_{\alpha_l}\alpha_l - Ml_{\alpha_l} + M \binom{l - l_{\alpha_l}}{\sum_{i=1}^{l-l_{\alpha_l}} \alpha_i} \\
&= d' - Ml_{\alpha_l} + M \binom{l}{\sum_{i=1}^l \alpha_i}
\end{aligned}$$

for some constant $d' = d'(M, \alpha_l) > 0$. Therefore, the number of subgroups isomorphic to \mathbb{Z}_p^M in G_p is, at least:

$$\begin{aligned}
c' p^{d' - Ml_{\alpha_l} + M(\sum_{i=1}^l \alpha_i)} p^{l_{\alpha_l} M} &= c' p^{d' + M(\sum_{i=1}^l \alpha_i)} = c_{1,p} p^{M(\sum_{i=1}^l \alpha_i)} \\
&= c_{1,p} |G_p|^M,
\end{aligned}$$

for $c_{1,p} = c_{1,p}(M, p^{\alpha_l}) > 0$.

Lets compute an upper bound for the number of subgroups isomorphic to \mathbb{Z}_p^{M-1} in $G_p/\mathbb{Z}_p^{\alpha_l}$. We use Proposition 5.10 to see that if G_p is of the type

$$(\alpha_1, \alpha_2, \dots, \alpha_{l-1}, \alpha_l),$$

then $G_p/\mathbb{Z}_p^{\alpha_l}$ is of the type

$$(\alpha_1, \alpha_2, \dots, \alpha_{l-1}).$$

Therefore,

$$\begin{aligned}
\left| \mathbb{Z}_p^{M-1} \subset G_p/\mathbb{Z}_p^{\alpha_l} \right| &= \frac{1}{\prod_{\nu=1}^{M-1} (p^{\nu} - 1)} p^H \prod_{i=1}^{M-1} (p^{(l-1) - (l-l_{\alpha_i}) - i + 1} - 1) \\
&\leq c'' p^H \prod_{i=1}^{M-1} p^{l_{\alpha_i} - i} \leq c''' p^H p^{l_{\alpha_l}(M-1)}
\end{aligned}$$

for some $c'' = c''(M, p) > 0$ and $c''' = c'''(M, p) > 0$. On the other side:

$$\begin{aligned}
H &= \sum_{i=1}^{M-1} ((l-1) - (l-l_{\alpha_i}) + 1 - 2i)(\alpha_i - 1) + 0 + \sum_{i=1}^{M-1} \sum_{\mu=0}^{\nu_i} k_{\mu} \\
&= \sum_{i=1}^{M-1} (l_{\alpha_i} - 2i)(\alpha_i - 1) + (M-1) \left(\sum_{i=1}^{l-l_{\alpha_l}} \alpha_i \right) \\
&\leq \sum_{i=1}^{M-1} l_{\alpha_i}(\alpha_i - 1) + (M-1) \left(\sum_{i=1}^{l-l_{\alpha_l}} \alpha_i \right) \\
&= (M-1)l_{\alpha_l}\alpha_l - (M-1)l_{\alpha_l} + (M-1) \left(\sum_{i=1}^{l-l_{\alpha_l}} \alpha_i \right) \\
&= -(M-1)l_{\alpha_l} + (M-1) \left(\sum_{i=1}^l \alpha_i \right)
\end{aligned}$$

Putting both together:

$$\begin{aligned}
\left| \mathbb{Z}_p^{M-1} \subset G_p / \mathbb{Z}_p^{\alpha_l} \right| &\leq c''' p^H p^{l_{\alpha_l}(M-1)} \\
&\leq c''' p^{-(M-1)l_{\alpha_l} + (M-1)(\sum_{i=1}^l \alpha_i)} p^{l_{\alpha_l}(M-1)} = c_{2,p} p^{(M-1)(\sum_{i=1}^l \alpha_i)} \\
&= c_{2,p} G_p^{M-1}
\end{aligned}$$

for some $c_{2,p} = c_{2,p}(M, p^{\alpha_l}) > 0$.

Recall that the number of subgroups isomorphic to $\mathbb{Z}_{n_s}^M \subset G$ is

$$\left| \mathbb{Z}_{n_s}^M \subset G \right| = \prod_{p|n_s} \left| \mathbb{Z}_p^{M_{\alpha_i p}} \subset G_p \right| \geq \prod_{p|n_s} (c_{1,p} |G_p|^M)$$

Therefore,

$$\left| \mathbb{Z}_{n_s}^M \subset G \right| \geq c_1 |G|^M,$$

and we have show the first part. To show the second part we use the same argument but with $c_{2,p}$ instead of $c_{1,p}$ and $(M-1)$ instead of M to obtain the desired result:

$$\left| \mathbb{Z}_{n_s}^{M-1} \subset G / \mathbb{Z}_{n_s} \right| \leq c_2 |G|^{M-1}.$$

□

If we particularize Proposition 5.11 in our case, for a given n ,

$$\left| \mathbb{Z}_n^M \subset \mathbb{Z}_n^N \right| \geq c_1 n^{MN}, \tag{5.2}$$

and

$$\left| \mathbb{Z}_n^{(M-1)} \subset \mathbb{Z}_n^{(N-1)} \right| \leq c_2 \frac{n^{MN}}{n^{N^2}}, \tag{5.3}$$

In particular, we can say that the number of subgroups of the type $(\mathbb{Z}_p^{\alpha})^M \subset G_p$ has “positive probability”: if we pick M elements uniformly at random from G_p , x_1, \dots, x_M , we have a positive probability that $\langle x_1, \dots, x_M \rangle \cong \mathbb{Z}_p^M$.

As our intention is to use the Removal Lemma for Abelian groups, Theorem 4.1, let us introduce two technical propositions that allow us to circumvent the condition regarding the primality between $d_k(A)$ and n .

Proposition 5.12 (from n -columns condition to conditions over \mathbb{Z}). *Let A be a $k \times m$ integer matrix and let n be a positive integer. Assume A fulfills the n -columns condition. Then, there exists a $k \times m$ integer matrix \bar{A} such that:*

- $A_{\mathbb{Z}_n} = \bar{A}_{\mathbb{Z}_n}$: A is equivalent to \bar{A} if we look the matrices in \mathbb{Z}_n , instead of looking them as matrices in \mathbb{Z} .
- \bar{A} fulfills the \mathbb{Z} -columns condition.

Proof. [Proof of Proposition 5.12] Assume that $1 \leq k_1 < \dots < k_t = m$ define the partition of the columns that bears the n -columns partition. The first equation reads:

$$S^1 = \sum_{j=1}^{k_1} A^j \stackrel{\mathbb{Z}_n}{=} 0$$

Thus, there exists an integer column $\mu_1 \in \mathbb{Z}^k$ such that

$$S^1 = \sum_{j=1}^{k_1} A^j \stackrel{\mathbb{Z}}{=} 0 + n\mu_1$$

Now let $\bar{A}^i = A^i$, $i \in [1, k_1 - 1]$ and $\bar{A}^{k_1} = A^{k_1} - n\mu_1$. Observe that $\bar{A}^{k_1} \stackrel{\mathbb{Z}_n}{=} A^{k_1}$.

We build the new columns inductively, assume that we have built up to column l and we want to build the $(l+1)$ -th. If $l+1 \neq \{k_i\}_{i \in [1, t]}$, then $\bar{A}^{l+1} = A^{l+1}$. Suppose $l+1 = k_s$, for some $s > 1$, then

$$\begin{aligned} \sum_{i=k_{s-1}+1}^{k_s} A^i \stackrel{\mathbb{Z}_n}{=} \sum_{i=1}^{k_{s-1}} \lambda_{i,s} A^i &\Rightarrow A^{k_s} + \sum_{i=k_{s-1}+1}^{k_s-1} \bar{A}^i \stackrel{\mathbb{Z}_n}{=} \sum_{i=1}^{k_{s-1}} \lambda_{i,s} \bar{A}^i \Rightarrow \\ A^{k_s} + \sum_{i=k_{s-1}+1}^{k_s-1} \bar{A}^i &\stackrel{\mathbb{Z}}{=} \sum_{i=1}^{k_{s-1}} \lambda_{i,s} \bar{A}^i + n\mu_s \end{aligned}$$

for a certain integer column μ_s with k components. Thus, defining $\bar{A}^{k_s} \stackrel{\mathbb{Z}}{=} A^{k_s} - n\mu_s$, we obtain:

$$\sum_{i=k_{s-1}+1}^{k_s} \bar{A}^i \stackrel{\mathbb{Z}}{=} \sum_{i=1}^{k_{s-1}} \lambda_{i,s} \bar{A}^i \Rightarrow \sum_{i=k_{s-1}+1}^{k_s} \bar{A}^i \stackrel{\mathbb{Z}_n}{=} \sum_{i=1}^{k_{s-1}} \lambda_{i,s} \bar{A}^i$$

By observing that the $\lambda_{i,s}$ can be considered integers, we have built a matrix \bar{A} that satisfies the conditions in the proposition. \square

Proposition 5.13. *Let A be a $k \times m$ integer matrix. Let G be an Abelian group. Assume that $d_k(A) > 1$. Let $L = L(A)$ be the set of linear equations satisfied by the columns of A with rational coefficients. Then, there exists a $k \times m$ integer matrix \bar{A} with:*

1. If $x = (x_1, \dots, x_m)$, $x_i \in G$, is a solution of $\bar{A}x = 0$, then x is also a solution to $Ax = 0$.
2. The k -th determinantal of \bar{A} is one, $d_k(\bar{A}) = 1$.
3. The columns of \bar{A} satisfy the same linear equations as the columns of A : $L(\bar{A}) = L(A)$.

Proof. [Proof of Proposition 5.13] Lets consider the Smith Normal Form of A : there exist two matrices U and V such that

$$UAV = (D|0)$$

where 0 is a $k \times (m-k)$ all-zero matrix, D is a $k \times k$ diagonal integer matrix with d_1, \dots, d_k in the main diagonal with $\prod_{i=1}^k d_i = d_k(A)$. Moreover, U and V are square unimodular integer matrices: U represents the row operations and V the column operations from A to $(D|0)$.

Let us consider $A' = (D|0)V^{-1} = UA$. As A' is built from A by integer row operations, A' is a system equivalent to A , which means that (1), (2) and (3) hold. Since V^{-1} is unimodular, it represents integer linear combinations of columns of $(D|0)$ and we can observe that all the coefficients in the i -th row of A' is a multiple of d_i , the i -th element of D .

Consider now the matrix \overline{A} obtained from A' by dividing the i -th row by d_i , for all $i \in [1, k]$. The matrix \overline{A} fulfills (2) as we can consider $\overline{A}V = (I|0)$ the Smith Normal Form of \overline{A} , so $d_k(\overline{A}) = \det(I) = 1$. To check that it also fulfills (1), consider $x = (x_1 \dots, x_m)$ a solution to $\overline{A}x = 0$, for all i , the i -th equation reads out as:

$$\overline{a}_{i,1}x_1 + \dots + \overline{a}_{i,m}x_m = 0,$$

so, multiplying by d_i :

$$d_i\overline{a}_{i,1}x_1 + \dots + d_i\overline{a}_{i,m}x_m = a'_{i,1}x_1 + \dots + a'_{i,m}x_m = d_i0 = 0.$$

Hence, x is a solution of $A'x = 0$ and, as A' is equivalent to A , it is also a solution to the homogeneous system defined by A so (1) follows.

Let \mathcal{L} be a linear equation fulfilled by the columns of A' :

$$\sum_{i=1}^m \lambda_i A'^i = 0,$$

with $\lambda_i \in \mathbb{Q}$. If we look at the j -th component, we observe that

$$0 = \sum_{i=1}^m \lambda_i a'_{j,i} = d_j \sum_{i=1}^m \lambda_i \overline{a}_{j,i},$$

hence, we have

$$\sum_{i=1}^m \lambda_i \overline{a}_{j,i} = 0,$$

for all i . Therefore \mathcal{L} is fulfilled by \overline{A} with the same coefficients. The argument is reversible: if $d_i \neq 0$ we divide by d_i , otherwise, we also have that $\sum_{i=1}^m \lambda_i a'_{j,i} = 0$. Therefore, we have $L(\overline{A}) = L(A')$ and since A' is equivalent to A as a system of equations, (3) is shown. \square

Together, Proposition 5.13 and Proposition 5.12 allow us to show Theorem 5.7 and Theorem 5.28 by proving them in the case of systems with $d_k(A) = 1$. Indeed, if A has $d_k(A) > 1$, then we first use Proposition 5.12 to obtain an equivalent, from the point of view of the group \mathbb{Z}_n^N , matrix \overline{A} . Moreover, the equations that determine the n -columns condition for A can be thought to be in \mathbb{Z} for \overline{A} , and thus with rational coefficients.

Afterwards, we apply Proposition 5.13 to obtain a new matrix \overline{A}' with $d_k(\overline{A}') = 1$ and also fulfilling the n -columns condition, as the equations that determine the n -columns conditions for \overline{A} have not changed.

Finally, we observe that the number of solutions for \overline{A}' is less than the ones for \overline{A} . However, the number of solutions for \overline{A} is the same as the number of solutions for A since both matrices are the same in \mathbb{Z}_n , and the ring operations in \mathbb{Z}_n^N just require the coefficients in the matrix to be in \mathbb{Z}_n , not in \mathbb{Z} . Summarizing: the solution set of \overline{A}' is a subset of the solution set of A , therefore, if the number of solutions is as large as we might expect for matrices with $d_k(A) = 1$, then the result for matrices with $d_k(A) > 1$ also holds.

Proof. [Proof of Theorem 5.7] Let G be an Abelian group isomorphic to \mathbb{Z}_n^N . Let A be a $k \times m$ integer matrix such that $d_k(A) = 1$ and satisfying the n -columns condition. Let $\chi : G \rightarrow [r]$ be an r -coloring. Let M be the value coming from Lemma 5.8 for which we can be sure that in any r -coloring of \mathbb{Z}_n^M there exists a monochromatic solution. By Proposition 5.11 we have that:

$$|\mathbb{Z}_n^M \subset \mathbb{Z}_n^N| \geq c_1 n^{MN},$$

and

$$\left| \mathbb{Z}_n^{(M-1)} \subset \mathbb{Z}_n^{(N-1)} \right| \leq c_2 \frac{n^{MN}}{n^N},$$

Let G_1, \dots, G_r be the partition defined by the coloring of G . Let $\delta_{\text{Theorem 4.1}} > 0$ be the value coming from Theorem 4.1 applied with: $\epsilon = \frac{c_1}{m2rc_2}$ as the portion of elements removed, $X_1 = X_2 = \dots = X_m = G_i$ for each $i \in [1, r]$ as the sets in which we are considering the solutions –the monochromatic solutions–, A as the matrix defining the system and G as the ambient group. Let G'_i be the set of elements removed from G_i .

As we want to apply Lemma 5.8, we are interested in solutions where all the elements have order n . Therefore, we can consider that $\bigcup_{i=1}^r G'_i$ contains only elements of order n .

Let $a \in \bigcup_{i=1}^r G'_i$ be a removed element. The number of subgroups isomorphic to \mathbb{Z}_n^M that contains a is the same as the number of subgroups inside $G/\langle a \rangle \cong \mathbb{Z}_n^{N-1}$ isomorphic to \mathbb{Z}_n^{M-1} . The reason is that, given a generic $\overline{G} \cong \mathbb{Z}_n^M$ and $a \in \overline{G}$ with order n , then we can use Proposition 5.10 to $\overline{G}/\langle a \rangle$ and obtain a subgroup isomorphic to \mathbb{Z}_n^{M-1} ($\{a\}$ can be extended to a subset of M elements that generate \overline{G}). Therefore, the number of subgroups isomorphic to \mathbb{Z}_n^M that survive after removing all the elements in $\bigcup_{i=1}^r G'_i$ is, at least,

$$c_1 n^{MN} - mr\epsilon n^N c_2 \frac{n^{MN}}{n^N} \geq c_1 n^{MN} - r \frac{c_1}{2rc_2} n^N c_2 \frac{n^{MN}}{n^N} = \frac{c_1}{2} n^{MN}$$

subgroups isomorphic to \mathbb{Z}_n^M inside \mathbb{Z}_n^N .

Therefore, if $\frac{c_1}{2} n^{MN}$ is larger than one, there exists a subgroup isomorphic to \mathbb{Z}_n^M inside \mathbb{Z}_n^N from which no element of order n has been removed. By Theorem 4.1, we should find a monochromatic solution. This contradicts the assumption concerning the upper bound on the number of monochromatic solutions, meaning that there are more than $\delta_{\text{Theorem 4.1}} (n^N)^{m-k}$ monochromatic solutions.

Let $n_0 > M$ is the smallest positive integer for which $\frac{c_1}{2} n_0^{MN} > 1$. By defining

$$c = \min \left(\frac{1}{2n_0^{MN}}, \delta_{\text{Theorem 4.1}} \right)$$

we obtain the desired result. Notice that, by Theorem 4.1, the dependency of c is with respect to A , not just m .

However, all the operations with A are equivalent as if A were a matrix with coefficients in \mathbb{Z}_n . Since n is finite, the number of matrices with $k \times m$ entries in \mathbb{Z}_n is bounded. Thus, the dependency of c on A can be thought as a dependency on m and n . \square

5.3.3 Proof of Lemma 5.8

In this section, we prove the Ramsey-type result, Lemma 5.8, that looks for solutions outside the finite fields. This result allows us to show Theorem 5.7 by finding more solutions than the ones coming from the results in [8] or Theorem 5.5.

The proof of Lemma 5.8 will follow the one in [8]: first we will observe that we can find solutions from a certain form.

Lemma 5.14 ([20], also in [8], for groups of the form \mathbb{Z}_n^N). *Let A be a $k \times m$ integer matrix satisfying the n -column condition. Let $G \cong \mathbb{Z}_n^N$ and let x_1, \dots, x_m be m elements in G such that $\langle x_1, \dots, x_m \rangle \cong \mathbb{Z}_n^m$. Let $F(x_1, \dots, x_m) = \{x_i + \sum_{j=i+1}^m a_{i,j}x_j : i \in \{1, \dots, m\} \text{ and each } a_{i,j} \in \mathbb{Z}_n\}$. If A fulfills the n -columns condition, then there exist elements y_1, \dots, y_m in $F(x_1, \dots, x_m)$ with $A(y_1, \dots, y_m)^T = 0$.*

Proof. Since A satisfies the n -columns condition, we assume that the columns of A are ordered in such a way that there exist $1 \leq k_1 < k_2 < \dots < k_t = m$ with:

$$(i) \quad S^1 = \sum_{j=1}^{k_1} A^j = 0,$$

$$(ii) \quad \text{for } 1 < i \leq t, \quad S^i = \sum_{j=k_{i-1}+1}^{k_i} A^j, \text{ can be expressed as a linear combination of the columns } A^1, \dots, A^{k_{i-1}} \text{ with coefficients in } \mathbb{Z}_n.$$

Let $S^i = \sum_{j=1}^{k_{i-1}} \lambda_{i,j} A^j$, with $\lambda_{i,j} \in \mathbb{Z}_n$ be the linear combination of S^i in terms of A^j . Then, we can create solutions (y_1, \dots, y_m) in $F(x_1, \dots, x_m)^m$ recursively.

We start the recursion by setting $y_i^1 = x_1$ for $i \in [1, k_1]$. Assume that $\bar{y}_i = (y_1^i, \dots, y_{k_i}^i)$ is a solution for $\bar{A}_i \bar{y}_i = 0$, where $\bar{A}_i = (A^1, \dots, A^{k_i})$ is the matrix of the first k_i columns. Then we assign $y_j^{i+1} = x_{i+1}$, for $j \in [k_i + 1, k_{i+1}]$ and we modify y_j^i to y_j^{i+1} , with $j \in [1, k_i]$, by letting $y_j^{i+1} = y_j^i - \lambda_{i+1,j} x_{i+1}$.

Let us show that $\bar{A}_{i+1} \bar{y}_{i+1} = 0$. Notice that we can separate

$$\bar{y}_{i+1} = (y_1^i, \dots, y_{k_i}^i, \overbrace{0, \dots, 0}^{k_{i+1}-k_i}) + (-\lambda_{i+1,1}x_{i+1}, \dots, -\lambda_{i+1,k_i}x_{i+1}, \overbrace{x_{i+1}, \dots, x_{i+1}}^{k_{i+1}-k_i}).$$

Therefore

$$\begin{aligned} \bar{A}_{i+1} \bar{y}_{i+1} &= \bar{A}_{i+1} (y_1^{i+1}, \dots, y_{k_i}^{i+1}, 0, \dots, 0) \\ &\quad + \bar{A}_{i+1} (-\lambda_{i+1,1}x_{i+1}, \dots, -\lambda_{i+1,k_i}x_{i+1}, \overbrace{x_{i+1}, \dots, x_{i+1}}^{k_{i+1}-k_i}). \end{aligned}$$

However, both summands are zero; the first one by induction –the column property behaves well under induction–, and the second is zero because of the n -column property: we can express the sum of the last $k_{i+1} - k_i$ columns of \bar{A}_{i+1} using the first k_i elements by means of the λ 's.

Notice that, when the induction finishes, we obtain an element of $F(x_1, \dots, x_t)^m$ where t is the number of classes of the partition of the columns. \square

Definition 5.15 (Echeloned generators). *Let x_1, \dots, x_s be elements of $G \cong \mathbb{Z}_n^M$ such that $\langle x_1, \dots, x_s \rangle \cong \mathbb{Z}_n^s$ and let $\mathcal{B} = \mathcal{B}(\mathbb{Z}_n^M)$ be a base: a set of M elements of order n in \mathbb{Z}_n^M such*

that $\langle \mathcal{B} \rangle = \mathbb{Z}_n^M$. Then $\{x_i\}_{i \in [1, s]}$ are said to be echeloned generators in \mathcal{B} or, echeloned for short, if x_i can be written, in the base \mathcal{B} as:

$$x_i = (0, \dots, 0, \overbrace{1}^{j_i}, \dots)_\mathcal{B}$$

with $j_\alpha > j_\beta$ if and only if $\alpha > \beta$.

Definition 5.16 ((t, n) -skeleton). Let G be an Abelian group isomorphic to \mathbb{Z}_n^N . Let \mathcal{B} be a base for G . Let $\{x_1, \dots, x_t\}$ be echeloned generators of a $G_0 \cong \mathbb{Z}_n^t$ then the (t, n) -skeleton of $\{x_1, \dots, x_t\}$ is the set of cyclic subgroups generated by y , $y \in \{x_i + \sum_{j=i+1}^t a_{i,j} x_j : i \in \{1, \dots, t\} \text{ and each } a_{i,j} \in \mathbb{Z}_n\}$. The parameter t is said to be the dimension of the skeleton.

Lemma 5.17 (Graham-Leeb-Rotschild for groups). For any r, s, n , there exists an $MG(r, s, n)$ such that, for any coloring of the cyclic subgroups $X \cong \mathbb{Z}_n$ of $\mathbb{Z}_n^{MG(r, s, n)}$, using r colors, there exists a monochromatic (s, n) -skeleton.

The proof is an adaptation of the one found in Graham-Rotschild-Spencer's book Ramsey Theory [40], which is a simplification of Spencer's proof [72] of the Graham-Leeb-Rotchild's result [39]. We use a similar notation as in [40]. We also follow the same scheme: first we prove the affine version, and then we show the desired version.

Lemma 5.18 (Affine version). For any r, s, n , there exists an $MGA(r, s, n)$ such that, for any coloring of the affine cyclic subgroups $y_0 + X$ of $\mathbb{Z}_n^{MGA(r, s, n)}$ using r colors, with $X \cong \mathbb{Z}_n$ and $y_0 \in \mathbb{Z}_n^{MGA(r, s, n)}$, there exist a point $x_0 \in \mathbb{Z}_n^{MGA(r, s, n)}$ and an (s, n) -skeleton S , such that all $x_0 + y$, $y \in S$, are monochromatic.

Using the notation from Graham-Rotschild-Spencer's book, we denote by $[V]^1$ the 1-translates of $V \cong \mathbb{Z}_n^u$; this is, the affine subgroups of dimension one: $[V]^1 = \{a + C \mid a \in V, C \cong \mathbb{Z}_n \text{ and } a + C \subset V\}$. More generally, we can define the i -translate of V by $[V]^i = \{a + C \mid a \in V, C \cong \mathbb{Z}_n^i \text{ and } a + C \subset V\}$.

Let B be a $(u + 1)$ -translate inside \mathbb{Z}_n^l , for some l , $B \in [\mathbb{Z}_n^l]^{u+1}$, and let $\bar{p} : B \rightarrow \mathbb{Z}_n^u$ be a surjective projection. Let T be a 1-translate of B , $T \in [B]^1$, then $\bar{p}(T)$ can be either a 1-translate, a 0-translate (a point) or something in between –the group part of the affine subgroup could be a strict subgroup of \mathbb{Z}_n –. If $\bar{p}(T)$ is a 1-translate, we say that T is transverse with respect to \bar{p} ; if $\bar{p}(T)$ is a 0-translate, we call T vertical with respect to \bar{p} , and if $\bar{p}(T)$ is none of them, we call it degenerated.

Definition 5.19. Let p be a projection from a $(u + 1)$ -translate B to \mathbb{Z}_n^u . A coloring $\chi : [B]^1 \rightarrow [r]$ is called special, relative to χ and p , if the color of a transverse 1-translate is determined by its projection. More formally, if T_1 and T_2 are two 1-translates in B such that $p(T_1) = p(T_2)$, then $\chi(T_1) = \chi(T_2)$.

Lemma 5.20. For all u, r , there exists $w = W(u : r)$ with the following property. Fix $p : \mathbb{Z}_n^{u+w} \rightarrow \mathbb{Z}_n^u$, the projection onto the first u coordinates. For any coloring $\chi : [\mathbb{Z}_n^{u+w}]^1 \rightarrow [r]$, there exists a $(u + 1)$ -translate B special with respect to p and χ .

Recall the definition of a combinatorial line. We define C_t^n the n -cube over t elements by

$$C_t^n = \{(x_1, \dots, x_n) : x_i \in \{0, 1, \dots, t - 1\}\}.$$

A *combinatorial line* in C_t^n is a t -set of suitably ordered points in the cube C_t^n , x_0, \dots, x_{t-1} , $x_i = (x_{i,1}, \dots, x_{i,n})$ so that, in each coordinate j , $1 \leq j \leq n$, either

$$x_{0,j} = x_{1,j} = \dots = x_{t-1,j}$$

or

$$x_{s,j} = s \text{ for } 0 \leq s < t,$$

and the latter occurs for at least one j . Let us recall the Hales-Jewett Theorem [44].

Theorem 5.21 (Hales-Jewett Theorem [44]). *For all r, t there exists $HJ(t, r) \in \mathbb{Z}^+$ so that, for $N \geq HJ(t, r)$ the following holds: if the vertices of C_t^N are r -colored, there exists a monochromatic line.*

Now, we proceed to the proof of Lemma 5.20.

Proof. [Proof of Lemma 5.20] Let F_u denote the family of u -variable affine linear functions: $f(x_1, \dots, x_u) = c_0 + c_1x_1 + \dots + c_u x_u$ with $c_0, c_1, \dots, c_u \in \mathbb{Z}_n$. We proof the lemma for

$$w = HJ(|F_u|, r^v),$$

where v is the number of 1-subtranslates of a u -translate, and HJ is the Hales-Jewett function. Fix $\chi : [\mathbb{Z}_n^{u+w}]^1 \rightarrow [r]$.

Let $\bar{f} = (f_1, \dots, f_w)$, $f_i \in F_u$. Define the lifting

$$\bar{f} : \mathbb{Z}_n^u \rightarrow \mathbb{Z}_n^{u+w}$$

by

$$\bar{f}(x_1, \dots, x_u) = (x_1, \dots, x_u, y_1, \dots, y_w), \quad y_i = f_i(x_1, \dots, x_u).$$

We have that \bar{f} is injective and it is also linear –indeed, affine, as the c_0 's are fixed–. Moreover, it is inverse to p in the sense that $p(\bar{f}(x_1, \dots, x_u)) = (x_1, \dots, x_u)$. We define (and this is the critical step) a coloring χ' on $(F_u)^w$ by

$$\chi'(\bar{f}) = \chi'(\bar{g}) \text{ iff, for all } T \in [\mathbb{Z}_n^u]^1, \chi(\bar{f}(T)) = \chi(\bar{g}(T)).$$

This is, color the lifting \bar{f} by the coloring of the range $\bar{f}(\mathbb{Z}_n^u)$ of the lift: if T_1, \dots, T_v are the 1-translates of \mathbb{Z}_n^u , then the color of $\chi'(\bar{f})$ is $(\chi(\bar{f}(T_1)), \dots, \chi(\bar{f}(T_v)))$. This coloring is well defined as $\bar{f}(T_i)$ is a 1-translate of \mathbb{Z}_n^{u+w} .

As χ' is an r^v -coloring, there exists a combinatorial line L in $(F_u)^w$ monochromatic under χ' –here is when we use the Hales-Jewett theorem–. By renumbering coordinates, we may write

$$L = \{(f, \dots, f, f_{\eta+1}, \dots, f_w) : f \in F_u\}, \quad (5.4)$$

where $f_{\eta+1}, \dots, f_w$ are fixed. We set

$$\begin{aligned} B &= \bigcup_{\bar{f} \in L} \bar{f}(\mathbb{Z}_n^u) \\ &= \{(x_1, \dots, x_u, y_1, \dots, y_w) : y_i = y_1, 2 \leq i \leq \eta, \\ &\quad y_i = f_i(x_1, \dots, x_u), \eta < i \leq w\}. \end{aligned}$$

B is the desired $(u+1)$ -translate.

Claim 5.22. *B is a $(u+1)$ -translate.*

Proof. [Proof of Claim 5.22] For B to be a $(u+1)$ -translate of the form $B = b + C$, with $b \in \mathbb{Z}_n^{u+w}$ and $C \cong \mathbb{Z}_n^{u+1} \subset \mathbb{Z}_n^{u+w}$, we can exhibit a b and $u+1$ elements from \mathbb{Z}_n^{u+w} that generate C . Let us denote those generators by x_i , $i \in [1, u+1]$. Let e_i be the element in \mathbb{Z}_n^u with a 1 in the i -th coordinate and 0 in the other coordinates. Then

$$x_i = (\overbrace{0, \dots, 0, 1, 0, \dots, 0}^u, \overbrace{0, \dots, 0}^\eta, f_{\eta+1}(e_i) - c_0^{\eta+1}, \dots, f_w(e_i) - c_0^w),$$

for $i \in [1, u]$. Also,

$$x_{u+1} = (\overbrace{0, \dots, 0}^u, \overbrace{1, \dots, 1}^\eta, 0, \dots, 0).$$

Finally,

$$b = (\overbrace{0, \dots, 0}^{u+\eta}, c_0^{\eta+1}, \dots, c_0^w).$$

Together, b and $\{x_i\}_{i \in [1, u+1]}$ configure a set of generators. \square

Observation 5.23. *If $y_1 \dots, y_t$ are echeloned generators of a subgroup isomorphic to $\mathbb{Z}_n^t \subset \mathbb{Z}_n^u$, then $y_1 \dots, y_t, x_{u+1}$ form also an echeloned set of generators of a subgroup isomorphic to $\mathbb{Z}_n^{t+1} \subset \mathbb{Z}_n^{u+w}$. This also holds without reordering the last w coordinates of \mathbb{Z}_n^{u+w} to better exhibit the combinatorial line as done in (5.4).*

If T is a transverse 1-translate $T \subset \mathbb{Z}_n^{u+w}$, it can be written as $T = \bar{g}(p(T))$, for some \bar{g} . Indeed, assume $T = t_0 + C$, with $C \cong \mathbb{Z}_n$, $t_0 \in \mathbb{Z}_n^{u+w}$. Since p is linear, we have that $p(T) = p(t_0) + p(C)$. As T is transverse, $p(C) \cong \mathbb{Z}_n$. A general \bar{g} can be written as $\bar{g} = (g_1, \dots, g_w)$ with $g_i = c_0^i + c_1^i x_1 + \dots + c_u^i x_u$. If we let $t_0 = (t_0^1, \dots, t_0^u, t_0^{u+1}, \dots, t_0^w)$, we have that $c_0^i = t_0^{u+i}$, which solves the part for the t_0 . As for C , we look for a generator γ of C where some of the first u coordinates, say $\alpha_1, \dots, \alpha_l$, are such that $\gcd(\gamma_{\alpha_1}, \dots, \gamma_{\alpha_l}) = 1$: since T is transverse then we have that $p(C) \cong \mathbb{Z}_n$, therefore the first u coordinates of γ should generate a cyclic group of order n , which implies that such coordinates, $\alpha_1, \dots, \alpha_l$, have to exist. Therefore, $p(\gamma)$ will also have the same coordinates, $\alpha_1, \dots, \alpha_l$, with greatest common divisor equal 1. Using Bezout's identity, let β_1, \dots, β_l be l integers such that $\beta_1 \gamma_{\alpha_1} + \dots + \beta_l \gamma_{\alpha_l} = 1$, in particular, if we consider them modulo n we obtain the same equality. By setting $c_{\alpha_j}^i = \gamma_{u+i} \beta_j$, for $j \in [1, l]$, where γ_{u+i} is the $(u+i)$ -th coordinate of γ , and $c_j^i = 0$ for all $j \notin \{\alpha_i\}_{i \in [1, l]} \cup \{0\}$, we obtain the desired result, as C is generated by this element.

When the transverse 1-translate T is such that $T \subset B$, then $T = \bar{g}(p(T)) = \bar{f}(p(T))$, where $\bar{f} = (g_1, \dots, g_1, f_{\eta+1}, \dots, f_w) \in L$. Indeed, as we already know that there exists a \bar{g} such that $T = \bar{g}(p(T))$, the only thing left to show is that $\bar{g} \in L$. As the last coordinates are completely determined by $f_{\eta+1}, \dots, f_w$, we can set $g_{\eta+1} = f_{\eta+1}, \dots, g_w = f_w$. Since the first η coordinates added are exactly the same in T as the first one T_{u+1} , the equalities $g_1 = g_2 = \dots = g_\eta$ have to be fulfilled, hence we can choose $\bar{g} = \bar{f} = (g_1, \dots, g_1, f_{\eta+1}, \dots, f_w) \in L$.

Let $T_1, T_2 \in [B]^1$ be such that $p(T_1) = p(T_2) = T \in [\mathbb{Z}_n^u]^1$. Then $T_1 = \bar{f}_1(p(T_1))$ and $T_2 = \bar{f}_2(p(T_2))$ with $\bar{f}_1, \bar{f}_2 \in L$. Hence

$$\chi(T_1) \stackrel{T_1 = \bar{f}_1(p(T_1))}{=} \chi(\bar{f}_1(p(T_1))) \stackrel{T = p(T_1)}{=} \chi(\bar{f}_1(T)) \stackrel{\chi' \text{ mon. in } L}{=} \chi(\bar{f}_2(T)) = \chi(\bar{f}_2(p(T_2))) = \chi(T_2),$$

and B is special with respect to p and χ . \square

After proving the technical Lemma 5.20, we are almost ready to prove Lemma 5.18, again we follow the lines of [40]. Recall the Extended Hales-Jewett theorem, which follows from Hales-Jewett Theorem by a simple trick.

Theorem 5.24 (Extended Hales-Jewett). *For all n, t, r there is an $EHJ(n, t, r)$ so that, for $N \geq EHJ(n, t, r)$, the following holds: if the points of C_t^N are r -colored, there exists a monochromatic n -dimensional combinatorial subspace.*

Observation 5.25 (Extended Hales-Jewett and skeletons). *The Extended Hales-Jewett theorem and the skeleton structure are compatible.*

Proof. [Proof of Observation 5.25] Let (k_1, \dots, k_r) be an r -tuple and let ξ be the value coming from the Extended Hales-Jewett with $\xi = EHJ(n, \max_{i \in [1, r]}(k_i), r)$. Thus, for any (k_1, \dots, k_r) , if we paint the elements of \mathbb{Z}_n^ξ using r colors, we can ask for a monochromatic translate of the group isomorphic to $\mathbb{Z}_n^{k_s}$, name it \overline{G}_{k_s} , for some $s \in [1, r]$. Moreover, if $\mathcal{B} = \{e_i\}_{i \in [1, \xi]}$ is the canonical base of \mathbb{Z}_n^ξ with

$$e_i = (0, \dots, 0, \overbrace{1}^i, 0, \dots, 0),$$

then there exist a point x_0 and k_s nonempty and pairwise disjoint subsets of $[1, \xi]$, $\{I_j\}_{j \in [1, k_s]}$, such that: $\overline{G}_{k_s} = x_0 + \langle y_1, \dots, y_{k_s} \rangle$ with

$$y_j^i = \begin{cases} 1 & \text{if } i \in I_j, \\ 0 & \text{otherwise} \end{cases},$$

where y_j^i denotes the i -th coordinate of y_j . Even more, the coordinates $x_0^t = 0$, for $t \in \cup_{j \in [1, k_s]} I_j$. In particular, if \mathcal{K} is an affine skeleton in \overline{G}_{k_s} with respect to the base (y_1, \dots, y_{k_s}) , then \mathcal{K} is an affine skeleton in \mathbb{Z}_n^ξ with respect to the canonical base. \square

Proof. [Proof of Lemma 5.18] We prove this result by doing an induction on the dimension of each color's larger monochromatic affine skeleton: instead of the original result, let us show a slight generalized Lemma 5.18. For all k_1, \dots, k_r , there exists an $\text{mga} = MGA(k_1, \dots, k_r)$ so that: if the 1-translates of $\mathbb{Z}_n^{\text{mga}}$ are r -colored, there exists, for some $1 \leq i \leq r$, a (k_i, n) -skeleton colored i . The proof uses be a double induction: first on the dimension of the painted translates –from points to 1-translates–, and then on the dimensions of the colored affine skeletons (k_1, \dots, k_r) .

Denote by $N' = N'(k_1^0, \dots, k_r^0, k_1^1, \dots, k_r^1)$ a number for which, given $N \geq N'$, if we paint the 0-translates and the 1-translates of \mathbb{Z}_n^N with r colors, then:

- for some i , there exists an affine subgroup of dimension k_i^0 where all its elements have color i .
- for some j , there exists a monochromatic (k_j^1, n) -skeleton.

The induction is the following: assume that N' exist for all (k_1^0, \dots, k_r^0) –fixing $(\overline{k}_1^1, \dots, \overline{k}_r^1)$ – and for all $(\overline{k}_1^1, \dots, \overline{k}_r^1) < (k_1^1, \dots, k_r^1)$, then we want to show that there exists an $N' = N'(k_1^0, \dots, k_r^0, k_1^1, \dots, k_r^1)$. Using the Extended Hales-Jewett, this induction is well based as we can always find large monochromatic affine subgroups (painting the elements) for any given $(\overline{k}_1^1, \dots, \overline{k}_r^1)$. Thus, we just have to concentrate our efforts on showing

the induction step: if we have found a monochromatic $((\bar{k}_1^1, \dots, \bar{k}_r^1), n)$ -skeleton for all $(\bar{k}_1^1, \dots, \bar{k}_r^1) < (k_1^1, \dots, k_r^1)$, we want to show that it is possible to find a monochromatic $((k_1^1, \dots, k_r^1), n)$ -skeleton –at the expense of, maybe, increase the size of the group \mathbb{Z}_n^\bullet .

We set

$$\begin{aligned} s &= \max_{1 \leq i \leq r} MGA(k_1, \dots, k_i - 1, \dots, k_r), \\ u &= EHJ(s, n, r), \\ w &= W(u : r), \\ N &= u + w, \end{aligned} \tag{5.5}$$

where $EHJ(s, n, r)$ is the function given by the Extended Hales-Jewett being \mathbb{Z}_n the base-space, we paint the elements of \mathbb{Z}_n^u with r colors and we aim to find an s -translate. $W(u : r)$ is the output of Lemma 5.20 about the existence of an special projection from $\mathbb{Z}_n^{u+W(u:r)}$ to \mathbb{Z}_n^u .

Color the 1-translates from \mathbb{Z}_n^N arbitrarily by the coloring function χ . By the definition of w , using Lemma 5.20, there is a $(u + 1)$ -translate B that is special under a projection $p : B \rightarrow \mathbb{Z}_n^u$. Induce a coloring χ' of the elements in \mathbb{Z}_n^u , the 0-translates, by $\chi'(T) = \chi(p^{-1}(T))$, where $p^{-1}(T)$ is the unique vertical 1-translate in B that collapses completely onto T .

By the definition of u , there exists an s -translate $S \subset \mathbb{Z}_n^u$ monochromatic, say of color 1, under χ' . Then $p^{-1}(S) \subset B$, thought as the pre-images of single points in S , is an special $(s + 1)$ -translate where all the vertical 1-translates are colored 1. We define a coloring χ'' in $[S]^1$ by

$$\chi''(T) = \chi(T'), \text{ where } p(T') = T, T' \text{ transversal, } T' \in B \text{ with } p(T') \in [S]^1.$$

Since p is special, this is a well defined coloring. As $s \leq MGA(k_1 - 1, k_2, \dots, k_r)$ (we apply induction on (k_1, \dots, k_r)), there exists an skeleton $\mathcal{W}' \subseteq S$ so that either

$$(i) \dim(\mathcal{W}') = k_1 - 1; \mathcal{W}' \text{ has color 1 under } \chi'',$$

or

$$(ii) 2 \leq i \leq r, \dim(\mathcal{W}') = k_i; \mathcal{W}' \text{ has color } i \text{ under } \chi''.$$

In case (ii), there exists a k_i -translate $\mathcal{W} \subset p^{-1}(\mathcal{W}')$ such that $p(\mathcal{W}) = \mathcal{W}'$. If $\mathcal{W} = v + C$, then we can choose $\mathcal{W} = p^{-1}(v) + C$, for some pre-image of v under p , as $p(C) = C$ whenever $C \subset \mathbb{Z}_n^u$. \mathcal{W} is an affine (k_i, n) -skeleton monochromatic with color i under the coloring χ in B , thus it is a monochromatic affine (k_i, n) -skeleton in the whole group. We use Observation 5.25 to check that the generators of the skeleton are, indeed, echeloned.

In case (i), we construct the set inductively, we set $\mathcal{W} = p^{-1}(\mathcal{W}')$. We can use the following observation:

Let $\mathcal{W}' = \overline{x_0} + \langle \overline{x_1}, \dots, \overline{x_k} \rangle$ be an affine skeleton, where $\overline{x_0}, \overline{x_1}, \dots, \overline{x_k} \in \mathbb{Z}_n^u$ and $\langle \overline{x_1}, \dots, \overline{x_k} \rangle \cong \mathbb{Z}_n^k$, and the $\langle \cdot \rangle$ means that we generate only the skeleton, not all the cyclic groups of length n . By reordering the w added coordinates, we can write B as

$$\begin{aligned} B &= \bigcup_{\overline{f} \in L} \overline{f}(\mathbb{Z}_n^u) \\ &= \{(x_1, \dots, x_u, y_1, \dots, y_w) : y_i = y_1, 2 \leq i \leq \eta, \\ &\quad y_i = f_i(x_1, \dots, x_u), \eta < i \leq w\}, \end{aligned}$$

where $f_i(x_1, \dots, x_u) = c_0^i + c_1^i x_1 + \dots + c_u^i x_u$. If $p : B \rightarrow \mathbb{Z}_n^u$, then $\mathcal{W} = p^{-1}(\mathcal{W}')$ is a $(k+1, n)$ -skeleton with the translation point

$$\overline{x}_0' = (x_0^1, \dots, x_0^u, \overbrace{0, \dots, 0}^\eta, f_{\eta+1}(\overline{x}_0), \dots, f_w(\overline{x}_0)),$$

and with generating vectors x_1', \dots, x_{k+1}' such that

$$\overline{x}_i' = (x_i^1, \dots, x_i^u, \overbrace{0, \dots, 0}^\eta, f_{\eta+1}(\overline{x}_i) - c_0^{\eta+1}, f_{\eta+2}(\overline{x}_i) - c_0^{\eta+2}, \dots, f_w(\overline{x}_i) - c_0^w),$$

for $i \in [1, k]$, and

$$\overline{x}_{k+1}' = (\overbrace{0, \dots, 0}^u, \overbrace{1, \dots, 1}^\eta, \overbrace{0, \dots, 0}^{w-\eta}).$$

Observe that all the cyclic subgroups generated by $\Omega = \{\overline{x}_i' + \sum_{j=i+1}^{k+1} a_{i,j} \overline{x}_j' : i \in \{1, \dots, k+1\} \text{ and each } a_{i,j} \in \mathbb{Z}_n\}$ and translated by \overline{x}_0' , the ones forming the skeleton inside \mathcal{W} , are either vertical or transverse. Notice that if $i \in [1, k]$, then $p(\overline{x}_i' + \sum_{j=i+1}^{k+1} a_{i,j} \overline{x}_j') = \overline{x}_i + \sum_{j=i+1}^k a_{i,j} \overline{x}_j$ which is of order $|\mathbb{Z}_n|$, hence $x_0' + \langle \overline{x}_i' + \sum_{j=i+1}^{k+1} a_{i,j} \overline{x}_j' \rangle$ is transverse and, if $i = k+1$, then $p(\overline{x}_{k+1}') = 0$, which has order 0, thus $\overline{x}_0' + \langle \overline{x}_{k+1}' \rangle$ is vertical.

Moreover, the generators $\overline{x}_1', \dots, \overline{x}_k', \overline{x}_{k+1}'$ are echeloned: although the coordinates have been reordered, if $\overline{x}_1, \dots, \overline{x}_k$ were echeloned in \mathcal{W}' , then $\overline{x}_1', \dots, \overline{x}_k'$ are also echeloned in \mathcal{W} . In addition, \overline{x}_{k+1}' is also echeloned with respect to $\overline{x}_1', \dots, \overline{x}_k'$. Therefore, we have seen that \mathcal{W} is, indeed, a $(k+1, n)$ -skeleton. \mathcal{W} is monochromatic with color 1 since, if T is a 1-translate in the skeleton and it is vertical, then $\chi(T) = \chi'(p(T)) = 1$ because $p(T)$ is a 0-translate in S . If T is in the skeleton and is a transverse 1-translate, then $\chi(T) = \chi(p(T)) = 1$ by (i). As all the 1-translates in the skeleton \mathcal{W} are either vertical or transverse, we have shown the induction step and finished the proof. \square

Once Lemma 5.18 has been shown, we proceed to the proof of Lemma 5.17.

Proof. [Proof of Lemma 5.17] Let χ be a coloring of the cyclic subgroups and let $MG(r, s, n) = MGA(r, s, n)$. We paint the 1-translates of \mathbb{Z}_n^{MG} using the color of the associated subgroup. By Lemma 5.18, there exists a monochromatic affine (s, n) -skeleton, hence, the associated (s, n) -skeleton structure has to be monochromatic. Thus, Lemma 5.17 is shown. \square

At this point, we follow the strategy in [8] to prove Lemma 5.8.

Proof. [Proof of Lemma 5.8] Let $N = MG(r, m, n)$ and let $\chi : \mathbb{Z}_n^N \rightarrow [r]$ be a coloring of the elements of \mathbb{Z}_n^N with r colors. Choose in \mathbb{Z}_n^N the canonical base \mathcal{B} and say that $(a_1, \dots, a_N) < (b_1, \dots, b_N)$ if there exists a $j \geq 1$ such that $a_i = b_i$ for $i \in [1, j-1]$ and $a_j < b_j$, considering the coordinates in $[0, n-1]$.

Let T be a cyclic group. Let γ be the generator of T with minimum coordinates. We define a coloring on the cyclic groups isomorphic to \mathbb{Z}_n , χ' , by $\chi'(T) = \chi(\gamma)$. We apply Lemma 5.17 to know that there exist x_1, \dots, x_m echeloned with respect to \mathcal{B} such that all the cyclic subgroups generated by $\{x_i + \sum_{j=i+1}^m a_{i,j} x_j, a_{i,j} \in \mathbb{Z}_n\}$ are monochromatic. Now we apply Lemma 5.14 to show that there exists a collection $\{y_i\}_{i \in [1, m]}$ with $y_i \in \{x_i + \sum_{j=i+1}^m a_{i,j} x_j, a_{i,j} \in \mathbb{Z}_n\}$ such that $y = (y_1, \dots, y_m)$ is a solution of $Ay = 0$.

Since $x_i, i \in [1, m]$ are echeloned, if $y_i \in \{x_i + \sum_{j=i+1}^k a_{i,j} x_j, a_{i,j} \in \mathbb{Z}_n\}$, then $\chi'(\langle y_i \rangle) = \chi(y_i)$, therefore, y is a monochromatic solution. \square

5.4 Extension for bounded torsion groups

In this section we combine Theorem 5.7 and Theorem 5.9 to show that the number of monochromatic solutions of an homogeneous linear system, for finite Abelian groups with bounded order, is as large as we might expect, Theorem 5.26.

Theorem 5.26 (Number of solutions for bounded torsion groups). *Let b be a positive integer. Let G be a finite Abelian group of order n such that $\max_{g \in G} (|g|) = b$. Let r be a positive integer and let A be a $(k \times m)$ matrix with integer entries. Assume that A satisfies the b -columns condition. There is a constant $c = c(r, b, m) > 0$ such that every r -coloring of $G \setminus \{0\}$ has at least $\lfloor c|G|^{m-k} \rfloor$ monochromatic solutions of the equation $Ax = 0$ in $G \setminus \{0\}$.*

Proof. [Proof of Theorem 5.26] Let r be the number of colors. Let $\text{div}(b) = \{\mu \mid \mu \text{ divides } b\} = \{\mu_i\}_{i \in [1, |\text{div}(b)|]}$. Let M_{μ_i} be the output of Lemma 5.8 to ensure that $\mathbb{Z}_{\mu_i}^{M_{\mu_i}}$ contains a monochromatic solution, then set $\overline{M}_b = \max_{i \in [1, |\text{div}(b)|]} \{M_{\mu_i}\}$. Let $\epsilon^{-1} = \prod_{i=1}^{|\text{div}(b)|} \mu_i^{\overline{M}_b} = \epsilon(r, b, A)$. If the order of the Abelian group G is large enough, we can be sure that G contains a $\mathbb{Z}_{\mu_i}^{M_{\mu_i}}$, for some μ_i divisor of b .

Let $\overline{\mu} = \max\{\mu_i \mid \mathbb{Z}_{\mu_i}^{M_{\mu_i}} \subset G\}$ and consider $G_{\overline{\mu}} = \mathbb{Z}_{\overline{\mu}}^{\overline{M}_b}$. Let G_0 be the maximal subgroup in G such that $G_{\overline{\mu}} \subset G_0$ and with $\max_{g \in G_0} |g| = \overline{\mu}$, then

$$G_0 \cong \prod_{i=1}^s \mathbb{Z}_{n_i}$$

for some s , with $n_s = \overline{\mu}$ and $n_1 | n_2 | \cdots | n_s$.

Notice that $|G_0| > \epsilon|G|$. Indeed, if

$$G \cong \prod_{i=1}^h \mathbb{Z}_{m_i}$$

for some h , with $m_h = b$ and $m_1 | m_2 | \cdots | m_h$, then the subgroup G_0 can be found in G by picking, for each factor \mathbb{Z}_{m_i} in G , the subgroup $\mathbb{Z}_{\text{gcd}(m_i, \overline{\mu})}$ so that:

$$G_0 \cong \prod_{i=1}^h \mathbb{Z}_{\text{gcd}(m_i, \overline{\mu})}.$$

Thus, it can be checked that $|G_0| > \epsilon|G|$ as, for all m with $m|b$, we should have, either $\mathbb{Z}_m^{\overline{M}_b} \subset G$, and then $m|\overline{\mu}$, or $\mathbb{Z}_m^{\overline{M}_b} \not\subset G$, and then m contribute less than the $|\mathbb{Z}_m^{\overline{M}_b}|$ to the shrinking of G_0 with respect to G .

We can successfully apply the same machinery to G_0 as we have done for Theorem 5.7 to find a positive proportion of monochromatic solutions in G_0 .

First, assuming there are not many monochromatic solutions in G_0 , we apply the removal lemma to delete not many elements of order $\overline{\mu}$ in G_0 . However, there is a significant amount of subgroups isomorphic to $\mathbb{Z}_{\overline{\mu}}^{\overline{M}_b}$ in G_0 by Theorem 5.9 and Proposition 5.11. Since each time we remove an element of order $\overline{\mu}$, we do not puncture many subgroups, then there exist a subgroup $\mathbb{Z}_{\overline{\mu}}^{\overline{M}_b}$ with no elements removed. In this group with no elements of order $\overline{\mu}$ removed, we use Lemma 5.8 to find a monochromatic solution were all the x_i have order $\overline{\mu}$

inside G_0 . This reaches a contradiction with the assumption of the removal lemma, hence we find at least $\lfloor c|G|^{m-k} \rfloor$ monochromatic solutions, for some $c = c(r, b, m)$. Again c only depends on m as the number of $k \times m$ matrices on \mathbb{Z}_b is bounded. \square

Observe that we could also ask for groups G where the orders of the elements of G have an absolute bound b , but this b might not be attained. Due to the variety of groups included in this notion, we have to restrict ourselves to systems that satisfy, roughly, the $b!$ -columns condition. To be more precise, we might just want to multiply the largest primer-powers that divide b .

5.5 Density case

In this section we show the characterization of integer matrices such that, for all finite Abelian groups and for all sets with positive density, they contain, asymptotically, as many solutions as we might expect. This result is similar to the version of Varnavides [78] of the Szemerédi Theorem [74] on arbitrarily long arithmetic progressions in dense sets of the integers: there is a constant times N^2 k -arithmetic progressions in any dense set of $[1, N]$.

We say that a $(k \times m)$ matrix A with integer coefficients and $m \geq k + 2$ is *density regular* if, for every $\epsilon > 0$ there is $n(\epsilon) \in \mathbb{N}$ such that the following holds: for every Abelian group G of order $n \geq n(\epsilon)$ and every subset $X \subset G$ such that $|X| \geq \epsilon n$, there is a nontrivial solution of the homogeneous linear system $Ax = 0$ with all coordinates in X . Here by trivial solution we mean one with all coordinates equal to the same common value.

In the terminology of Rado's characterization of partition regular matrices, we say that the $k \times m$ integer matrix A , with $m \geq k + 2$, verifies the *strong column* condition if the sum of the columns is the zero vector in \mathbb{Z}^k . Our main result is the following:

Theorem 5.27. *A matrix is density regular if and only if it satisfies the strong column condition.*

In particular we show:

Theorem 5.28 (Counting for dense sets). *Let A be a $k \times m$ integer matrix. For every $\epsilon > 0$, there exists a $\delta = \delta(\epsilon, A) > 0$ such that: for every finite Abelian group G and for every set $X \subset G$ with $|X| \geq \epsilon|G|$, $Ax = 0$ has $\lfloor \delta|G|^{m-k} \rfloor$ solutions with $x \in X^m$ if and only if A satisfy the strong column condition.*

The proof of Theorem 5.28 is similar as the one in Theorem 5.7: we use Theorem 4.1, the fact that each element in $s \in X$ bears a trivial solution $x = (s, \dots, s)$ and Proposition 5.13 to remove the technical condition of the determinantal. For the “if” part, we give a sequence of groups G_i and associated dense subsets X_i such that, if A does not satisfy the strong columns condition, then there is no solution at all to $Ax = 0$, $x \in X_i^k$.

Proof. [Proof of Theorem 5.28] Assume A satisfies the strong column condition. Assume also that $d_k(A) = 1$. Let G be an Abelian group. Let $\epsilon > 0$ be a real number and let X be a set with density larger than ϵ in G . Then we can find the so-called trivial solutions to $Ax = 0$, namely $x = (x_0, \dots, x_0)$, for each $x_0 \in X$.

We use the removal lemma for Abelian groups, Theorem 4.1, to know that there exists a $\delta_{\text{Theorem 4.1}} = \delta(\epsilon m^{-1}/2, A) > 0$ such that: if there are less than $\delta_{\text{Theorem 4.1}}|G|^{m-k}$

solutions to $Ax = 0$, $x \in X^m$, then we can destroy all these solutions by remove at most $\epsilon/2|G|$ elements from X . However, as we have not removed all the elements from X , there are, still, some trivial solutions. Therefore, the total number of solutions has to be larger than $\delta_{\text{Theorem 4.1}}|G|^{m-k}$.

Observe that the strong column condition, $\sum_{i=1}^m A^i = 0$, can be expressed as a linear combination with integer coefficients of the columns of A . Therefore, if $d_k(A) > 1$, we use Proposition 5.13 to obtain a matrix A' with $d_k(A') = 1$. Moreover, A' satisfies the strong column condition and its solution set is a subset of solution set of $Ax = 0$. Hence, we use the reasoning from the preceding paragraph with the matrix A' to check that the number of solutions is proportional to $|G|^{m-k}$, which implies that the number of solutions to $Ax = 0$, $x \in X^m$ is, at least, proportional to $|G|^{m-k}$.

For the only if part, suppose that there is one equation $a_1x_1 + \dots + a_mx_m = 0$ with $\sum_i a_i = \alpha \neq 0$. Take a sufficiently large positive integer r and consider G to be the cyclic group $\mathbb{Z}/r\mathbb{Z}$. Let $X \subset \mathbb{Z}/r\mathbb{Z}$ consists of the elements whose representatives in $[0, r]$ are congruent to 1 modulo $|\alpha| + 1$ and lie in an initial segment $[0, r_0]$, where $r_0 = r/(mt)$ and $t = \max_i |a_i|$. Thus $|X| \geq r/(mt(|\alpha| + 1))$. Every element in X^m is of the form $u' = u(|\alpha| + 1) + \mathbf{1}$, where u is an integer valued m -vector and $\mathbf{1}$ is the all ones vector. Hence, if $a = (a_1, \dots, a_m)$ and $u' \in X^m$, we have

$$a \cdot u' = (a, u)(|\alpha| + 1) + \alpha,$$

which, since (a, u) is an integer, $a \cdot u'$ cannot be equal to zero. Moreover, u' is nonzero modulo r because the elements in X are in $[1, r_0]$, so $(a, u') \in [-r + 1, r - 1]$. Thus the equation $a_1x_1 + \dots + a_mx_m = 0$ has no solutions in X . \square

5.6 Remarks

In light of Bergelson, Deuber, Hindman's [8, Theorem 2.4] and Rado's [59] characterization results, Lemma 5.8 does not characterize those systems in which, for every r -coloring, we can find a monochromatic solution in groups of the type \mathbb{Z}_n^N where all the components have order n .

Using the same techniques displayed in [8, Lemma 2.2] or in [20], we can see that if the monochromatic solution is inside an (m, n) -skeleton then the system has to fulfill the n -columns condition. However, the condition "to be inside an (m, n) -skeleton" seems too strong. Indeed, if a solution is inside an (m, n) -skeleton, then it behaves as if it would have lain in a finite field. The structure of the skeleton is the translation of the n -columns condition from the system onto the solution set.

The n -columns condition is, however, a natural generalization of the F_q -columns, which is the characterization from Bergelson, Deuber and Hindman result: if we restrict the group to be a power of a prime-order cyclic group \mathbb{Z}_p^N , both conditions become the same.

Moreover, there are weaker generalizations of the F -columns condition that do not work. For instance, if the system fulfills the column condition for all $p|n$ in the case of the group \mathbb{Z}_n^N , then there are examples of systems, groups and colorings for which no monochromatic solution with all the elements of order n is found. Hence, an stronger condition is needed.

Example: Let $n = 4$ and

$$\mathbf{A} = \begin{pmatrix} 1 & 0 & -1 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix}.$$

This matrix fulfills the 2-column condition, as the sum of the columns is 0 with the coefficients in \mathbb{Z}_2 . However, there is no solution with $x_4 \in \mathbb{Z}_4^s \setminus \mathbb{Z}_2^s$, as that variable should be such that $2 \cdot x_4 = 0 \pmod{4}$ in all the s coordinates, but this is precisely the condition for an element in \mathbb{Z}_4^s to be in the subgroup \mathbb{Z}_2^s .

To finish this subsection, let us recall Voigt's result on groups that fulfill the partition property [79]. This result generalizes Graham, Leeb and Rotschild's result [39] from finite fields to Abelian groups. With the idea of weakening the n -columns condition from Lemma 5.8, the role of Graham, Leeb and Rotschild's result in the proof could be replaced by Voigt's theorem and ideas. However, this process does not seem to be straightforward.

Final remarks and future work

In this chapter we discuss some final remarks and comments as well as some open problems and future work that have been left open in the other chapters.

In particular, Section 6.1 is devoted to a discussion of a possible removal lemma for polynomial systems and the differences of this framework with respect to the linear systems setting. Section 6.2 present infinite families of examples of combinatorial structures that, even though they resemble a linear system of equations in some aspects, they do not fulfill a removal lemma statement.

In Section 6.3 we discuss about the dimension of the hypergraph needed to represent the system using our techniques. In particular, we present an example of a linear system that, although a hypergraph is needed, it can be represented by a hypergraph with less dimension than what is required by the general construction. In Section 6.4 we briefly explain how our results fit in Szegedy's framework of the Symmetry-preserving removal lemma [73]. Finally, Section 6.5 discuss some open problems and future work related with this thesis.

6.1 Extension to Polynomials

In Chapters 2-4, we have seen how the Removal Lemma for one equation for Abelian groups can be generalized to systems of equations for finite fields and, then, to systems of equations for Abelian groups. A possible generalization could be to extend the context of applications: for example, we could ask for a removal lemma for linear systems in non-abelian finite groups, or a removal lemma for compact Abelian groups (see [15] for the result for the circle). Instead, we could rise the question of a possible removal lemma for polynomial equations or, in other words, for algebraic varieties.

Question 6.1 (Removal lemma for polynomials). *Let P_1, \dots, P_k be k polynomials on m variables, x_1, \dots, x_m of maximal degree d and let $X_i \subset \mathbb{F}_q$, with $i \in [1, m]$. Assume that*

$$\left| \left\{ (x_1, \dots, x_m) \in \prod_{i=1}^m X_i \subset \mathbb{F}_q^m : P_1(x_1, \dots, x_m) = \dots = P_k(x_1, \dots, x_m) = 0 \right\} \right| = o(q^{m-k}).$$

Are there some subsets $X'_i \subset X_i$, $i \in [1, m]$, with $|X'_i| = o(q)$ such that there is no solution to $P_1(x_1, \dots, x_m) = 0, \dots, P_r(x_1, \dots, x_m) = 0$ with $x_i \in X_i \setminus X'_i$?

We should observe that this problem has a slightly different nature than the one for linear systems of equations. In the case of a linear system A , there is an integer $N = N(A)$ and

a constant $1 \geq \delta = \delta(A) > 0$ such that, for any \mathbb{Z}_n , $n \geq N$, there are subsets X_1, \dots, X_m such that:

- $|X_i| \geq \delta n$, for all $i \in [1, m]$,
- there is no solution to $Ax = 0$ with $x_i \in X_i$.

Moreover, if the sum of the columns of A is not the zero-vector (so A does not satisfy the strong columns condition,) then we can further assume that $X_i = X_1$, for all $i \in [1, m]$.

The removal lemma can be seen as a 2-options result: either the set $\prod_{i=1}^m X_i$ has many solutions, of the order of $|G|^{m-k}$, or it is $o(|G|)$ -close to a solution-free product set. The observation above shows that these options are non-trivial in the following sense: there are plenty of large solution-free product sets.

On the other side, a similar observation does not work even for simple polynomials. For instance, Theorem 6.2 by Cilleruelo exhibits three examples of polynomial equations where the number of solutions is large for any choice of sets with positive density. Thus, the hypothesis that the polynomial equation has a small number of solutions turns to be equivalent to the case in which some of the sets are small (at least for the polynomial equations in Theorem 6.2.)

Theorem 6.2 ([18]). *Let $X_1, X_2, X_3, X_4 \subset \mathbb{F}_q^*$, then the number of solutions to any of the equations:*

$$x_1 + x_2 = (x_3 + x_4)^2, \quad x_i \in X_i, \quad (6.1)$$

or

$$x_1 x_2 = x_3 + x_4, \quad x_i \in X_i, \quad (6.2)$$

or

$$x_2 x_3 - x_1 x_4 = 1, \quad x_i \in X_i, \quad (6.3)$$

is

$$S_i = \frac{|X_1||X_2||X_3||X_4|}{q} + \theta_j \sqrt{|X_1||X_2||X_3||X_4|q},$$

where $|\theta_j| \leq 1 + o(1)$, $j \in \{1, 2, 3\}$ is a constant depending on each equation and sets, but absolutely bounded independently of q .

Let us notice that we have a similar behavior in the case of linear systems fulfilling the strong columns condition; in particular, any dense set will have many solutions. However, the polynomial equations described in Theorem 6.2 behave differently from the strong columns condition case in the following sense: if $x_1 = x_2 = x_3 = x_4$, then, in general, (x_1, x_1, x_1, x_1) is not a solution to any of the polynomial equations (6.1), (6.2) or (6.3) and, in Theorem 6.2, we can choose a different set for each variable.

The fact that the product of large arbitrary sets have large intersection with the algebraic variety defined by the system of polynomial equations, as illustrated in Theorem 6.2 for one equation, seems to be a general phenomenon [12]. However, to characterize the polynomial systems that have large solution-free sets seems to be a hard problem. Besides, although this behaviour shows that the case for linear equations is substantially different than the one for algebraic varieties, it is not clear if a removal lemma for polynomials might hold.

6.2 Orthogonal Arrays

In previous chapters we have seen some algebraic structures that fulfill a Removal Lemma-like statement. For example: graphs (undirected, directed and edge-colored), hypergraphs, or linear systems with integer coefficients over finite Abelian groups.

However, certain combinatorial structures, which can be regarded as generalizations of the solution space of a linear equation, do not fulfill a removal lemma statement.

Our example is an orthogonal array $(2, n)$. So, it is a $n^2 \times 3$ matrix where each pair of columns contains all the ordered pairs with elements in $[1, n]$. In particular: for any pair of indices of columns $(i_1, i_2) \in \binom{[1, 2, 3]}{2}$ and for any given value to $(x_{i_1}, x_{i_2}) \in [1, n]^2$, there exists a unique assignment for the third element x_{i_3} , where i_3 is the remaining index. This behavior is similar as the one found in any linear equation with three variables over some group G with n elements, and where the greatest common divisor of the coefficients of the equation is coprime with n .

6.2.1 General considerations

Before starting the construction, we make a simple observation about how some orthogonal arrays of this type can be built up.

Consider the ordered set $[1, n]$ with the natural increasing order. Furthermore, consider the triples of the orthogonal arrays to be packed in chunks, depending on the value of the first column: the i -chunk will be the collection of rows in which the element of the first column is i .

Without loss of generality we consider that each chunk is ordered lexicographically considering the first two columns: the i -chunk is ordered $(i, 1), (i, 2), \dots, (i, n)$.

Let us start the construction by picking, in the first two columns, all the ordered $[1, n]^2$. Let us pick a permutation σ of $[1, n]$ so that $\sigma(i)$ is the third element of the triple for the first chunk. We also select an n -cycle permutation c . For the second chunk we use the permutation σc ; in general, for the $(i + 1)$ -th chunk, we use the σc^i permutation as the element of the third column.

Proposition 6.3. *The construction outlined above creates an orthogonal array for any σ , c and any order of the n chunks.*

Proof. By construction, the first two columns contain all the pairs. The first and the third column fulfill the property by the same reason: as σ and c are permutations, the i -th chunk will see all the pairs (i, j) , with $j \in [1, n]$.

The second and the third columns see all the pairs because the powers of the cycle c : once $\sigma(j)$ is fixed, we perform an n -cycle permutation $i \rightarrow \sigma c^{i-1} \sigma^{-1} \sigma(j)$. Indeed, for a fixed element in the second column j , when we go through the i -chunks, we have the triples $(i, j, \sigma c^{i-1} \sigma^{-1} \sigma(j))$, with $i \in [1, n]$. If $\sigma c^{i_1-1} \sigma^{-1} \sigma(j) = \sigma c^{i_2-1} \sigma^{-1} \sigma(j)$, then $i_1 = i_2$. Thus, we do go through all the pairs (j, \cdot) . As we can do this for every j , we obtain all the possible pairs using the second and the third column. As we have seen that each pair of columns sees all the pairs, we obtain an orthogonal array. \square

6.2.2 Construction

The following construction is illustrated by an example in Section 6.2.3. The reader may want to check it while reading the general construction below.

We consider orthogonal arrays with 3 columns and n^2 rows similar to the previous one but where n is a multiple of 6 and, if $n > 12$, then n is not a multiple of 4. We use an initial permutation σ but, instead of a cycle of length n , we use two disjoint cycles of length $n/2$, c_1 and c_2 , and the permutation

$$p_t = (c_{1_1}, c_{2_1})(c_{1_2}, c_{2_2}) \cdots (c_{1_i}, c_{2_i}) \cdots (c_{1_{n/2}}, c_{2_{n/2}}).$$

Given some initial permutation of the chunks, Ord , and the permutation σ , the third column of the i -th chunk in the orthogonal array, which might not be the i -chunk, is constructed as follows:

- If $i \in [1, n/2]$ then the permutation of the elements in the i -th chunk's third column is $\sigma c_2^{i-1} c_1^{i-1}$ (notice that c_1 and c_2 are disjoint and, hence, mutually independent.)
- If $i \in [n/2 + 1, n]$ then the permutation of the elements in the i -th chunk's third column is $\sigma c_2^{i-n/2-1} c_1^{i-n/2-1} p_t$.

In our case we use $c_1 = (1, 2, \dots, n/2 - 2, n/2 - 1, n/2)$ and $c_2 = (n/2 + 1, \dots, n - 1, n)$, so $p_t = (1, n/2 + 1)(2, n/2 + 2) \cdots (i, n/2 + i) \cdots (n/2, n)$.

The construction takes place in two steps. In the first one we obtain a configuration and a set S with $|S| = \Omega(n)$ and with no solution at all. Then we modify the construction to obtain $\theta(n)$ pairwise disjoint solutions. Recall that a solution is a row of the orthogonal array (x_1, x_2, x_3) , where $x_i \in S$ for all $i \in \{1, 2, 3\}$, and that two solutions, (x_1, x_2, x_3) and (y_1, y_2, y_3) , are said to be disjoint if $\{x_1, x_2, x_3\} \cap \{y_1, y_2, y_3\} = \emptyset$.

The goal is to obtain only the solutions of the form $(3k - 2, 3k - 1, 3k)$, with $k \in [1, n/6]$.

We define two bijective functions: M , the matching function, and Ord , to determine the order of the chunks.

$$\begin{aligned} M : [1, n] &\rightarrow [1, n] \\ &i \rightarrow i + n/2 \pmod n, \text{ with the usual identification } 0 \equiv n, \\ Ord : [1, n] &\rightarrow [1, n] \\ &i \rightarrow Ord(i) = \text{the position of the } i\text{-chunk.} \end{aligned}$$

The initial order of the chunks is the following:

- If $i \in [1, n/6]$ the i -th chunk is such that $Ord^{-1}(i) = 3i - 2$.
- If $i \in [n/6 + 1, 2n/6]$ the i -th chunk is such that $Ord^{-1}(i) = 3(i - n/6) - 1$.
- If $i \in [2n/6 + 1, n/2]$ the i -th chunk is such that $Ord^{-1}(i) = 3(i - 2n/6)$.
- If $i \in [n/2, n]$ the i -chunk has the position $Ord(i) = Ord(M(i)) + n/2$.

We divide each chunk into two pieces depending on the value of the second column. The part with the second component between 1 and $n/2$ is called the upper part, and if the second component is between $n/2 + 1$ and n we say that it is the lower part of the chunk.

The σ is constructed in such a way that, for $j \in [1, n/6]$, the element $3j$ appears in the position $3j - 1 + j - 1 \pmod{n/2}$ (with the identification $n/2 \equiv 0$) of the lower part (this means that it really goes to $(3j - 1 + j - 1 \pmod{n/2}) + n/2$.) Then we fill the lower part with the remaining elements in $[1, n/2]$ arbitrarily. The upper part is filled with the matching elements of the lower part: if j is the i -th element in the lower part, then $M(j)$ is the i -th element in the upper part.

The purpose of putting the element $3j$ in the $4j - 2$ spot is that, when we cycle the permutation σ with the cycles c_1 and c_2 , we would get the element $3j$ in the right spots for this stage. This is, after the required modification through the cycles (because we are moving the permutation between chunks), we will obtain the quasi-wanted solutions $(3j - 2, 3j - 1, 3j + n/2)$.

To summarize, for $j \in [1, n/6]$, $\sigma((3j - 1 + j - 1 \pmod{n/2}) + n/2) = 3j$ and for $j \in [n/2 + 1, n/2 + n/6]$, $\sigma(3j - 1 + j - 1 \pmod{n/2}) = 3j + n/2$. For the remaining i not treated before we distribute them arbitrarily with the restrictions that, if $i \in [1, n/2]$ then $\sigma(i) \in [n/2 + 1, n]$ and $\sigma(i + n/2) = \sigma(i) - n/2$. Thus, in the first chunk, the triples are $(1, i, \sigma(i))$, for $i \in [1, n]$.

Due to the condition on the n , the operation $j \rightarrow 3j - 1 + j - 1 = 4j - 2$ is bijective over $\mathbb{Z}_{n/2}$, thus making σ a valid permutation. Due to the disjointedness of c_1 and c_2 , a similar proposition as Proposition 6.3 can be shown, making this construction an orthogonal array. The difficult part is to see that the second and the third column contain all the pairs in $[1, n]^2$. If the value of the second column is $i \in [1, n/2]$ then, for the first $n/2$ chunks, we have all the pairs (i, j) , $j \in [n/2 + 1, n]$, in the last two columns. In the first chunk, all the rows with the third coordinate in the range $[n/2 + 1, n]$ are in some spot of the upper part of the chunk, and the cycle c_1 move them in solidarity through the upper part of the first $n/2$ chunks. The pairs with $i \in [1, n/2]$ and $j \in [1, n/2]$ are located in the last $n/2$ chunks as we apply the permutation p_t , which allows us to exchange the roles that the $j \in [n/2 + 1, n]$ played in the first $n/2$ chunks, for the ones that now the $j \in [n/2 + 1, n]$ plays in the last $n/2$ chunks and we use c_2 instead of c_1 . The cases in which $i \in [n/2 + 1, n]$ are treated similarly but, in the first $n/2$ chunks, we see all the pairs with $j \in [1, n/2]$ and the cycle acting is c_2 ; if $i \in [n/2 + 1, n]$ and $j \in [n/2 + 1, n]$ we should use p_t , the cycle c_1 and the last $n/2$ chunks to find the remaining pairs.

At this point we make two observations. The first one is that, if $S = [1, n/2]$ then S is solution-free, meaning that there is no triple (x_1, x_2, x_3) in the orthogonal array with $x_i \in S$, for all i . This is clear for the last $n/2$ chunks, since the first column contains an element between $n/2 + 1$ and n . For the first $n/2$ chunks we have that the upper part of them is formed by element in S but they are paired with elements not in S in the last column. Similarly, in the lower part, the elements of the second column are not in S . So, this construction is solution-free.

The second observation is that, as we have commented before, instead of the triples $(3i - 2, 3i - 1, 3i)$ we have the triples $(3i - 2, 3i - 1, M(3i))$, for $i \in [1, n/6]$. Moreover, we have also the triples $(3i - 2, M(3i - 1), 3i)$, for $i \in [1, n/6]$. Observe that we have put the element $3j$ paired with the $3j - 1 + j - 1 \pmod{n/2}$ in the lower part. Once we move ourselves through $j - 1$ chunks, the element $3j$ will be moved by c_2 $j - 1$ times, so that it will be placed in the $(3j - 1)$ -th position of the lower part. This means that we have the

triple $(3i - 2, M(3i - 1), 3i)$; therefore, in the pair $(3i - 2, 3i - 1, \cdot)$, we should have the matching of $3i$, which is $M(3i)$.

At this point we make a further modification to the orthogonal array. In the i -th chunk for $i \in [1, n/6]$, which is the $Ord(i)$ -chunk, we exchange the element $3i$ in the third column with $M(3i)$. Since the property of being an orthogonal array has to be preserved, we would also exchange $3i$ with $M(3i)$ in the $M(i)$ -th chunk. Indeed, in the $M(i)$ -th chunk we find the duplicity of the pairs $(3i - 1, 3i)$ and $(M(3i - 1), M(3i))$ corresponding to the second and third columns. Once this double exchange has been done, the second and the third column maintain the property of seeing all the pairs $(i, j) \in [1, n] \times [1, n]$. As the columns 1 and 2 have remained untouched, they maintain the orthogonal array property. Also, the columns 1 and 3 still contain all the pairs $(i, j) \in [1, n]^2$ because all the changes have been done inside the chunks. Thus, the whole thing is, still, an orthogonal array.

With this exchange, we have created the solutions $(3i - 2, 3i - 1, 3i)$, with $i \in [1, n/6]$, and no more. Thus the number of solutions is $n/6$, yet, in order to erase all the solutions, we have to delete one element in each of them. Thus, we have to delete $n/6$ elements from S , which is not a function in $o(n)$, although the number of solutions grows as $o(n^2)$.

Thus, in general, we could not have a removal lemma statement, similar of the ones showed in previous chapters, for orthogonal arrays.

6.2.3 Example

In this section we illustrate the construction with a small example, with $n = 12$.

The order of the chunks is: 1, 4, 2, 5, 3, 6, 7, 10, 8, 11, 6, 12.

Then the initial σ makes the first chunk look like:

1	1	7
1	2	<u>9</u>
1	3	8
1	4	10
1	5	11
1	6	<u>12</u>
1	7	1
1	8	<u>3</u>
1	9	2
1	10	4
1	11	5
1	12	<u>6</u>

where the underlined entries indicate that those are the first placed elements. Those will be switched with its partner in a later step as they will produce the solutions for the system. The non-underlined elements have been put arbitrarily in the lower part of the chunk, and matched with the upper part.

The most significant set of chunks are the first $n/6$ ones, as they will carry the solutions. From the second chunk till the $(n/2)$ -th, the last column is produced by cycling the initial permutation within each part, lower or upper. For the last ones, we would produce them by exchanging the role of the elements i and $i + n/2$ from the last column.

1st			7th			3rd			9th			5th			11th		
1	1	7	7	1	1	2	1	8	8	1	2	3	1	11	6	1	5
1	2	<u>9</u>	7	2	<u>3</u>	2	2	10	8	2	4	3	2	12	6	2	6
1	3	8	7	3	2	2	3	11	8	3	5	3	3	7	6	3	1
1	4	10	7	4	4	2	4	12	8	4	6	3	4	9	6	4	3
1	5	11	7	5	5	2	5	7	8	5	1	3	5	8	6	5	2
1	6	<u>12</u>	7	6	<u>6</u>	2	6	9	8	6	3	3	6	10	6	6	4
1	7	1	7	7	7	2	7	2	8	7	8	3	7	5	6	7	11
1	8	<u>3</u>	7	8	<u>9</u>	2	8	4	8	8	10	3	8	6	6	8	12
1	9	2	7	9	8	2	9	5	8	9	11	3	9	1	6	9	7
1	10	4	7	10	10	2	10	6	8	10	12	3	10	3	6	10	9
1	11	5	7	11	11	2	11	1	8	11	7	3	11	2	6	11	8
1	12	<u>6</u>	7	12	<u>12</u>	2	12	3	8	12	9	3	12	4	6	12	10
2nd			8th			4th			10th			6th			12th		
4	1	<u>9</u>	10	1	<u>3</u>	5	1	10	11	1	4	6	1	12	12	1	6
4	2	8	10	2	2	5	2	11	11	2	5	6	2	7	12	2	1
4	3	10	10	3	4	5	3	12	11	3	6	6	3	9	12	3	3
4	4	11	10	4	5	5	4	7	11	4	1	6	4	8	12	4	2
4	5	<u>12</u>	10	5	<u>6</u>	5	5	9	11	5	3	6	5	10	12	5	4
4	6	7	10	6	1	5	6	8	11	6	2	6	6	11	12	6	5
4	7	<u>3</u>	10	7	<u>9</u>	5	7	4	11	7	10	6	7	6	12	7	12
4	8	2	10	8	8	5	8	5	11	8	11	6	8	1	12	8	7
4	9	4	10	9	10	5	9	6	11	9	12	6	9	3	12	9	12
4	10	5	10	10	11	5	10	1	11	10	7	6	10	2	12	10	8
4	11	<u>6</u>	10	11	<u>12</u>	5	11	3	11	11	9	6	11	4	12	11	10
4	12	1	10	12	7	5	12	2	11	12	8	6	12	5	12	12	11

To finish we switch the underlined elements in its respective chunk to obtain the final example. The boldface type expose the solutions.

1st			7th			3rd			9th			5th			11th		
1	1	7	7	1	1	2	1	8	8	1	2	3	1	11	6	1	5
1	2	3	7	2	<u>9</u>	2	2	10	8	2	4	3	2	12	6	2	6
1	3	8	7	3	2	2	3	11	8	3	5	3	3	7	6	3	1
1	4	10	7	4	4	2	4	12	8	4	6	3	4	9	6	4	3
1	5	11	7	5	5	2	5	7	8	5	1	3	5	8	6	5	2
1	6	12	7	6	6	2	6	9	8	6	3	3	6	10	6	6	4
1	7	1	7	7	7	2	7	2	8	7	8	3	7	5	6	7	11
1	8	<u>9</u>	7	8	<u>3</u>	2	8	4	8	8	10	3	8	6	6	8	12
1	9	2	7	9	8	2	9	5	8	9	11	3	9	1	6	9	7
1	10	4	7	10	10	2	10	6	8	10	12	3	10	3	6	10	9
1	11	5	7	11	11	2	11	1	8	11	7	3	11	2	6	11	8
1	12	6	7	12	12	2	12	3	8	12	9	3	12	4	6	12	10
2nd			8th			4th			10th			6th			12th		
4	1	9	10	1	3	5	1	10	11	1	4	6	1	12	12	1	6
4	2	8	10	2	2	5	2	11	11	2	5	6	2	7	12	2	1
4	3	10	10	3	4	5	3	12	11	3	6	6	3	9	12	3	3
4	4	11	10	4	5	5	4	7	11	4	1	6	4	8	12	4	2
4	5	6	10	5	<u>12</u>	5	5	9	11	5	3	6	5	10	12	5	4
4	6	7	10	6	1	5	6	8	11	6	2	6	6	11	12	6	5
4	7	3	10	7	9	5	7	4	11	7	10	6	7	6	12	7	12
4	8	2	10	8	8	5	8	5	11	8	11	6	8	1	12	8	7
4	9	4	10	9	10	5	9	6	11	9	12	6	9	3	12	9	12
4	10	5	10	10	11	5	10	1	11	10	7	6	10	2	12	10	8
4	11	<u>6</u>	10	11	<u>6</u>	5	11	3	11	11	9	6	11	4	12	11	10
4	12	1	10	12	7	5	12	2	11	12	8	6	12	5	12	12	11

Notice that we just have exchanged the elements of the third column in the $[1, n/6]$ first chunks and have done the same, but reversed, in its respective pairs of chunks (the $[n/2 + 1, n/2 + n/6]$ ones.) The remaining chunks do not suffer any change.

6.3 On the dimension of the representation

In this section, we aim to discuss a technical question about which hypergraph dimension we need to represent each system of equations.

Following the works in Chapter 2, we can observe that some systems can be represented by graphs. Using our representation techniques (where each cycle in the graph represents a valid equation of the graph,) Theorem 2.4 is best possible in the following sense: every finite graph where each edge is painted with a different colour represents a unique system. This comment is further developed in Section 6.3.1.

As we have seen in the proof of Theorem 3.1, a similar representation idea applies for

systems over finite fields. A clear example occurs when every $k \times k$ submatrix of a $k \times m$ integer matrix A is invertible: the hypergraph representation turns to be a Hamiltonian k -cycle¹ hypergraph on m vertices. Indeed, we can mix both ideas to reduce the dimension of the hypergraph that represents the system of equations by increasing the number of cycles in the hypergraph. We show this reduced dimension with an example in Section 6.3.2.

It is not clear how a general statement regarding the dimension of the representation might look like. However, the role that the variables play in different equations seems to be the determining factor responsible for increasing the dimension (at least using the techniques showed in this work.) In the case of systems representable by graphs, each variable plays a similar role in each of the equations in which appears. In the example of Section 6.3.2, we mimic this behavior with a group of variables in order to reduce the dimension of the hypergraph, but we still need a hypergraph to represent such a system.

6.3.1 The case of graphs (for systems in Abelian groups)

From a graph to a system In a given graph² \mathcal{G} with edges labelled with $x_1, \dots, x_{|E|}$ and where each edge has a given orientation, each circuit c on the graph represents a natural equation with the variables $x_1, \dots, x_{|E|}$. This representation is the one presented in Chapter 2 and follows the representation of a single equation by a cycle.

Notice that the leaves³ of graph are useless with respect to the system of equations. Thus, we might consider just graphs without leaves for the graphs that represent systems of equations.

Notice that each circuit c comes with an orientation. Let x_{i_1}, \dots, x_{i_t} be the edges of the circuit starting at v and following the circuit orientation. The equation that c outputs in the Abelian group is:

$$\epsilon_{i_1} x_{i_1} + \dots + \epsilon_{i_t} x_{i_t} = 0$$

where $\epsilon_{i_j} = 1$ if the edge labelled by x_{i_j} has the same orientation in the graph and in the circuit, or $\epsilon_{i_j} = -1$ if the orientation of the edge x_{i_j} is opposite to the circuit orientation.

For a given graph, the orientation of the edges is fixed and consistent throughout all the circuits and the multiplication by ± 1 is an automorphism in any given Abelian group. Following the lines of Chapter 2, the set of all the circuits form a system of equations. In this way, any graph outputs a unique system of linear equations over an Abelian group. The uniqueness is up to automorphisms like changing the name of the variables.

If we just use the circuits that form a base of fundamental cycles FC instead of all the circuits, we can create all the possible circuits and equations in a coherent way by generating integer linear combinations of the elements in FC . Indeed, the system generated by the fundamental cycles and the one that contains all the circuits are equivalent in any given Abelian group (see [21]).

Therefore, we can obtain the system from a given graph \mathcal{G} by picking a base of fundamental cycles. This representation is unique in the sense that, for a given graph, there exists just one linear system represented by the graph. The uniqueness is up to change of base by integer unimodular matrices. This graph representation depends on how we have translated the cycles into equations.

¹see [45]

²may be a multigraph

³Edges that do not belong to any proper cycle.

From a system to a graph So far we have shown how a graph represents a system. In the following paragraphs we show how we can construct a graph from a graph-representable system.

For this section, we say that an $k \times m$ integer linear system of equations S with variables x_1, \dots, x_m is graph-representable by \mathcal{G} if: there is an orientation and a labeling of the edges with x_1, \dots, x_m and an spanning tree T for which $S = U \cdot Inc$ where Inc is the vertex-edges incidence matrix of the fundamental cycles given by T and U is a unimodular matrix⁴.

Recall that a base of fundamental cycles with respect to an spanning tree T , contains the unique cycles formed by adding an edge e to an spanning tree T , with $e \in E \setminus T$. Since each fundamental cycle base can be obtained from another via a unimodular change of base (see [11]), all of them are equivalent in any Abelian group.

If a system is graph-representable by \mathcal{G} , then the deletion of any edge e in \mathcal{G} , together with any leaf in $\mathcal{G} \setminus e$, creates a new graph and a new system associated with it. To be precise, it has one less equation and, if we have not created any leafs by erasing e from \mathcal{G} , also one less variable.

We can go from the system induced by \mathcal{G} to the one induced by $\mathcal{G} \setminus e$ and back by choosing an spanning tree $T \subset \mathcal{G}$ which do not contain e as an edge. So e is an edge that closes one fundamental cycle in \mathcal{G} . The set of fundamental cycles/equations in $\mathcal{G} \setminus e$ is obtained by deleting the fundamental cycle/equation that e creates in T . The other way around also holds: the equations in \mathcal{G} are the ones from $\mathcal{G} \setminus e$ together with the new equation created by adding e to $\mathcal{G} \setminus e$. Instead of a single edge we might have to add a path, depending on whether the deletion of e creates leafs or not in $\mathcal{G} \setminus e$.

With this two considerations in mind, we can build the graph from the bottom-up: starting by just one equation and add new equations and variables in a consistent way. If, at some point, we are not able to add another equation to the system, then the system is not graph representable.

For a system to be graph representable, there has to be an equivalent system with a totally unimodular system matrix A . Moreover, we can assume that A contains a $k \times k$ identity matrix.

Assume that A is graph representable using \mathcal{G} and that the identity matrix contains 1's for the variables x_1, \dots, x_k . These variables are represented by the edges e_1, \dots, e_k with the property that $\mathcal{G} \setminus \{e_1, \dots, e_k\}$ is a tree.

Indeed, the rows A_i represent some circuits. Since we are considering that cycles represent equations and the A_i generates the equation space, then the circuits coming from the rows A_i generate the cycles space of \mathcal{G} . Moreover, as the matrix is totally unimodular, we can represent the cycle/equation space by integer linear combinations of the A_i . Let us assume A_i is not a proper cycle containing e_i but a circuit. Then the circuit induced by A_i contains more, shorter cycles c_1, \dots, c_{a_i} . Since A has the $k \times k$ identity matrix as a submatrix, there is a distinguished cycle in A_i , namely the one that contains the edge that represents x_i , let it be c_1 . As none of the other cycles c_j , $j \in [2, a_i]$ contains the edge representing x_s , $s \in [1, k]$, they cannot be generated by an integer linear combination of the A_i . This means that the cycle space is larger than the assumption, reaching a

⁴This definition can be extended to other pairs of (system, group) depending on the order of the group and the determinantal of the matrix S . The definition here posted works for any group, regardless of its order. See Chapter 2 for more details.

contradiction. Therefore A_i has to appear in the graph as a genuine cycle.

Thus, we can start to build our graph starting with a the first equation in A : it should be a (directed) cycle, but the order of the variables and the edges might not be the same as the order in the system A . We continue to add equations with some variables in common with the previous set of equations. Since any subset of equations from A , S , should be also graph-representable, if we cannot build a graph with one more equation, $S \cup \{A_i\}$, then the final system is not graph representable for all the groups. Notice that we might have to change some initial assignation of the variables to the edges in the previously built graph in order to accommodate the new cycle.

In this way we can determine, for a given graph, which system does it represents. Moreover, given a system of equations, we can find a graph that represents it or determine that the system is not graph-representable.

This discussion arose by a question of Candela [14].

6.3.2 Example of a system with lower dimension

In this subsection we present an example of a system that, even though it has four equations, it can be represented using a 3-uniform hypergraph.

The system is:

$$\left. \begin{array}{rcl} x_1 - 2x_2 + x_3 & & = 0 \\ & x_2 - 2x_3 + x_4 & = 0 \\ x_1 & - x_5 + x_6 + 3x_7 & = 0 \\ & x_5 - 2x_6 + x_7 & = 0 \end{array} \right\}, \quad (6.4)$$

and its matrix is A .

We can represent this system with a 3-uniform hypergraph on 5 vertices, $\{v_1, \dots, v_5\}$ and 7 edges:

Edge representing x_1 : $\{v_1, v_2, v_3\}$. Color 1.
 Edge representing x_2 : $\{v_2, v_3, v_4\}$. Color 2.
 Edge representing x_3 : $\{v_3, v_4, v_1\}$. Color 3.
 Edge representing x_4 : $\{v_4, v_1, v_2\}$. Color 4.
 Edge representing x_5 : $\{v_2, v_3, v_5\}$. Color 5.
 Edge representing x_6 : $\{v_3, v_5, v_1\}$. Color 6.
 Edge representing x_7 : $\{v_5, v_1, v_2\}$. Color 7.

Let H be this coloured hypergraph.

Recall that, to translate the values in the sets X_i to the hypergraph, we use a matrix C . In this case:

$$C = \begin{pmatrix} -1 & -2 & -3 & 0 & 0 \\ 0 & -1 & -2 & -1/3 & 0 \\ 1 & 0 & -1 & -2/3 & 0 \\ 2 & 1 & 0 & -1 & 0 \\ 0 & -1/2 & -6 & 0 & -7 \\ 1/7 & 0 & -3 & 0 & -4 \\ 2/7 & 1/2 & 0 & 0 & -1 \end{pmatrix}.$$

To build this matrix we have gone through the procedure described in Chapter 3 for the first two equations first and, afterwards, for the last two. Since they have to share the edge corresponding to x_1 , we have had to adapt the coefficients of the matrix accordingly. Originally, the coefficients $C_{5,2}$ and $C_{7,2}$ were -1 and 1 respectively; also $C_{5,3}$ and $C_{6,3}$ were previously -2 and -1 . Notice that we might need to be in a field, or a group, where the division by 2, 3 and 7 is allowed.

Once the matrix C has been built, the construction of the graph is exactly the same as explained in Chapter 3 for the case of finite fields.

To show the relation between solutions and copies of the hypergraph we should notice that, for each single value g_i in X_i , there will be $|G|^2$ edges labelled g_i . If we see a hypergraph with all the edges colored in different colors, then it is supported with one vertex in each class corresponding with v_i . Assume the label of the vertices are $\nu = (\nu_1, \dots, \nu_5)$. Then its edges bear a solution as $AC = 0$ and the edges with label $g := (g_1, \dots, g_7)$ fulfill $C\nu^T = g^T$.

To see that each solution (g_1, \dots, g_7) generates $|G|^2$ edge-disjoint copies of H , we pick an index i and an edge with the label g_i . Notice that we have already $|G|^2$ choices for this. This edge, along with the full solution, determines a whole copy as we can split the system according to the first two equations or the second pair. Assume $i \in [1, 4]$, the other case is similar. As (g_1, g_2, g_3, g_4) fulfill the first pair of equations, and the hypergraph that uses the vertices v_1, v_2, v_3, v_4 can be seen as the representation that would correspond with the system induced by the first pair of equations, there should be a unique vertex, outside the i -th edge, that completes a quadruple v_1, v_2, v_3, v_4 and that supports both the edge and bears the solution (g_1, g_2, g_3, g_4) . Once we have x_1 , we complete the graph with an appropriate vertex v_5 using the same argument, but with the second pair of equations. This completes the construction of the example.

6.3.3 Comments on the complexity of the system

The graph/hypergraph dimension seems to be related to the way the variables are tied together in the equations of the system and not to the structure of the solution set. However, it is not clear which is the minimal dimension of a hypergraph representing a linear system and how to determine it. Also, the way of representing a linear system using a hypergraph is not unique; see [68] for a different representation.

The use of hypergraphs allows us to capture the entangled relations between the variables that a general linear system of equations can impose. However, it is not clear which dimension is needed. In our construction, we need uniform hypergraphs whose edges have dimension $k + 1$, where k is the number of equations. However, this can be reduced to dimension 2, or graphs, for some systems where the behavior of the variables is essentially the same in all the equations (Theorem 2.4 and Theorem 2.5). Also, when the variables do not behave much different from one equation to the other, it can be reduced below $k + 1$. This dimension is, therefore, very dependent on the technique used to represent a system using a hypergraph. It is not at all clear whether the proof of these algebraic removal lemmas need the hypergraph removal lemma (or the graph removal lemma at all.)

Moreover, when we represent the system in an Abelian group using the techniques from Chapter 5, we may use higher dimension than $k + 1$ to avoid problems with the divisibility by the denominators of the rational numbers that appear in the construction of the matrix C . However, this extra dimensions might not be necessary as the construction by Szegedy in [73] shows for the $k + 2$ -term arithmetic progression system.

In several works, like Green and Tao [42] or Gowers and Wolf [38], the authors introduce notions of complexity of the system related with the solution set of the system. This notion of complexity reflects the difficulty of problems. For instance, Goldbach’s conjecture or the twin prime conjecture has “infinite” complexity (the corresponding system has only two variables). Whereas arithmetic progressions are defined by low complexity systems. It was hinted to us that there might be a connection between this notion of complexity and the dimension of hypergraphs, but some examples indicate that the two notions are of different nature (see also the work by Candela [13].)

6.4 Symmetry-preserving lemma

In [73], Szegedy presented a Symmetry-preserving removal lemma. The proof of this result is similar to the ones found in Theorem 3.1 or Theorem 4.1: a pigeonhole principle is combined with the automorphism group G in the hypergraph is used to ensure that the edges removed form a union of orbits of G .

Along with the general framework of the Symmetry-preserving removal lemma, a particular construction involving a generalization of Cayley graphs for hypergraphs is presented and the automorphism group is specified. In order to use the result, certain conditions regarding certain applications and group quotients should be checked. This framework can be connected with the work presented in this thesis by observing that the construction of the matrices C , as well as the edges of the hypergraph H in Theorem 3.1 or Theorem 4.1, fulfill the conditions required to apply the Symmetry-preserving removal lemma either for linear systems in finite fields or in finite Abelian groups. Thus the present work, although independent from the work of Szegedy, can be seen as a formal completion of his framework.

6.5 Open problems and future work

Removal lemma for non-abelian groups From the perspective of Chapter 2, it seems a natural question to ask about a removal lemma for “linear” systems on non-abelian finite groups as presented in Chapter 2. This question seems hard in a general scenario and the results in Chapter 2 regarding removal lemmas in non-necessarily Abelian groups, Theorem 2.1 or Theorem 2.5, deal, apparently, with the systems and equations that behave closely to the Abelian case.

Deletion of the condition on the k -th determinantal A natural question that has remained open from Theorem 4.1 of Chapter 4 is whether the condition concerning the determinantal of the system is necessary. Under this condition, the behaviour of the linear system with respect to the group resembles the case in which the group is a finite field. As we have seen, if $d_k(A) = d > 1$, then there is an equivalent system in which all the coefficients of a row are multiples of some divisor of d . This observation indicates that, even though the condition might seem artificial, it is somewhat natural.

Notice that, if $\gcd(d_k(A), |G|) > 1$, then the solution set $S(A, G)$ of a full rank $k \times m$ integer linear system A in G^m does not have size $|G|^{m-k}$. Indeed, $|S(A, G)|$ is a larger function depending A and G as shown in Section 4.5. Therefore, two notions of “small number of solutions” appear: the easier case, where we ask for at most $\delta|G|^{m-k}$ solutions,

or the more natural bound $\delta|S(A, G)|$. Even though the answer seems to be affirmative in both cases, additional ideas are needed.

The condition over the determinantal can be relaxed in some cases. If the group is close to a finite field, like $\mathbb{Z}_3^n \times \mathbb{Z}_2$, then some equations can be omitted as they are almost superfluous. Even more, when the proportion of the subgroup $S(A, G) \subset G^m$ with respect to G^{m-k} is a constant, then the determinantal condition can be bypassed. Notice that these two examples represent the extremal cases: either the subgroup $|S(A, G)|/|G^{m-k}|$ is a constant fraction of G^i for some i , or $|S(A, G)|/|G^{m-k}|$ has constant size.

For instance, let $G = \mathbb{Z}_n$. By Proposition 4.13 the system $Ax = b$, $x \in \mathbb{Z}_n^m$ is equivalent to $\overline{A}x = \overline{b}$ where, for some rows r_1, \dots, r_s , we have $\prod_{i=1}^s \gcd(\overline{A}_{r_i,1}, \dots, \overline{A}_{r_i,m}) = d$. Therefore, the solution set of $\overline{A}x = \overline{b}$ is, either the empty set, or the union of the solutions sets of the systems $A'x = b_i$ for some $\{b_i\}_{i \in [1, |G|/|d \cdot G|]}$ where $d \cdot G$ is the group $\{dg : g \in G\}$ and $d_k(A') = 1$. Notice that $|G|/|d \cdot G| < d$. If the number of solutions to $\overline{A}x = \overline{b}$, with $x \in \prod_{i=1}^m X_i$, is less than $\delta d|G|^m$ then, for each of the b_j , the number of solutions to $A'x = b_j$ is at most $\delta d|G|^m$; using Theorem 4.1, we obtain sets $X_{i,j}$ with $|X_{i,j}| < \epsilon/d|G|$ such that $A'x = b_j$ with $x \in \prod_{i=1}^m X_i \setminus X_{i,j}$ is solution-free. Therefore, there is no solution to $Ax = b$ with $x \in \prod_{i=1}^m X_i \setminus (\cup_{j=1}^{|G|/|d \cdot G|} X_{i,j})$. Since $|\cup_{j=1}^{|G|/|d \cdot G|} X_{i,j}| < \epsilon|G|$, a removal lemma is shown.

Although it seems that the determinantal condition might be omitted from Theorem 4.1 and the hypotheses of the theorem changed accordingly, the construction presented in Chapter 4 does not seem to work in the general case. A deeper understanding of the structure of the solution set might be needed in order to show a similar result.

Number of monochromatic solutions for all finite groups In Chapter 5 we have shown that, under some conditions on the system, the number of monochromatic solutions in certain groups is asymptotically as large as one can expect. The cases treated are those where all the elements in the ambient group have bounded order. Moreover, using [28, Theorem 1] and an standard argument we can show:

Theorem 6.4. *Let A be a $k \times m$ full rank integer matrix with $m > k$. Let $r > 0$ be an integer and $\epsilon > 0$ be a positive number. Let G be an Abelian group of order n . Assume A fulfills the columns property in \mathbb{Q} and there exists a $g \in G$ and $|g|/|G| > \epsilon$. Then, there exist a $\delta = \delta(\epsilon, r, A) > 0$ such that for any r -coloring, the number of monochromatic solutions to $Ax = 0$ is, at least,*

$$\lfloor \delta G^{m-k} \rfloor.$$

Therefore, we have sufficient conditions to ensure a large number of monochromatic solutions in the two opposite cases of finite Abelian groups: large products of small cyclic groups and essentially a unique cyclic group. Even though the proofs of Theorem 5.26 and [28, Theorem 1] rely heavily on the structure of the group, it seems plausible to ask if a similar result might hold for any Abelian group, but new ideas seem to be needed.

Characterization for solutions of maximal order Based on the comments in Section 5.6, one may wonder about a characterization of the systems for which any r -coloring contains solutions where all its elements have maximal order, or, at least, relax the conditions of Lemma 5.8.

This is, Lemma 5.8 presents the n -columns condition as sufficient to ensure, for any r -coloring, monochromatic solutions where all the elements have maximal order. However, it is not clear which condition might be necessary and sufficient to obtain the same conclusion.

Monochromatic solutions on non-abelian groups In the case of monochromatic solutions of linear systems in non-abelian groups very little is known. For example [43, Theorem 3.3.1] shows that: for every r , there exists an $N = N(r)$ such that if a group G has more than N elements, then every r -coloring contains a monochromatic solution to the equation $x + y = z$.

However, a general sufficient condition (of those systems for which any finite group, with large enough order, have a monochromatic solution for every r -coloring) is missing. A more concrete question could be: find characterizations of systems for certain families of groups, like Deuber's result [20] did for finite Abelian groups.

On bounds of the removal lemma The relation between the number of copies of a (hyper)graph and the number of (hyper)edges to be removed makes these results not valid for practical purposes. The original proof of the removal lemma for graphs involved the use of the Szemerédi Regularity Lemma. This makes the above relation be tower-like: let K be a graph on n vertices, if the number of edges to be deleted is ϵn^2 , then the upper bound on the number of graphs H on h vertices that we can erase is

$$\frac{1}{2^{\underbrace{2^{2^2 \dots}}_{\approx 1/\epsilon^5}}} n^h.$$

This bound comes from the number of pairs needed to obtain an ϵ -regular partition of a graph. Since Gowers [35] showed that for some graphs

$$\frac{1}{2^{\underbrace{2^{2^2 \dots}}_{\approx \log(1/\epsilon)}}}$$

partitions are needed, the tower-type bound is unavoidable.

In [27], Fox showed another proof of the removal lemma for graphs without using Szemerédi's regularity lemma and could lower the upper bound from a tower-type of height $1/\epsilon^5$ to

$$\frac{1}{2^{\underbrace{2^{2^2 \dots}}_{\approx h^4 \log(1/\epsilon)}}}.$$

However, an upper bound for the removal lemma which is not of tower-type is not known. Moreover, the bounds for the hypergraph removal lemma are even worse than for graphs and no improvement is known.

Removal lemma for compact Abelian groups A recent result of Candela and Sissak [15] shows that a removal lemma for linear systems of equations holds for the circle. The proof can be extended to the groups \mathbb{T}^n , but the relation between the measure of the set of solutions and the measure of the set of elements to be removed depends on the topological dimension of the compact Abelian group, n in \mathbb{T}^n . As it is suggested in [15], it seems interesting to try to remove such dependence. Szegedy and the author are working on that direction.

Bibliography

- [1] M. Ajtai and E. Szemerédi. Sets of lattice points that form no squares. *Stud. Sci. Math. Hungar.*, 9:9–11 (1975), 1974.
- [2] N. Alon, R. A. Duke, H. Lefmann, V. Rödl, and R. Yuster. The algorithmic aspects of the regularity lemma. *J. Algorithms*, 16(1):80–109, 1994.
- [3] N. Alon, E. Fischer, M. Krivelevich, and M. Szegedy. Efficient testing of large graphs. In *40th Annual Symposium on Foundations of Computer Science (New York, 1999)*, pages 656–666. IEEE Computer Soc., Los Alamitos, CA, 1999.
- [4] N. Alon, E. Fischer, M. Krivelevich, and M. Szegedy. Efficient testing of large graphs. *Combinatorica*, 20(4):451–476, 2000.
- [5] N. Alon and A. Shapira. Testing subgraphs in directed graphs. *J. Comput. System Sci.*, 69(3):353–382, 2004.
- [6] N. Alon and A. Shapira. Every monotone graph property is testable. *SIAM J. Comput.*, 38(2):505–522, 2008.
- [7] T. Austin and T. Tao. Testability and repair of hereditary hypergraph properties. *Random Structures Algorithms*, 36(4):373–463, 2010.
- [8] V. Bergelson, W. A. Deuber, and N. Hindman. Rado’s theorem for finite fields. In *Sets, graphs and numbers (Budapest, 1991)*, volume 60 of *Colloq. Math. Soc. János Bolyai*, pages 77–88. North-Holland, Amsterdam, 1992.
- [9] N. Biggs. *Algebraic graph theory*. Cambridge University Press, London, 1974. Cambridge Tracts in Mathematics, No. 67.
- [10] J. Bourgain. Roth’s theorem on progressions revisited. *Journal d’Analyse Mathématique*, 104(1):155–192.
- [11] R. A. Brualdi and H. J. Ryser. *Combinatorial matrix theory*, volume 39 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1991.
- [12] B. Bukh. Personal communication. 2011.
- [13] P. Candela. *Developments at the interface between combinatorics and Fourier analysis*. PhD thesis, University of Cambridge, 2009.

- [14] P. Candela. Personal communication. 2011.
- [15] P. Candela and O. Sisask. A removal lemma for linear configurations in subsets of the circle. *Proceedings of the Edinburgh Mathematical Society*, (accepted).
- [16] F. R. K. Chung. Regularity lemmas for hypergraphs and quasi-randomness. *Random Structures Algorithms*, 2(2):241–252, 1991.
- [17] F. R. K. Chung, R. L. Graham, and R. M. Wilson. Quasi-random graphs. *Combinatorica*, 9(4):345–362, 1989.
- [18] J. Cilleruelo. Combinatorial problems in finite fields and sidon sets. *Combinatorica*, (to appear).
- [19] D. Conlon and J. Fox. Bounds for graph regularity and removal lemmas. *arXiv*, 07 2011.
- [20] W. Deuber. Partition theorems for abelian groups. *J. Combinatorial Theory Ser. A*, 19:95–108, 1975.
- [21] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, third edition, 2005.
- [22] G. Elek and B. Szegedy. Limits of hypergraphs, removal and regularity lemmas. a non-standard approach. *arXiv*, 05 2007.
- [23] G. Elek and B. Szegedy. A measure-theoretic approach to the theory of dense hypergraphs. *arXiv*, 10 2008.
- [24] P. Erdős, P. Frankl, and V. Rödl. The asymptotic number of graphs not containing a fixed subgraph and a problem for hypergraphs having no exponent. *Graphs Combin.*, 2(2):113–121, 1986.
- [25] P. Erdős and A. H. Stone. On the structure of linear graphs. *Bull. Amer. Math. Soc.*, 52:1087–1091, 1946.
- [26] M. Fang. On the completion of a partial integral matrix to a unimodular matrix. *Linear Algebra Appl.*, 422(1):291–294, 2007.
- [27] J. Fox. A new proof of the graph removal lemma. *Ann. of Math. (2)*, 174(1):561–579, 2011.
- [28] P. Frankl, R. L. Graham, and V. Rödl. Quantitative theorems for regular systems of equations. *J. Combin. Theory Ser. A*, 47(2):246–261, 1988.
- [29] P. Frankl and V. Rödl. The uniformity lemma for hypergraphs. *Graphs Combin.*, 8(4):309–312, 1992.
- [30] P. Frankl and V. Rödl. Extremal problems on set systems. *Random Structures Algorithms*, 20(2):131–164, 2002.
- [31] A. Frieze and R. Kannan. Quick approximation to matrices and applications. *Combinatorica*, 19(2):175–220, 1999.

- [32] Z. Füredi. Extremal hypergraphs and combinatorial geometry. In *Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994)*, pages 1343–1352, Basel, 1995. Birkhäuser.
- [33] H. Furstenberg and Y. Katznelson. An ergodic Szemerédi theorem for commuting transformations. *J. Analyse Math.*, 34:275–291 (1979), 1978.
- [34] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *J. ACM*, 45(4):653–750, 1998.
- [35] W. T. Gowers. Lower bounds of tower type for Szemerédi’s uniformity lemma. *Geom. Funct. Anal.*, 7(2):322–337, 1997.
- [36] W. T. Gowers. Quasirandomness, counting and regularity for 3-uniform hypergraphs. *Combin. Probab. Comput.*, 15(1-2):143–184, 2006.
- [37] W. T. Gowers. Hypergraph regularity and the multidimensional Szemerédi theorem. *Ann. of Math. (2)*, 166(3):897–946, 2007.
- [38] W. T. Gowers and J. Wolf. The true complexity of a system of linear equations. *Proc. Lond. Math. Soc. (3)*, 100(1):155–176, 2010.
- [39] R. L. Graham, K. Leeb, and B. L. Rothschild. Ramsey’s theorem for a class of categories. *Advances in Math.*, 8:417–433, 1972.
- [40] R. L. Graham, B. L. Rothschild, and J. H. Spencer. *Ramsey theory*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons Inc., New York, second edition, 1990. A Wiley-Interscience Publication.
- [41] B. Green. A Szemerédi-type regularity lemma in abelian groups, with applications. *Geom. Funct. Anal.*, 15(2):340–376, 2005.
- [42] B. Green and T. Tao. Linear equations in primes. *Ann. of Math. (2)*, 171(3):1753–1850, 2010.
- [43] M. Hablicsek. Sum-free sets. Master’s thesis, Eötvös Loránd University, 2009.
- [44] A. W. Hales and R. I. Jewett. Regularity and positional games. *Trans. Amer. Math. Soc.*, 106:222–229, 1963.
- [45] H. Hàn and M. Schacht. Dirac-type results for loose Hamilton cycles in uniform hypergraphs. *J. Combin. Theory Ser. B*, 100(3):332–346, 2010.
- [46] Y. Ishigami. A simple regularization of hypergraphs. *arXiv*, 2009.
- [47] J. Komlós, A. Shokoufandeh, M. Simonovits, and E. Szemerédi. The regularity lemma and its applications in graph theory. In *Theoretical aspects of computer science (Tehran, 2000)*, volume 2292 of *Lecture Notes in Comput. Sci.*, pages 84–112. Springer, Berlin, 2002.
- [48] J. Komlós and M. Simonovits. Szemerédi’s regularity lemma and its applications in graph theory. In *Combinatorics, Paul Erdős is eighty, Vol. 2 (Keszthely, 1993)*, volume 2 of *Bolyai Soc. Math. Stud.*, pages 295–352. János Bolyai Math. Soc., Budapest, 1996.

- [49] D. Král, O. Serra, and L. Vena. A removal lemma for linear systems over finite fields. In *Sixth Conference on Discrete Mathematics and Computer Science (Spanish)*, pages 417–423. Univ. Lleida, Lleida, 2008.
- [50] D. Král, O. Serra, and L. Vena. A combinatorial proof of the removal lemma for groups. *J. Combin. Theory Ser. A*, 116(4):971–978, 2009.
- [51] D. Král, O. Serra, and L. Vena. On the removal lemma for linear systems over abelian groups. *arXiv*, 06 2011.
- [52] D. Král, O. Serra, and L. Vena. A removal lemma for systems of linear equations over finite fields. *Israel J. Math.*, 187(1):193–207, 2012.
- [53] C. Liebchen. Finding short integral cycle bases for cyclic timetabling, 2003.
- [54] C. Liebchen and L. Peeters. On cyclic timetabling and cycles in graphs. Technical report, Technische Universität Berlin, 2002.
- [55] L. Lovász and B. Szegedy. Szemerédi’s lemma for the analyst. *Geom. Funct. Anal.*, 17(1):252–270, 2007.
- [56] B. Nagle, V. Rödl, and M. Schacht. The counting lemma for regular k -uniform hypergraphs. *Random Structures Algorithms*, 28(2):113–179, 2006.
- [57] M. Newman. *Integral matrices*. Academic Press, New York, 1972. Pure and Applied Mathematics, Vol. 45.
- [58] D. Piguet and M. J. Stein. An approximate version of the loebl-komlos-sos conjecture. *arXiv*, 08 2007.
- [59] R. Rado. Studien zur Kombinatorik. *Math. Z.*, 36(1):424–470, 1933.
- [60] V. Rödl and M. Schacht. Regular partitions of hypergraphs: counting lemmas. *Combin. Probab. Comput.*, 16(6):887–901, 2007.
- [61] V. Rödl and M. Schacht. Regular partitions of hypergraphs: regularity lemmas. *Combin. Probab. Comput.*, 16(6):833–885, 2007.
- [62] V. Rödl and J. Skokan. Regularity lemma for k -uniform hypergraphs. *Random Structures Algorithms*, 25(1):1–42, 2004.
- [63] V. Rödl and J. Skokan. Applications of the regularity lemma for uniform hypergraphs. *Random Structures Algorithms*, 28(2):180–194, 2006.
- [64] K. F. Roth. On certain sets of integers. *J. London Math. Soc.*, 28:104–109, 1953.
- [65] I. Z. Ruzsa and E. Szemerédi. Triple systems with no six points carrying three triangles. In *Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976)*, Vol. II, volume 18 of *Colloq. Math. Soc. János Bolyai*, pages 939–945. North-Holland, Amsterdam, 1978.
- [66] O. Serra and L. Vena. On the number of monochromatic solutions of integer linear systems on abelian groups. *arXiv*, 03 2012.
- [67] A. Shapira. Green’s conjecture and testing linear-invariant properties. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 159–166, 2009.

- [68] A. Shapira. A proof of Green's conjecture regarding the removal properties of sets of linear equations. *J. Lond. Math. Soc. (2)*, 81(2):355–373, 2010.
- [69] J. Solymosi. Note on a generalization of Roth's theorem. In *Discrete and computational geometry*, volume 25 of *Algorithms Combin.*, pages 825–827. Springer, Berlin, 2003.
- [70] J. Solymosi. A note on a question of Erdős and Graham. *Combin. Probab. Comput.*, 13(2):263–267, 2004.
- [71] J. Solymosi. Roth type theorems in finite groups. *arXiv*, 01 2012.
- [72] J. H. Spencer. Ramsey's theorem for spaces. *Trans. Amer. Math. Soc.*, 249(2):363–371, 1979.
- [73] B. Szegedy. The symmetry preserving removal lemma. *Proc. Amer. Math. Soc.*, 138(2):405–408, 2010.
- [74] E. Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arith.*, 27:199–245, 1975. Collection of articles in memory of Juriĭ Vladimirovič Linnik.
- [75] E. Szemerédi. Regular partitions of graphs. In *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*, volume 260 of *Colloq. Internat. CNRS*, pages 399–401. CNRS, Paris, 1978.
- [76] T. Tao. Szemerédi's regularity lemma revisited. *Contrib. Discrete Math.*, 1(1):8–28 (electronic), 2006.
- [77] T. Tao. A variant of the hypergraph removal lemma. *J. Combin. Theory Ser. A*, 113(7):1257–1280, 2006.
- [78] P. Varnavides. On certain sets of positive density. *J. London Math. Soc.*, 34:358–360, 1959.
- [79] B. Voigt. The partition problem for finite abelian groups. *J. Combin. Theory Ser. A*, 28(3):257–271, 1980.
- [80] Y. Yeh. On prime power Abelian groups. *Bull. Amer. Math. Soc.*, 54:323–327, 1948.
- [81] X. Zhan. Completion of a partial integral matrix to a unimodular matrix. *Linear Algebra Appl.*, 414(1):373–377, 2006.
- [82] Y. Zhao. Proof of the $(n/2 - n/2 - n/2)$ conjecture for large n . *Electron. J. Combin.*, 18(1):Paper 27, 61, 2011.

Index

- t -partite graph, 9
- adjugate of a matrix, 34
- characteristic vector
 - cycle, 13
 - equation, 12
- circuit, 9
- circular unimodular matrices, 30
- coloring, 9
- columns condition
 - over a field, 45
 - over \mathbb{Q} , 43
 - over \mathbb{Z}_n , 47
- copy of a subgraph, 9
- cycle, 9
 - space, 13
- determinantal, 9
- determinantal divisor, *see* determinantal
- directed graph, 9
- edge-disjoint copies, 9
- fundamental cycle, 13
- fundamental theorem finite Abelian groups, 10
- graph, 9
 - directed, 9
 - undirected, 9
- graph homomorphism, 9
- graph representable
 - strongly, 16
- graph representable system, 12
- Hales Jewett Theorem, 57
 - extended, 59
- leaf of a graph, 74
- matrix extension, 35
- monochromatic solutions
 - over \mathbb{Z} , 44
 - over \mathbb{Z}_n^N , 47
 - over bounded torsion groups, 62
 - over finite fields, 45
- number subgroups of a p -group, 48
- quasi-random graph, 1
- regular partition, 44
- Regularity Lemma
 - colored graphs, 2
 - directed graphs, 2
 - hypergraphs, 2
- Removal Lemma
 - Abelian groups, 5
 - colored directed graphs, 14
 - colored graphs, 3
 - colored hypergraphs, 4, 21
 - directed graphs, 3, 11
 - directed hypergraphs, 4
 - finite Abelian group graph representable, 12
 - finite group graph representable, 13
 - finite groups, 10
 - graphs, 3
 - hypergraphs, 3
 - linear systems finite Abelian groups, 29
 - linear systems finite fields, 19
- restricted system, 34
 - extension, 34
- Roth theorem for finite groups, 10
- skeleton, 56
- Smith Normal Form, 35
- subgraph, 9
- Szemerédi Theorem, 5
 - for finite fields, 20
- uniform hypergraph, 9