



UNIVERSIDAD DE BUENOS AIRES
Facultad de Ciencias Exactas y Naturales
Departamento de Matemática

Teoremas inversos discretos

Tesis presentada para optar al título de Doctor de la Universidad de Buenos Aires
en el área Ciencias Matemáticas

Miguel Nicolás Walsh

Director de tesis y consejero de estudios: Román Sasyk

Buenos Aires, 2012

Teoremas inversos discretos

Resumen

La presente tesis estudia dos instancias diferentes de teoremas inversos discretos, la primera relacionada con cuestiones de convergencia en la teoría ergódica y la segunda con problemas de distribución local en la teoría de números.

El primer resultado nos permite caracterizar aquellas funciones que pueden formar promedios ergódicos no convencionales grandes en la norma L^2 . Discutimos luego descomposiciones abstractas de estructura y aleatoriedad y las extendemos al contexto de la teoría ergódica, habilitando la posibilidad de estudiar varios niveles de estructura en forma simultánea. Combinando estas herramientas con el teorema inverso previamente mencionado y un proceso inductivo adecuado, logramos demostrar que los promedios ergódicos polinomiales múltiples provenientes de la acción de un grupo nilpotente de transformaciones que preservan la medida en un espacio de probabilidad siempre convergen en norma. Esto responde una conjetura de Bergelson y Leibman.

El segundo resultado concierne la distribución de conjuntos en clases residuales. Introducimos los conceptos de conjuntos característicos y genéricos, y los aplicamos en el marco de la criba de Gallagher para mostrar que si un conjunto grande de puntos enteros $S \subseteq \{1, \dots, N\}^d$, $d > 1$, ocupa pocas clases residuales modulo p , para muchos primos p , entonces debe estar esencialmente contenido en el conjunto de soluciones de una ecuación polinomial de grado acotado. Esto resuelve una pregunta de Helfgott y Venkatesh.

Palabras claves: Promedios ergódicos no convencionales, teoremas de descomposición, secuencias polinomiales, problema inverso de criba, cribas multidimensionales.

Discrete inverse theorems

Abstract

The present thesis studies two different instances of discrete inverse theorems, the first one pertaining convergence issues in ergodic theory and the second one problems of local distribution in number theory.

The first result allows us to characterize those functions that can form nonconventional ergodic averages with large L^2 norm. We then discuss abstract structure-randomness decompositions and extend them to the context of ergodic theory, by allowing for different levels of structure to be handled simultaneously. Combining these tools with the aforementioned inverse theorem and an adequate induction procedure, we are able to show that multiple polynomial ergodic averages arising from nilpotent groups of measure preserving transformations of a probability space always converge in norm. This answers a conjecture of Bergelson and Leibman.

The second result concerns the distribution of sets in residue classes. We introduce the concepts of characteristic and generic sets, and apply them in the framework of the larger sieve to show that if a big set of integer points $S \subseteq \{1, \dots, N\}^d$, $d > 1$, occupies few residue classes mod p , for many primes p , then it must essentially lie in the solution set of some polynomial equation of low degree. This settles a question of Helfgott and Venkatesh.

Keywords: Nonconventional ergodic averages, decomposition theorems, polynomial sequences, inverse sieve problem, high dimensional sieves.

A mis padres

Índice general

Introducción	8
Promedios ergódicos no convencionales	9
El problema inverso de criba	10
Introduction	13
Nonconventional ergodic averages	14
The inverse sieve problem	15
1 Decompositions and nonconventional ergodic averages	18
1.1 Introduction	18
1.1.1 Nonconventional ergodic averages	18
1.1.2 Historical background	21
1.1.3 Overview of the proof	22
1.2 Structure-randomness decompositions	23
1.2.1 Applying the Hahn-Banach theorem	23
1.2.2 Nested structures	26
1.3 Norm convergence for systems of finite complexity	29
1.3.1 Reducibility	31
1.3.2 Proof of Theorem 1.13	35
1.4 The complexity of polynomial systems	36
1.5 Further results	41
1.6 Some examples of reductions	42
2 The inverse sieve problem	47
2.1 Introduction	47
2.1.1 Sieves and algebraic structure	48
2.1.2 Organization	50
2.2 A conditional proof of Theorem 1.1	50
2.2.1 General notation	50
2.2.2 Characteristic sets	51
2.3 Applying the larger sieve in high dimensions	54
2.3.1 A review of the larger sieve	54
2.3.2 Genericity	56

2.4	The proof of Proposition 2.2	59
2.4.1	Summary of notation	59
2.4.2	Building good sections	60
2.4.3	Construction of the characteristic set	63
2.5	Ill-distributed sets with no algebraic structure	65
2.6	Further results and conjectures	66
2.6.1	A generalization of Theorem 1.1	66
2.6.2	Approximate reduction	67
2.6.3	The inverse sieve problem in one dimension	67
	Bibliografía	70

Introducción

La presente tesis se concentra en el estudio de estructuras discretas en la teoría ergódica y la teoría de números. Para explicar este concepto, es importante notar que muchos de los problemas centrales en estas áreas buscan mostrar que los objetos de interés se comportan acorde a estimaciones probabilísticas y principios de equidistribución. El objetivo entonces es entender en qué forma estos objetos pueden agruparse dentro de estructuras discretas que conspiran contra tales heurísticas.

Este punto de vista lleva naturalmente a la consideración de teoremas inversos discretos. El objetivo de estos teoremas inversos es proveer una caracterización de aquellos elementos que no se comportan en forma aleatoria con respecto a nuestro problema de interés. Un tema notable que es recurrente en esta dirección, y efectivamente aparecerá en esta tesis, es que tales obstrucciones a la aleatoriedad frecuentemente manifiestan lo que puede llamarse 'rigidez algebraica'. Una primera aproximación a este fenómeno es la afirmación de que únicamente elementos prescritos a satisfacer fuertes restricciones algebraicas pueden comportarse en forma anormal.

Por supuesto, un principio de esta generalidad es aplicable a una gran variedad de problemas, pero muchas de sus manifestaciones pueden ser unificadas bajo el título de Combinatoria Aritmética. Ejemplos notables de teoremas inversos discretos en esta área incluyen teoremas de tipo Freiman (ver por ejemplo [6, 24, 29, 46] y la exposición [21]), teoremas inversos para la norma de Gowers [8, 25, 50] y la teoría inversa de Littlewood-Offord [48, 49].

Esta tesis contiene dos nuevas instancias de teoremas inversos discretos, una correspondiente a la teoría ergódica y la segunda a la teoría de números. En el primer caso, caracterizamos aquellas funciones que dan lugar a promedios ergódicos no convencionales con valores grandes en la norma L^2 . Luego aplicamos esta caracterización para establecer la convergencia de estos promedios. En el segundo caso, caracterizamos aquellos conjuntos en dimensiones altas que están mal distribuidos en clases residuales módulo p , para muchos primos p . Ambos resultados evidencian el tipo de rigidez algebraica discutida con anterioridad. En el caso ergódico, estas funciones especiales están restringidas por relaciones algebraicas entre las transformaciones involucradas. En el segundo caso, los elementos del conjunto deben estar esencialmente contenidos en el conjunto de soluciones de una ecuación polinomial

de grado acotado.

A continuación proveemos una breve descripción de ambos resultados.

Promedios ergódicos no convencionales

Sea $T : X \rightarrow X$ una transformación que preserve la medida de un espacio de probabilidad (X, Σ, μ) . El teorema ergódico clásico de von Neumann nos dice que los promedios ergódicos

$$\frac{1}{N} \sum_{n=1}^N f(T^n x),$$

siempre convergen en $L^2(X)$, para toda elección de $f \in L^2(X)$. En mayor generalidad, supongamos que nos es dado un grupo G de transformaciones que preservan la medida de (X, Σ, μ) . ¿Podemos garantizar también que los promedios de la forma

$$\frac{1}{N} \sum_{n=1}^N f_1(T_1^{p_1(n)} x) f_2(T_2^{p_2(n)} x) \dots f_l(T_l^{p_l(n)} x),$$

siempre convergerán en $L^2(X)$, para toda elección de $T_1, \dots, T_l \in G$, $f_1, \dots, f_l \in L^\infty(X)$ y polinomios a valores enteros $p_1(n), \dots, p_l(n) : \mathbb{Z} \rightarrow \mathbb{Z}$?

El estudio de estos promedios ergódicos 'no convencionales' se origina con el trabajo de Furstenberg sobre el teorema de Szemerédi [18] y efectivamente, existe un gran número de trabajos sobre este problema, motivados en parte por las conexiones con la combinatoria y la teoría de números (ver §1.1). De todas formas, el problema de convergencia se mantuvo abierto aún en el caso abeliano.

El objetivo del Capítulo 1 es establecer la convergencia de estos promedios para todo grupo nilpotente G . Precisamente, probamos el siguiente resultado.

Teorema ([53]). *Sea G un grupo nilpotente de transformaciones que preservan la medida de un espacio de probabilidad (X, \mathcal{X}, μ) . Entonces, para todo $T_1, \dots, T_l \in G$, los promedios*

$$\frac{1}{N} \sum_{n=1}^N \prod_{j=1}^d \left(T_1^{p_{1,j}(n)} \circ \dots \circ T_l^{p_{l,j}(n)} \right) f_j,$$

siempre convergen en $L^2(X, \mathcal{X}, \mu)$, para todo $f_1, \dots, f_d \in L^\infty(X, \mathcal{X}, \mu)$ y todo conjunto de polinomios a valores enteros $p_{i,j}$.

Esto responde en forma afirmativa una conjetura de Bergelson y Leibman [7], quienes también demostraron que incluso los promedios ergódicos no convencionales más simples pueden divergir si G sólo se asume soluble.

Un papel importante en la demostración es jugado por un sencillo teorema inverso discreto (inspirado en [46]) que nos permite caracterizar aquellas funciones que

pueden originar un promedio ergódico no convencional cuya norma L^2 es larga para algún tiempo N grande (ver Lema 1.15). Demostramos que una tal función debe estar restringida por relaciones algebraicas entre las transformaciones involucradas y esto lleva al concepto de funciones reducibles (ver §1.3.1).

La propiedad crucial de estas funciones reducibles es que dan lugar a sistemas que son en principio más accesibles. La tarea de establecer este hecho rigurosamente es llevada a acabo en §1.3. El problema se reduce entonces a demostrar que mediante una iteración del anterior proceso podemos trasladar los promedios de interés a sistemas más sencillos sobre los cuales es trivial establecer convergencia. Este proceso de inducción es desarrollado abstractamente en §1.4, con varios ejemplos concretos presentados en §1.6.

Para poder implementar el esquema anterior necesitamos introducir herramientas que nos permiten descomponer funciones arbitrarias en componentes 'estructurados' y 'aleatorios'. En general, para obtener aplicaciones útiles de teoremas inversos es deseable poder llevar a cabo tal descomposición. En principio, esto debería permitirnos distinguir el componente de una función que contiene el tipo de estructura señalada por el teorema inverso, de la parte aleatoria de esta función, permitiendo de esta forma la implementación de estrategias específicas en cada componente. Por ejemplo, en el problema discutido en esta sección, la parte estructurada corresponde a funciones reducibles, para las cuales un proceso inductivo adecuado es aplicado.

A pesar de que los teoremas de descomposición del tipo mencionado anteriormente poseen una larga historia, un tratamiento particularmente atractivo de estos temas fue dado recientemente por Gowers [20]. Allí, él hace la importante observación de que muchos de los resultados de descomposición presentes en la literatura se vuelven mucho más transparentes mediante la aplicación del teorema de Hahn-Banach. En §1.2 presentamos una discusión de estas ideas y las adaptamos al contexto de nuestro problema. En particular, para la aplicación a promedios ergódicos, necesitamos poder manejar diferentes niveles de estructura en forma simultánea y las herramientas a este efecto son desarrolladas en §1.2.2.

El problema inverso de criba

Nuestro segundo teorema inverso estudia la distribución de conjuntos en clases residuales a módulo primo. Este es un tópico importante de la teoría analítica de números. En general, es de esperar que los conjuntos de interés estén bien distribuidos en clases residuales. Por ejemplo, esperamos que los primos estén uniformemente distribuidos en clases residuales primitivas. Esto no sólo es una pregunta interesante en sí misma, sino que también suele conllevar aplicaciones significativas al ser combinado con métodos de criba.

Por supuesto, no todos los conjuntos resultan bien distribuidos tras reducirlos a módulos primos. Por ejemplo, si consideramos el conjunto de los cuadrados, es

bien sabido que este conjunto ocupa únicamente $(p + 1)/2$ clases residuales módulo p para todo primo impar p . Asimismo, sabemos por la desigualdad de Lang-Weil (la cual es equivalente a la Hipótesis de Riemann sobre cuerpos finitos en el caso de curvas) que el conjunto de soluciones en \mathbb{Z}^d a una ecuación polinomial de grado acotado puede ocupar únicamente $O(p^k)$ clases residuales al ser reducida módulo p , para cierto entero $k < d$.

Estas observaciones llevaron a Croot y Elsholtz [14] e, independientemente, a Helfgott y Venkatesh [30] en un contexto más general, a formular la notable conjetura de que la única manera en la cual un conjunto grande de puntos enteros puede estar mal distribuido en clases residuales es si este posee una fuerte estructura algebraica (ver §2.1.1 y la Conjetura 2.16).

En el Capítulo 2 probamos esta conjetura para toda dimensión $d \geq 2$. El enunciado preciso es el siguiente.

Teorema ([52]). *Sean $0 \leq k < d$ enteros y sean $\varepsilon, \alpha, \eta > 0$ números reales positivos. Entonces, existe una constante C dependiendo sólo de los parámetros anteriores, tal que para todo conjunto $S \subseteq [N]^d$ ocupando menos de αp^k clases residuales para todo primo p , por lo menos una de las siguientes afirmaciones es válida:*

- (*S es pequeño*) $|S| \ll_{d,k,\varepsilon,\alpha} N^{k-1+\varepsilon}$,
- (*S es fuertemente algebraico*) Existe un polinomio no nulo $f \in \mathbb{Z}[x_1, \dots, x_d]$ de grado a lo sumo C y coeficientes acotados por N^C que se anula en al menos $(1 - \eta)|S|$ puntos de S .

En §2.5 presentamos un número de ejemplos mostrando que este resultado es óptimo. Proveemos también varios refinamientos y generalizaciones de este resultado (ver el Teorema 2.5 y especialmente §2.6).

El caso $k = 1$ de este resultado fue resuelto por Helfgott y Venkatesh en [30]. Su método se basa en ideas de Bombieri y Pila [10] que no se extienden a dimensiones más altas. El problema, que es recurrente en este tipo de situaciones, es que los métodos de criba suelen basarse en argumentos de contar y en consecuencia requieren que el número de clases ocupadas por el conjunto sea pequeño, mientras que en la presente situación únicamente sabemos que la densidad de estas clases es pequeña. Para corregir esto, la estrategia presentada en el Capítulo 2 procede haciendo uso de la estructura específica del conjunto estudiado (ver §2.3.2).

Un papel crucial en la demostración es desempeñado por el concepto de conjuntos característicos (Definición 2.2). En §2.4 demostramos que un conjunto mal distribuido suficientemente 'genérico' debe admitir un pequeño subconjunto tal que si un polinomio de grado acotado se anula en este subconjunto, entonces también debe anularse en una proporción positiva de los puntos del conjunto original (ver la Proposición 2.3). Dado que cualquier conjunto suficientemente pequeño es el conjunto de soluciones de una ecuación polinomial de grado acotado (ver Lema 2.4)

esto se encarga ya del caso de conjuntos 'genéricos'. El resto del trabajo se reduce entonces a demostrar que cualquier conjunto grande mal distribuido admite un subconjunto genérico denso (lo cual se establece en §2.3.2 y §2.4.2).

Introduction

The present thesis is concerned with the study of discrete structures in ergodic theory and number theory. To understand what is meant by this, it is important to notice that many of the central problems in these areas seek to show that the objects of interest behave in accordance with probabilistic estimates and principles of equidistribution. The goal then is to understand in what way these objects may arrange themselves into discrete structures that conspire against these heuristics.

This point of view leads naturally into the consideration of discrete inverse theorems. The aim of these inverse theorems is to provide a characterization of those elements that fail to behave in a random manner with respect to our problem of interest. A remarkable theme that is recurrent in this direction, and will indeed appear in this thesis, is that such obstructions to randomness often manifest what may be termed as 'algebraic rigidity'. A first approximation to this phenomenon may be the statement that only elements prescribed to satisfy strong algebraic restrictions may behave in an abnormal manner.

Of course, such a general principle applies to a wide variety of problems, but many of its manifestations can be unified under the heading of Arithmetic Combinatorics. Notable examples of discrete inverse theorems in this area include Freiman type theorems (see for instance [6, 24, 29, 46] and the survey [21]), inverse theorems for the Gowers norm [8, 25, 50] and the inverse Littlewood-Offord theory [48, 49].

This thesis contains two further instances of discrete inverse theorems, one pertaining to ergodic theory and the second one to number theory. In the first case, we characterize functions giving rise to nonconventional ergodic averages with large L^2 -norm. We then apply this characterization to establish the convergence of these averages. In the second case, we characterize high dimensional sets that are badly distributed in residue classes mod p , for many primes p . Both results evidence the algebraic rigidity discussed above. In the ergodic case, these special functions are constrained by algebraic relations between the transformations involved. In the second case, the elements of the set must essentially be contained inside the solution set of some polynomial equation of low degree.

We provide below a brief discussion of both results.

Nonconventional ergodic averages

Let $T : X \rightarrow X$ be a measure preserving transformation of a probability space (X, Σ, μ) . The classical mean ergodic theorem of von Neumann tells us that the ergodic averages

$$\frac{1}{N} \sum_{n=1}^N f(T^n x),$$

always converge in $L^2(X)$, for every choice of $f \in L^2(X)$. More generally, suppose we are given a group G of measure preserving transformations of (X, Σ, μ) . Can we also guarantee that averages like

$$\frac{1}{N} \sum_{n=1}^N f_1 \left(T_1^{p_1(n)} x \right) f_2 \left(T_2^{p_2(n)} x \right) \dots f_l \left(T_l^{p_l(n)} x \right),$$

will always converge in $L^2(X)$, for every choice of $T_1, \dots, T_l \in G$, $f_1, \dots, f_l \in L^\infty(X)$ and integer valued polynomials $p_1(n), \dots, p_l(n) : \mathbb{Z} \rightarrow \mathbb{Z}$?

The study of these 'nonconventional' ergodic averages arises with the work of Furstenberg on Szemerédi's theorem [18] and indeed, there is a large body of work on the above problem, motivated in part by its connections with combinatorics and number theory (see §1.1). Nevertheless, the problem of convergence remained open even in the abelian case.

The aim of Chapter 1 is to establish the convergence of the above averages for every nilpotent group G . Precisely, we prove the following result.

Theorem ([53]). *Let G be a nilpotent group of measure preserving transformations of a probability space (X, \mathcal{X}, μ) . Then, for every $T_1, \dots, T_l \in G$, the averages*

$$\frac{1}{N} \sum_{n=1}^N \prod_{j=1}^d \left(T_1^{p_{1,j}(n)} \circ \dots \circ T_l^{p_{l,j}(n)} \right) f_j,$$

always converge in $L^2(X, \mathcal{X}, \mu)$, for every $f_1, \dots, f_d \in L^\infty(X, \mathcal{X}, \mu)$ and every set of integer valued polynomials $p_{i,j}$.

This answers in the affirmative a conjecture of Bergelson and Leibman [7], who also showed that even the simplest nonconventional ergodic averages may diverge if G is only assumed to be solvable.

An important part of the proof is played by an easy discrete inverse theorem (inspired by [46]) that allows us to characterize those functions that can originate nonconventional ergodic averages with big L^2 norm at some large time N (see Lemma 1.15). We show that such a function must be constrained by algebraic relations between the transformations involved and this leads to the concept of reducible functions (see §1.3.1).

The key feature of these reducible functions is that they give rise to systems that in principle are more amenable to study. The task of establishing this fact rigorously is accomplished in §1.3. The problem is then reduced to showing that by an iteration of the above procedure one can translate the averages of interest to simpler systems over which it is trivial to establish convergence. This induction process is performed abstractly in §1.4, with several concrete examples presented in §1.6.

In order to develop the above scheme we need to introduce tools that allow us to decompose arbitrary functions into a 'structured' and a 'random' component. In general, in order to obtain useful applications from inverse theorems it is desirable to be able to perform such decompositions. This should permit us to distinguish the component of a function that contains the kind of structure signaled by the inverse theorem, from the random part of this function, thus allowing for specific strategies to be implemented in each component. For instance, in the problem discussed in this section, the structured part corresponds to reducible functions, for which an adequate induction is to be applied.

Although decomposition theorems of the above kind have a long history, a particularly appealing treatment was given recently by Gowers [20]. There, he makes the important observation that many of the decomposition results in the literature become much clearer by means of the Hahn-Banach theorem. In §1.2 we present a discussion of these ideas and adapt them to the context of our problem. In particular, for the application to ergodic averages, we need to be able to handle different levels of structure simultaneously and the tools for this purpose are developed in §1.2.2.

The inverse sieve problem

Our second inverse theorem studies the distribution of sets in residue classes to prime moduli. This is an important topic of analytic number theory. In general, we expect our sets of interest to be well distributed in residue classes. For instance, we expect the primes to be uniformly distributed in primitive residue classes. This is not only an interesting question in itself but can also lead to striking applications when combined with sieve methods.

Of course, not all sets are well distributed when reduced to prime moduli. For example, if we consider the set of squares, it is well known that this set only occupies $(p+1)/2$ residue classes mod p for every odd prime p . Also, we know by the Lang-Weil inequality (which is equivalent to the Riemann Hypothesis over finite fields in the case of curves) that the set of solutions in \mathbb{Z}^d to a polynomial equation of low degree may only occupy $O(p^k)$ residue classes when reduced mod p , for some integer $k < d$.

These observations led Croot and Elsholtz [14] and, independently, Helfgott and

Venkatesh [30] in a more general context, to the remarkable conjecture that the only way a big set of integer points may be badly distributed in residue classes is for it to possess some strong algebraic structure (see §2.1.1 and Conjecture 2.16).

In Chapter 2 we prove this conjecture for every dimension $d \geq 2$. The statement reads as follows.

Theorem ([52]). *Let $0 \leq k < d$ be integers and let $\varepsilon, \alpha, \eta > 0$ be positive real numbers. Then, there exists a constant C depending only on the above parameters, such that for any set $S \subseteq [N]^d$ occupying less than αp^k residue classes for every prime p , at least one of the following holds:*

- (*S is small*) $|S| \ll_{d,k,\varepsilon,\alpha} N^{k-1+\varepsilon}$,
- (*S is strongly algebraic*) *There exists a nonzero polynomial $f \in \mathbb{Z}[x_1, \dots, x_d]$ of degree at most C and coefficients bounded by N^C vanishing at more than $(1 - \eta)|S|$ points of S .*

In §2.5 we give a number of examples showing that this result is sharp. We also provide several refinements and generalizations of this result (see Theorem 2.5 and especially §2.6).

The case $k = 1$ of this result was handled by Helfgott and Venkatesh in [30]. Their method is based on ideas of Bombieri and Pila [10] and fails to extend to higher dimensions. The problem, which is natural in such a situation, is that sieve methods are usually based on counting arguments and therefore require the number of classes occupied by the set to be small, while in the present situation we only know that the density of these classes is small. In order to remedy this, the strategy presented in Chapter 2 proceeds by making use of the specific structure of the set being studied (see §2.3.2).

A crucial role in the proof is played by the concept of a characteristic subset (Definition 2.2). In §2.4 we show that a sufficiently 'generic' set that is badly distributed must admit a small subset such that if a polynomial of low degree vanishes at this subset, then it must also vanish at a large proportion of the points of the original set (see Proposition 2.3). Since any sufficiently small set is the solution set of some polynomial of low degree (see Lemma 2.4) this already handles the case of 'generic' sets. The rest of the proof is then reduced to showing that any large set that is badly distributed admits a dense generic subset (this is accomplished in §2.3.2 and §2.4.2).

Chapter 1

Decompositions and nonconventional ergodic averages

1.1 Introduction

The aims of this chapter are two-fold. In Section §1.2 we introduce and discuss the general philosophy of decomposition theorems, the goal of which is to decompose an arbitrary object of interest into well-defined components that are easier to handle: the so called 'random' and 'structured' parts of the object. As we will see, this idea is intimately linked with the concept of inverse theorems. Then, in the rest of the chapter, we will see how this philosophy can be applied to ergodic theory, more precisely, to the study of nonconventional ergodic averages.

1.1.1 Nonconventional ergodic averages

The subject of ergodic theory is concerned with the study of a probability space (X, Σ, μ) and (possibly a family of) measure preserving transformations $T : X \rightarrow X$ on this space. In this context, the main interest lies in understanding how does the space X evolves along iterations of the map T , which could be regarded as the 'time' parameter. In the foundations of ergodic theory lies the following classical result of von Neumann.

Theorem 1.1 (von Neumann's mean ergodic theorem). *Let $T : X \rightarrow X$ be a measure preserving transformation of a probability space (X, Σ, μ) . Then, the ergodic averages*

$$\frac{1}{N} \sum_{n=1}^N f(T^n x),$$

converge in $L^2(X)$ for every choice of $f \in L^2(X)$.

Here $L^2(X)$ stands for the space of square integrable functions in the probability space (X, Σ, μ) . The question we will be interested in is the extent to which this

result can be extended to more general families of transformations. Thus, let G be a group of measure preserving transformations of the space (X, Σ, μ) . There are three distinct directions in which the desired generalization may be accomplished.

- (*multiple compositions*). Given $T_1, \dots, T_l \in G$, do the averages

$$\frac{1}{N} \sum_{n=1}^N (T_1^n \circ T_2^n \circ \dots \circ T_l^n) f,$$

converge in $L^2(X)$ for every $f \in L^\infty(X)$?

- (*multiple functions*). Given $T_1, \dots, T_l \in G$, do the averages

$$\frac{1}{N} \sum_{n=1}^N T_1^n f_1 \dots T_l^n f_l,$$

converge in $L^2(X)$ for every $f_1, \dots, f_l \in L^\infty(X)$?

- (*polynomial orbits*). What about the averages

$$\frac{1}{N} \sum_{n=1}^N T_1^{p_1(n)} f_1 T_2^{p_2(n)} f_2,$$

where $T_1, T_2 \in G$, $f_1, f_2 \in L^\infty(X)$ and $p_1, p_2 : \mathbb{Z} \rightarrow \mathbb{Z}$ are integer valued polynomials?

The goal here is to understand what conditions should be imposed on G in order to guarantee the convergence of the above averages. Notice that, in contrast with the statement of Theorem 1.1, we assume our functions to lie in $L^\infty(X)$ instead of $L^2(X)$. This in fact is necessary, as it can easily be seen from the product of functions being unbounded in the L^2 -norm (but see Theorem 1.20).

It is important to remark that while the limiting behavior of these averages is interesting from a strictly ergodic theoretical point of view, it turns out to be deeply related to interesting phenomena in number theory and combinatorics. The most remarkable manifestation of this connection involves the following fundamental result of Szemerédi [44].

Theorem 1.2 (Szemerédi's theorem). *Let $A \subseteq \mathbb{Z}$ be a subset of the integers of positive upper density, that is, such that*

$$\limsup_{N \rightarrow \infty} \frac{|A \cap [-N, N]|}{|N|} > 0.$$

Then, A contains arbitrarily long arithmetic progressions.

Shortly after this result was obtained, a different proof was found by Furstenberg [18], who connected it to the study of ergodic averages. In fact, he was able to prove the following result, and to show that it implies Szemerédi's theorem.

Theorem 1.3 (Furstenberg's multiple recurrence theorem). *Let $T : X \rightarrow X$ be a measure preserving transformation of a probability space (X, Σ, μ) . Let $A \in \Sigma$ be such that $\mu(A) > 0$ and write 1_A for the characteristic function of this set. Then, for every integer k ,*

$$\liminf_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \int T^n 1_A T^{2n} 1_A \dots T^{kn} 1_A d\mu > 0.$$

For a nice discussion of the fruitful interactions between ergodic theory, combinatorics and number theory originating from this work, the reader is referred to the survey articles [36, 45].

Given a group G and two elements $x, y \in G$, we write $[x, y] := x^{-1}y^{-1}xy$ for the commutator of x and y , and define the lower central series of G as $G_0 := G$ and $G_i := [G, G_{i-1}]$ for every $i \in \mathbb{N}$. We recall that a group G is said to be nilpotent if there exists some finite integer r such that $G_r = \{1_G\}$, where 1_G is the identity of G . The integer r is called the nilpotency class of G .

It turns out that the assumption of G being nilpotent is particularly relevant to the study of the ergodic averages mentioned above. Indeed, it was conjectured that this is the right assumption to be placed on G in order to guarantee convergence of these averages. The main purpose of this chapter is to prove this claim in the form of the following result.

Theorem 1.4. *Let G be a nilpotent group of measure preserving transformations of a probability space (X, \mathcal{X}, μ) . Then, for every $T_1, \dots, T_l \in G$, the averages*

$$\frac{1}{N} \sum_{n=1}^N \prod_{j=1}^d \left(T_1^{p_{1,j}(n)} \circ \dots \circ T_l^{p_{l,j}(n)} \right) f_j, \quad (1.1.1)$$

always converge in $L^2(X, \mathcal{X}, \mu)$, for every $f_1, \dots, f_d \in L^\infty(X, \mathcal{X}, \mu)$ and every set of integer valued polynomials $p_{i,j}$.

Notice that the averages in (1.1.1) do indeed generalize the three directions discussed above. This result was conjectured in the present form by Bergelson and Leibman, who also showed that even

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N T^n f S^n g,$$

need not exist if T and S only generate a solvable group [7].

1.1.2 Historical background

Partial results towards Theorem 1.4 have a rich history. Notice that when $d = l = 1$ and the polynomial is linear it reduces to Theorem 1.1. The only case of Theorem 1.4 that was fully settled is that in which $T_1 = \dots = T_l$, that is, when G is a cyclic group. The study of this case originated in the seminal work of Furstenberg [18] on Szemerédi's theorem, while a general solution when the polynomials are linear was later provided by Host and Kra [32] following the work of several authors (with a different proof subsequently found by Ziegler [55]). Convergence for general polynomials was established by Bergelson [5] under the assumption of weakly mixing, while the first unconditional non-linear result was obtained by Furstenberg and Weiss [19]. The general result for cyclic groups and arbitrary polynomials was finally settled by Host and Kra [33] and Leibman [39].

Another case of Theorem 1.4 that is known is that in which G is abelian and every polynomial is linear. Here, the case $d = 2$ was proven by Conze and Lesigne [13] and assuming extra ergodicity hypothesis on the transformations Zhang [54] gave a proof for $d = 3$ and Frantzikinakis and Kra [15] for general d . Without these assumptions, this result was established by Tao [46] and by now possesses several different proofs [2, 31, 51]. However, when G is abelian but the polynomials are arbitrary, very little was known. It was shown by Chu, Frantzikinakis and Host [11] that

$$\frac{1}{N} \sum_{n=1}^N T_1^{p_1(n)} f_1 \dots T_l^{p_l(n)} f_l \quad (1.1.2)$$

converges whenever the polynomials p_i have distinct degrees, but the convergence of (1.1.2) has remained open for arbitrary polynomials. Notice that (1.1.2) corresponds to taking $p_{i,j} = 0$ whenever $i \neq j$ in Theorem 1.4. More generally, very little was known until now for convergence of \mathbb{Z}^d actions along polynomials. A particular result in this direction is the convergence of the averages

$$\frac{1}{N} \sum_{n=1}^N T^{n^2} f \left(T^{n^2} S^n \right) g,$$

which was established by Austin [3, 4].

Finally, when G is only assumed to be nilpotent the results are much scarcer. Prior to the result discussed in this chapter, it was known by the work of Bergelson and Leibman [7] that the averages

$$\frac{1}{N} \sum_{n=1}^N T^n f S^n g,$$

always converge in L^2 , but even in the linear case no convergence result has been previously established for more than two transformations.

1.1.3 Overview of the proof

Our proof of Theorem 1.4 does not make use of the aforementioned results and therefore provides an alternative proof of these statements, which in many cases is substantially simpler than the original ones. In particular, we do not make use of the machinery of characteristic factors that is heavily used in previous literature. The price we pay in doing so is that we do not obtain any explicit description of the limits. In this sense, our approach is similar to that of Tao [46], in that we use a weak inverse theorem (see Lemma 1.15) to decompose our functions into the sum of a random component, which is easily treatable, and a structured one, which can be handled by an adequate induction. Interestingly, we find that our decomposition is best carried out by adapting ideas of Gowers related to the Hahn-Banach separation theorem [20] and this is done in §1.2. This is arguably the first time that these ideas are used in a purely ergodic theoretical context.

The main new ingredient of the proof is the concept of an L -reducible function (Definition 1.14), which will play the role of the structured component of our decompositions. We refer to §1.3 for precise definitions, but for now let us discuss what these are in the linear abelian case. Here, an L -reducible function σ with respect to a set of transformations T_1, \dots, T_j , is a function for which the behavior of $T_j^n \sigma$ can be somewhat recovered from that of the set $T_1^n b_1, \dots, T_{j-1}^n b_{j-1}$, for some prescribed set of functions b_i . This way, the problem of convergence for the set of transformations T_1, \dots, T_j is reduced to the analogous question for the smaller set T_1, \dots, T_{j-1} , and one may then proceed inductively. The details of these reductions are carried out in §1.3.

When either G is not abelian or the polynomials are not linear, the system of transformations to which L -reducible functions allow us to pass does not admit such a simple expression. In general, it will consist of twice as many transformations as the original one and the degree of the polynomials involved may not necessarily decrease, so that it may seem that we have not gained much with this procedure. As it turns out however, one can define a suitable notion of complexity for every set of transformations and show that the above process does indeed lead us to a set of lower complexity. The proof that every system of transformations of the type studied in Theorem 1.4 reduces in finitely many steps to one consisting only of the identity transformation $Ix = x$ is performed in §1.4, and this completes the proof of Theorem 1.4.

The methods discussed immediately shield some further convergence results and these are discussed in Section §1.5. We also include in Section §1.6 several examples of how the induction process mentioned in the previous paragraph works in some concrete cases.

1.2 Structure-randomness decompositions

Intrinsically related with the notion of an inverse theorem, is the idea of decomposing a given object into a structured and a random component. Indeed, since the purpose of an inverse theorem is to characterize those objects or functions that possess a particular property that is relevant to the problem at hand, perhaps as a potential obstruction to certain type of arguments, it is then desirable to be able to decompose an arbitrary object into a structured part that manifests this property in some way, and a random part that has no trace of such a potential obstruction. This reduces the general problem in two instances that allow for specific strategies to be applied.

The purpose of this section is to present some of the tools developed by Gowers in this direction and extend them to the context of our problem. A crucial observation of Gowers is that many of the decomposition arguments present in the literature become much cleaner by means of the Hahn-Banach theorem, and this will be the point of view taken here. For a better discussion of these ideas the reader is referred to Gowers' article [20].

1.2.1 Applying the Hahn-Banach theorem

We will concentrate on the study of a real Hilbert space \mathcal{H} with norm $\|\cdot\|$. In our application to ergodic averages, this space will be given by $\mathcal{H} = L^2(X, \mu)$, the space of square integrable functions in the probability space (X, μ) . Nevertheless, for the present discussion we allow \mathcal{H} to be arbitrary.

We let $\Sigma \subseteq \mathcal{H}$ be a bounded set in the norm $\|\cdot\|$. The choice of Σ is arbitrary, but once this is made we consider the elements of Σ as 'structured'. Of course, in order to obtain meaningful applications the choice of the distinguished set Σ should be appropriate for the problem being studied, but we emphasize that the arguments and results to be presented in this section are independent of such a choice, and this accounts for their usefulness. A fruitful example to keep in mind is that of Fourier analysis, where the set of structured elements Σ corresponds to the characters.

Our goal is to decompose an arbitrary element $f \in \mathcal{H}$ in the form $f = u + v$, where u resembles the elements of Σ , while v is a 'random' component, in the sense that it has a small correlation with these elements. Of course, a natural possibility would be to require in this decomposition that u itself belongs to Σ , but this turns out to be too weak for applications. In fact, it is often desirable to have some flexibility, since it is natural to expect that if $\sigma_1, \sigma_2 \in \Sigma$ are structured elements, then their sum $\sigma_1 + \sigma_2$ will also retain some of this structure even if this sum does not itself belong to Σ . Because of this, it is useful to allow u to be the sum of a few elements of Σ .

We can now give an informal statement of the type of decompositions we are looking for.

Theorem 1.5 (simple decomposition, informal statement). *Let Σ be some bounded set in \mathcal{H} . Then, every bounded $f \in \mathcal{H}$ may be written as*

$$f = \sum_{j=1}^k \lambda_j \sigma_j + v,$$

where

- $\sigma_j \in \Sigma$ for every j ,
- $\sum_{j=1}^k |\lambda_j|$ is bounded,
- $|\langle v, \sigma \rangle|$ is small for every $\sigma \in \Sigma$.

As we mentioned, there is a nice analogy with Fourier analysis, where we can decompose any bounded f into its large Fourier coefficients and an element v that has small correlation with every Fourier character.

An important observation of Gowers is that the decompositions of the type described above are closely related to the study of certain kind of norms. In order to make this observation precise, given some norm $\|\cdot\|_X$ on \mathcal{H} equivalent to $\|\cdot\|$, we define its dual norm by

$$\|f\|_X^* := \sup_{\|g\|_X \leq 1} |\langle f, g \rangle|.$$

Notice that $\|\cdot\|_X^*$ is then also equivalent to $\|\cdot\|$.

We have the following result.

Lemma 1.6 (cf. [20, Corollary 3.5]). *Let $\Sigma \subseteq \mathcal{H}$ be a bounded set and suppose the norm*

$$\|f\|_\Sigma := \inf \left\{ \sum_{j=0}^{k-1} |\lambda_j| : f = \sum_{j=0}^{k-1} \lambda_j \sigma_j, \sigma_j \in \Sigma \right\}, \quad (1.2.1)$$

is well defined and equivalent to $\|\cdot\|$. Then its dual norm is given by $\|f\|_\Sigma^ = \sup_{\sigma \in \Sigma} |\langle f, \sigma \rangle|$.*

Proof. Given some $f \in \mathcal{H}$ it is clear on one hand that

$$\sup_{\sigma \in \Sigma} |\langle f, \sigma \rangle| \leq \sup_{\|g\|_\Sigma \leq 1} |\langle f, g \rangle|.$$

On the other hand, for every $\epsilon > 0$, if $g = \sum_{j=0}^{k-1} \lambda_j \sigma_j$ with $\sum_{j=0}^{k-1} |\lambda_j| < 1 + \epsilon$, then $|\langle f, g \rangle| \leq (1 + \epsilon) \sup_{\sigma \in \Sigma} |\langle f, \sigma \rangle|$. The result follows. \square

What this lemma shows is that the definitions of structure and randomness suggested in Theorem 1.5 are dual to each other in a very precise sense. An element having a small norm $\|\cdot\|_\Sigma$ will be structured, while a small value of the dual norm $\|\cdot\|_\Sigma^*$ characterizes random objects. In particular, notice that this makes both sets convex.

In order to exploit this fact, we will make use of the classical Hahn-Banach theorem.

Theorem 1.7 (Geometric Hahn-Banach). *Let A be an open convex subset containing 0 of a real topological vector space V and suppose $v \in V$ does not lie in A . Then there exists some continuous linear functional $\phi : V \rightarrow \mathbb{R}$ such that $\phi(v) \geq 1$ and $\phi(w) < 1$ for every $w \in A$.*

The idea of Gowers to obtain decompositions can roughly be described as follows. While it may be difficult to check directly whether an arbitrary function can be described by the sum of a structured and a random component, if such a decomposition fails to exist an application of the Hahn-Banach theorem would allow us to find some large functional which does not correlate with random functions (therefore having a kind of structure itself) nor with structured functions (therefore also having some randomness). This way, we are only left with proving that no object can be random and structured at the same time, which generally tends to be an easier task.

To apply this scheme, we will need the following refinement of Theorem 1.7.

Corollary 1.8 (cf. [20, Corollary 3.2]). *Let A_1, \dots, A_n be open convex subsets containing 0 of some real Hilbert space \mathcal{H} . Let $c_1, \dots, c_n > 0$ be positive real numbers and suppose $f \in \mathcal{H}$ cannot be written as $\sum_{j=1}^n c_j f_j$ with $f_j \in A_j$. Then there exists some $\phi \in \mathcal{H}$ such that $\langle \phi, f \rangle \geq 1$ and $\langle \phi, g_i \rangle < c_i^{-1}$ for every $g_i \in A_i$.*

Proof. Since the set $A := \sum_{i=1}^n c_i A_i$ will be an open convex set in \mathcal{H} containing 0 but not f , it follows by the Hahn-Banach theorem that there exists some $\phi \in \mathcal{H}$ satisfying $\langle \phi, f \rangle \geq 1$ and $\langle \phi, g \rangle < 1$, for every $g \in A$. The result follows immediately, since $c_i g_i \in A$ for every $g \in A_i$. \square

We can now prove the following formal version of Theorem 1.5 (see [20, Proposition 3.6])

Theorem 1.9 (Simple decomposition, formal statement). *Let $\|\cdot\|_X$ be some norm on \mathcal{H} equivalent to $\|\cdot\|$ and let $C > 0$ be an arbitrary positive real number. Then, every $f \in \mathcal{H}$ with $\|f\| \leq 1$ may be written as $f = u + v$, with $\|u\|_X < C$ and $\|v\|_X^* < C^{-1}$.*

Proof. Write

$$A_1 := \{u \in \mathcal{H} : \|u\|_X < 1\},$$

and

$$A_2 := \{v \in \mathcal{H} : \|v\|_X^* < 1\},$$

for the respective unit balls and assume the claim fails. It follows from Corollary 1.8 that there exists some linear functional φ such that $\langle \varphi, f \rangle > 1$, $\langle \varphi, u \rangle \leq C^{-1}$ for every $u \in A_1$ and $\langle \varphi, v \rangle \leq C$ for every $v \in A_2$. But by Cauchy-Schwarz, the boundedness of f and the definition of dual norms, this implies that

$$1 < \langle \varphi, f \rangle \leq \|\varphi\| \leq \|\varphi\|^* \|\varphi\|^{**} \leq 1,$$

giving us the desired contradiction. \square

It follows from Lemma 1.6 that this indeed corresponds to Theorem 1.5. The only issue lies in the fact that we cannot ensure that the norms $\|\cdot\|_\Sigma$ are always well defined and equivalent to $\|\cdot\|$. However, this is not a serious problem as long as we are prepared to tolerate a small error term, since we can simply adjoin a small ball in the norm $\|\cdot\|$ to the set Σ (see (1.3.8)). The real inconvenient is that this decomposition is too weak as it stands. In fact, it would be desirable to have stronger bounds on the norms of the components, and also to be able to handle several notions of structure at the same time. In the next subsection, we will see how all these can be accomplished.

1.2.2 Nested structures

We now proceed to obtain a strong decomposition theorem as it will be needed for the study of nonconventional ergodic averages. There will be two main improvements over the kind of decompositions we obtained in §1.2.1. First of all, we will obtain much sharper bounds on the size of the norms. The second improvement is more suggestive and it allows us to study simultaneously a large family of distinguished sets. In fact, instead of considering a fixed set Σ we will study a nested family $\Sigma_1 \supseteq \Sigma_2 \supseteq \dots \supseteq \Sigma_n \supseteq \dots$. These should be understood as different levels of structure, with an element $\sigma \in \Sigma_n$ being more structured the larger the value of n is.

The reason why we need to consider such a family simultaneously instead of simply fixing a large value of n is that, while it is certainly advantageous for the elements σ_j in Theorem 1.5 to belong to Σ_n for some large n , the property of the element v in Theorem 1.5 of being random becomes weaker as n increases. In fact, since this amounts to avoiding high correlations with elements of Σ_n , this information is weaker the more restricted the set Σ_n is. Because of this, it is desirable in practice to arrange for the σ_j to lie in Σ_B , for some large value of B , while ensuring that the component v stays random with respect to Σ_A , for a much smaller value of A .

It is often the case that allowing a small error term permits one to obtain a sharper estimate for the main expression being studied. This is also the case here

and fortunately such a small error term can be easily dealt with in the applications. As before, we begin with an informal statement of the decomposition we are looking for.

Theorem 1.10 (strong decomposition, informal statement). *Let $\Sigma_1 \supseteq \Sigma_2 \supseteq \dots \supseteq \dots \supseteq \Sigma_n \supseteq \dots$ be a family of bounded subsets of \mathcal{H} and let $M > 0$ be some positive integer. Then, there exists a pair of integers A, B with A much smaller than M and B much larger than M , such that every bounded element $f \in \mathcal{H}$ may be written as*

$$f = \sum_{j=1}^k \lambda_j \sigma_j + f_2 + f_3,$$

where

- $\sigma_j \in \Sigma_B$ for every $1 \leq j \leq k$,
- $\sum_{j=1}^k |\lambda_j|$ is bounded,
- $|\langle f_2, \sigma \rangle|$ is very small for every $\sigma \in \Sigma_A$,
- f_3 is a small error term (i.e. $\|f_3\|$ is small).

Some remarks are in order. The reason we emphasize that the expression $|\langle f_2, \sigma \rangle|$ is very small is because we will indeed be able to make this arbitrarily small in terms of the bounds on $\sum_{j=1}^k |\lambda_j|$ (see Proposition 1.11). On the other hand, there is a slight inaccuracy in our choice of integers A, B since we will not be able to obtain such a decomposition for an arbitrary value of M , but instead we will deduce the assertion for some M in a bounded region. Finally, it is important to notice that the constants we shall obtain in the decomposition are absolute, in the sense that they *do not* depend on the choice of the family $\Sigma_1 \supseteq \Sigma_2 \supseteq \dots \supseteq \dots \Sigma_n \supseteq \dots$. This amusing phenomenon will be of much use later.

We now turn to the details. Given a positive real number δ and some decreasing function $\eta : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, we will consider the sequence of real numbers $C_1^{\delta, \eta}, \dots, C_{\lceil 2\delta^{-2} \rceil}^{\delta, \eta}$ defined recursively by

$$C_{\lceil 2\delta^{-2} \rceil}^{\delta, \eta} := 1, C_{n-1}^{\delta, \eta} := \max \{ C_n^{\delta, \eta}, 2\eta(C_n^{\delta, \eta})^{-1} \}. \quad (1.2.2)$$

We shall also write $C^{\delta, \eta} := C_1^{\delta, \eta}$. These constants will provide the parameters for the decomposition obtained below and the fact that they are independent of the specific family of sets, as anticipated in the previous paragraph, will allow us to do *a priori* modifications on our set of structured functions so that they are better suited to the resulting bounds.

By an appeal to Lemma 1.6 the problem we are studying reduces to the consideration of an infinite family of norms $(\|\cdot\|_N)_{N \in \mathbb{N}}$ measuring increasing rates of

structure and for which their dual norms $(\|\cdot\|_N^*)_{N \in \mathbb{N}}$ measure decreasing rates of randomness. As we have discussed, we are interested in studying this large family of norms *simultaneously*, so that if we know one of the components is random at a level A (that is, $\|\cdot\|_A^*$ is small), we need the other component to be structured at a much higher level B (that is, $\|\cdot\|_B$ must be small for some B much larger than A). This is accomplished by the following result.

Proposition 1.11 (Strong decomposition). *Let $(\|\cdot\|_N)_{N \in \mathbb{N}}$ be a family of norms on \mathcal{H} equivalent to $\|\cdot\|$ and satisfying $\|\cdot\|_{N+1}^* \leq \|\cdot\|_N^*$ for every N . Let $0 < \delta, c < 1$ be positive real numbers, $\eta : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ some decreasing function and $\psi : \mathbb{N} \rightarrow \mathbb{N}$ some function satisfying $\psi(N) \geq N$ for all N . Then, for every integer $M_\bullet > 0$, there exists a sequence*

$$M_\bullet \leq M_1 \leq \dots \leq M_{\lceil 2\delta^{-2} \rceil} \leq M^\bullet = O_{M_\bullet, \delta, c, \psi}(1),$$

which does not depend on the specific family of norms, with the property that for any $f \in \mathcal{H}$ with $\|f\| \leq 1$, we can find some $1 \leq i \leq \lceil 2\delta^{-2} \rceil$ and integers A, B with

$$M_\bullet \leq A < cM_i < \psi(M_i) \leq B,$$

such that we have the decomposition $f = f_1 + f_2 + f_3$ with

$$\|f_1\|_B < C_i^{\delta, \eta}, \|f_2\|_A^* < \eta(C_i^{\delta, \eta}), \|f_3\| < \delta.$$

Proof. Our proof is modeled on the proof of Proposition 3.5 of [20]. Set $A_1 := M_\bullet$, $M_1 := \lceil c^{-1}A_1 + 1 \rceil$ and $B_1 := \psi(M_1)$. If there is no decomposition of the desired form with these parameters and $C_1 := C_1^{\delta, \eta}$, we may apply Corollary 1.8 to obtain some $\phi_1 \in \mathcal{H}$ such that

- $\langle \phi_1, f \rangle \geq 1$,
- $\|\phi_1\|_{B_1}^* \leq C_1^{-1}$,
- $\|\phi_1\|_{A_1}^{**} \leq \eta(C_1)^{-1}$,
- $\|\phi_1\| \leq \delta^{-1}$,

where we are using the fact that if $\|\cdot\|_N$ is some norm equivalent to $\|\cdot\|$, then $\{f \in \mathcal{H} : \|f\|_N < 1\}$ is an open convex set in \mathcal{H} containing 0.

Recursively, if we cannot find a decomposition with parameters $A_{j-1}, M_{j-1}, B_{j-1}, C_{j-1}$ we set $A_j := B_{j-1}$, $M_j := \lceil c^{-1}A_j + 1 \rceil$, $B_j := \psi(M_j)$ and $C_j := C_j^{\delta, \eta}$. If no such decomposition exists with these parameters we can then use Corollary

1.8 to find some $\phi_j \in \mathcal{H}$ with properties analogous to the ones above. This way we construct a sequence of elements obeying the orthogonality relationships

$$\begin{aligned} |\langle \phi_j, \phi_i \rangle| &\leq \|\phi_j\|_{A_j}^{**} \|\phi_i\|_{A_j}^* \\ &\leq \|\phi_j\|_{A_j}^{**} \|\phi_i\|_{B_i}^* \\ &\leq \eta(C_j)^{-1} C_i^{-1} \\ &\leq 1/2, \end{aligned}$$

whenever $i < j$, by construction of C_k . But then, by the bounds on $\|\phi_i\|$, we obtain upon expanding the inner product

$$\|\phi_1 + \dots + \phi_r\|^2 \leq \delta^{-2}r + \frac{r^2 - r}{2}, \quad (1.2.3)$$

for each $r \leq \lceil 2\delta^{-2} \rceil$. On the other hand, the condition $\langle \phi_i, f \rangle \geq 1$ for all i implies that the left-hand side of (1.2.3) is at least r^2 . Since this is absurd for $r = \lceil 2\delta^{-2} \rceil$ the result follows. \square

1.3 Norm convergence for systems of finite complexity

From now on fix a nilpotent group G and a probability space X as in the statement of Theorem 1.4. By a G -sequence we shall mean a sequence $\{g(n)\}_{n \in \mathbb{Z}}$ taking values in G . An ordered tuple $\mathbf{g} = (g_1, \dots, g_j)$ of G -sequences will be called a system, and for each system one can ask whether the corresponding ergodic averages

$$\mathcal{A}_N^{\mathbf{g}}[f_1, \dots, f_j] := \mathbb{E}_{n \in [N]} \prod_{i=1}^j g_i(n) f_i, \quad (1.3.1)$$

converge in $L^2(X)$ for every $f_1, \dots, f_j \in L^\infty(X)$. Here, for a finite set A we write

$$\mathbb{E}_{x \in A} f(x) := \frac{1}{|A|} \sum_{x \in A} f(x),$$

and for every positive integer N it is $[N] := \{1, \dots, N\}$. We say two systems are equivalent if they consist of the same G -sequences, so for example if g, h are G -sequences then the system (h, g) is equivalent to the system (g, h) , and so is (g, h, h) . Clearly, the convergence of the averages of the form (1.3.1) for some system implies the convergence of the averages associated to every equivalent system, since $T(f_1)T(f_2) = T(f_1 f_2)$ for every $T \in G$ and $f_1, f_2 \in L^\infty(X)$.

To each pair of G -sequences g, h we will associate, for each positive integer m , the G -sequence

$$\langle g|h \rangle_m(n) := g(n)g(n+m)^{-1}h(n+m),$$

and we define the m -reduction of a system $\mathbf{g} = (g_1, \dots, g_j)$ to be the system

$$\mathbf{g}_m^* = (g_1, \dots, g_{j-1}, \langle g_j | 1_G \rangle_m, \langle g_j | g_1 \rangle_m, \dots, \langle g_j | g_{j-1} \rangle_m),$$

where by a slight abuse of notation we write 1_G for the G -sequence $1_G(n) := 1_G$, where 1_G is the identity of G . The main purpose of this section will be to show that one can deduce the convergence of the averages (1.3.1) for some system \mathbf{g} from knowing this (actually, the slightly stronger Theorem 1.13 below) for every reduction \mathbf{g}_m^* of \mathbf{g} . This leads us to define the complexity of a system.

Definition 1.12 (Complexity of a system). We say a system \mathbf{g} has complexity 0 if it is equivalent to the trivial system (1_G) (that is, the system consisting only on the sequence 1_G). Recursively, we say a system \mathbf{g} has complexity d , for some positive integer $d \geq 1$, if it is not of complexity d' for any $0 \leq d' < d$ and it is equivalent to some system \mathbf{h} for which every reduction \mathbf{h}_m^* has complexity $\leq d - 1$. We say a system has finite complexity if it has complexity d for some integer $d \geq 0$.

Given a system $\mathbf{g} = (g_1, \dots, g_j)$, some set of functions $f_1, \dots, f_j \in L^\infty(X)$ and a pair of integers N, N' , write

$$\mathcal{A}_{N, N'}^{\mathbf{g}}[f_1, \dots, f_j] := \mathcal{A}_{N'}^{\mathbf{g}}[f_1, \dots, f_j] - \mathcal{A}_N^{\mathbf{g}}[f_1, \dots, f_j].$$

We have the following result.

Theorem 1.13. *Let G and X be as above and let $d \geq 0$. Let $F : \mathbb{N} \rightarrow \mathbb{N}$ be some non-decreasing function with $F(N) \geq N$ for all N and let $\varepsilon > 0$ be some positive real number. Then, for every integer $M > 0$, there exists a sequence of integers*

$$M \leq M_1^{\varepsilon, F, d} \leq \dots \leq M_{K_{\varepsilon, d}}^{\varepsilon, F, d} \leq M^{\varepsilon, F, d} = O_{d, F, \varepsilon, M}(1), \quad (1.3.2)$$

for some $K_{\varepsilon, d} = O_{\varepsilon, d}(1)$, such that for every system $\mathbf{g} = (g_1, \dots, g_j)$ of complexity at most d and every choice of functions $f_1, \dots, f_j \in L^\infty(X)$ with $\|f_i\|_\infty \leq 1$, there exists some $1 \leq i \leq K_{\varepsilon, d}$ such that

$$\left\| \mathcal{A}_{N, N'}^{\mathbf{g}}[f_1, \dots, f_j] \right\|_{L^2(X)} \leq \varepsilon, \quad (1.3.3)$$

for every $M_i^{\varepsilon, F, d} \leq N, N' \leq F(M_i^{\varepsilon, F, d})$.

This type of statement already appears in the works of Tao [46] and of Avigad, Gerhardy and Towsner [1]. Clearly, Theorem 1.13 implies that the averages (1.3.1) converge in $L^2(X)$ for every system \mathbf{g} of finite complexity, since otherwise one could find some $\varepsilon > 0$ and some increasing function $F : \mathbb{N} \rightarrow \mathbb{N}$ such that

$$\left\| \mathcal{A}_{N, F(N)}^{\mathbf{g}}[f_1, \dots, f_d] \right\|_{L^2(X)} > \varepsilon,$$

for every integer N . The usefulness of Theorem 1.13 lies on its uniformity over all systems of a fixed complexity, which plays an important role in the inductive argument. In fact, the ergodic averages (1.3.1) associated to a system \mathbf{g} for which the reductions \mathbf{g}_m^* do not satisfy stability bounds which are uniform on m may not necessarily converge, even if the ergodic averages associated to each individual reduction \mathbf{g}_m^* do converge.

The rest of this section is devoted to the proof of Theorem 1.13. In §1.4 we will show that every system of the form given in Theorem 1.4 has finite complexity, thereby completing the proof of that theorem.

1.3.1 Reducibility

Since Theorem 1.13 is trivially true when $d = 0$, we may proceed by induction. Thus, let $d > 0$ be some positive integer and assume the result holds for every $d' < d$. Let F and $0 < \varepsilon < 1$ be as in the statement of the theorem and let $\mathbf{g} = (g_1, \dots, g_j)$ be some system of complexity at most d . Since it clearly suffices to prove the result for any system equivalent to \mathbf{g} , by definition of the complexity we may assume without loss of generality that \mathbf{g}_m^* has complexity $\leq d - 1$ for every positive integer m .

Let C^* denote the quantity $C^{\delta, \eta}$ defined in (1.2.2) associated to $\delta := \varepsilon/(2^5 3)$ and $\eta(x) := \varepsilon^2/(2^3 3^3 x)$, so that in particular C^* depends only on ε . We will sometimes use the shorthands $\|\cdot\|_\infty$ for $\|\cdot\|_{L^\infty(X)}$, $\|\cdot\|_2$ for $\|\cdot\|_{L^2(X)}$ and $\langle \cdot, \cdot \rangle$ for $\langle \cdot, \cdot \rangle_{L^2(X)}$. The following definition will be crucial.

Definition 1.14 (reducible functions). Given a positive integer L , we say $\sigma \in L^\infty(X)$, $\|\sigma\|_\infty \leq 1$, is an L -reducible function (with respect to \mathbf{g}), if there exists some integer $M > 0$ and a family $b_0, b_1, \dots, b_{j-1} \in L^\infty(X)$ with $\|b_i\|_\infty \leq 1$, such that for every positive integer $l \leq L$

$$\left\| g_j(l)\sigma - \mathbb{E}_{m \in [M]} (\langle g_j | 1_G \rangle_m(l)) b_0 \prod_{i=1}^{j-1} (\langle g_j | g_i \rangle_m(l)) b_i \right\|_{L^\infty(X)} < \frac{\varepsilon}{16C^*}.$$

Reducible functions will play a similar role than the one played by basic anti-uniform functions in [46]. We stress that we do not care for the value of M in Definition 1.14. We will show in Lemma 1.15 below that every function giving rise to a large average must resemble a reducible function. The main feature of these objects is that the role of the G -sequence g_j on the averages (1.3.1) can essentially be recovered by means of the set of G -sequences $\langle g_j | 1_G \rangle_m, \langle g_j | g_1 \rangle_m, \dots, \langle g_j | g_{j-1} \rangle_m$.

Lemma 1.15 (Weak inverse result for ergodic averages). *Assume the inequality*

$$\|\mathcal{A}_N^{\mathbf{g}}[f_1, \dots, f_{j-1}, u]\|_2 > \varepsilon/6,$$

holds for some $\|u\|_{L^\infty(X)} \leq 3C$, some $1 \leq C \leq C^$ and some $f_1, \dots, f_{j-1} \in L^\infty(X)$ with $\|f_i\|_\infty \leq 1$. Then, there exists some constant $0 < c_1 < 1$, depending only on ε ,*

such that for every positive integer $L < c_1 N$ there is an L -reducible function σ with $\langle u, \sigma \rangle > 2\eta(C)$.

Proof. We begin by noticing that $\|\mathcal{A}_N^{\mathbf{g}}[f_1, \dots, f_{j-1}, u]\|_2^2 = \langle u, h \rangle$, where

$$h := \mathbb{E}_{n \in [N]} g_j(n)^{-1} \mathcal{A}_N^{\mathbf{g}}[f_1, \dots, f_{j-1}, u] \prod_{i=1}^{j-1} g_j(n)^{-1} g_i(n) f_i. \quad (1.3.4)$$

We claim $\sigma := h/3C$ is an L -reducible function for every $L < c_1 N$ and some $0 < c_1 < 1$ depending only on ε , from where the result immediately follows since by the observation above it is $\langle u, \sigma \rangle > 2\eta(C)$.

It remains to prove this claim. Write $c_1 := \frac{\varepsilon}{96(C^*)^2}$ and assume $0 < l < c_1 N$. Then, if we shift $[N]$ to $[N] + l$ we see that the right hand side of (1.3.4) changes by a magnitude of at most $6lC^*/N < \varepsilon/(16C^*)$ (since $\|\mathcal{A}_N^{\mathbf{g}}[f_1, \dots, f_{j-1}, u]\|_\infty \leq 3C \leq 3C^*$) and thus

$$\begin{aligned} & \left\| h - \mathbb{E}_{n \in [N]} g_j(l+n)^{-1} \mathcal{A}_N^{\mathbf{g}}[f_1, \dots, f_{j-1}, u] \right. \\ & \quad \left. \times \prod_{i=1}^{j-1} g_j(l+n)^{-1} g_i(l+n) f_i \right\|_{L^\infty(X)} < \frac{\varepsilon}{16C^*}. \end{aligned}$$

Applying $g_j(l)$ we get

$$\begin{aligned} & \left\| g_j(l)h - \mathbb{E}_{n \in [N]} (\langle g_j | 1_G \rangle_n(l)) \mathcal{A}_N^{\mathbf{g}}[f_1, \dots, f_{j-1}, u] \right. \\ & \quad \left. \times \prod_{i=1}^{j-1} (\langle g_j | g_i \rangle_n(l)) f_i \right\|_{L^\infty(X)} < \frac{\varepsilon}{16C^*}. \end{aligned}$$

The claim then follows with $M := N$, $b_0 := \frac{1}{3C} \mathcal{A}_N[f_1, \dots, f_{j-1}, u]$ and $b_i := f_i$. \square

As mentioned early, the advantage of L -reducible functions is that they allow us to reduce the study of the ergodic averages of \mathbf{g} to the study of averages arising from the reductions \mathbf{g}_m^* , which we already know to satisfy uniform stability bounds by the induction hypothesis. This idea is carried out in the next proposition.

Proposition 1.16. *For every positive integer M_* there exists a sequence*

$$M_* \leq M_1 \leq \dots \leq M_{\tilde{K}} \leq M^* = O_{M_*, \varepsilon, d, F}(1), \quad (1.3.5)$$

depending only on M_*, ε, d and F , and with \tilde{K} depending only on ε and d , such that if $f_1, \dots, f_{j-1} \in L^\infty(X)$ with $\|f_i\|_\infty \leq 1$ and

$$f = \sum_{t=0}^{k-1} \lambda_t \sigma_t,$$

where $\sum_{t=0}^{k-1} |\lambda_t| \leq C^*$, and each σ_t is an L -reducible function for some $L \geq F(M^*)$, then there exists some $1 \leq i \leq \tilde{K}$ such that

$$\|\mathcal{A}_{N,N'}^{\mathbf{g}}[f_1, \dots, f_{j-1}, f]\|_{L^2(X)} \leq \varepsilon/4,$$

for every pair $M_i \leq N, N' \leq F(M_i)$.

Proof. For every σ_t let $M^{(t)}$ be the integer coming from the definition of an L -reducible function and let $b_i^{(t)} \in L^\infty(X)$ be the corresponding family of functions. It follows from the definition of an L -reducible function that for every $N \leq L$ and every $0 \leq t \leq k-1$ we may replace $\mathcal{A}_N^{\mathbf{g}}[f_1, \dots, f_{j-1}, \sigma_t]$ by

$$\mathbb{E}_{m \in [M^{(t)}]} \left[\mathbb{E}_{n \in [N]} \left(\prod_{i=1}^{j-1} g_i(n) f_i \right) \left((\langle g_j | 1_G \rangle_m(n)) b_0^{(t)} \right) \right. \\ \left. \times \left(\prod_{i=1}^{j-1} (\langle g_j | g_i \rangle_m(n)) b_i^{(t)} \right) \right],$$

at the cost of an L^∞ error of at most $\varepsilon/(16C^*)$. Therefore, we get by Minkowski's inequality that for $N, N' \leq L$, the expression

$$\|\mathcal{A}_{N,N'}^{\mathbf{g}}[f_1, \dots, f_{j-1}, f]\|_2,$$

is bounded by

$$\left(\sum_{t=0}^{k-1} |\lambda_t| \mathbb{E}_{m \in [M^{(t)}]} \left\| \mathcal{A}_{N,N'}^{\mathbf{g}_m^*} \left[f_1, \dots, f_{j-1}, b_0^{(t)}, b_1^{(t)}, \dots, b_{j-1}^{(t)} \right] \right\|_{L^2(X)} \right) + \varepsilon/8. \quad (1.3.6)$$

We are thus given a large family of averages coming from the lower complexity systems \mathbf{g}_m^* . Write $\gamma := \frac{\varepsilon}{16C^*}$. Clearly, it would suffice to find a suitable interval on which each of this lower dimensional averages is bounded by γ . Although this will not be possible, we will indeed show by repeated applications of the induction hypothesis that we can get such a bound for all but a negligible subset of these averages. In order to do this, consider non-decreasing functions $F_1, \dots, F_r : \mathbb{N} \rightarrow \mathbb{N}$, for some $r = O_{\varepsilon,d}(1)$ to be specified, defined recursively by $F_r := F$ and

$$F_{i-1}(N) := \max_{1 \leq M \leq N} F_i(M^{\gamma, F_i, d-1}),$$

where we are using the notation in the statement of Theorem 1.13. Also, let $K := K_{\gamma, d-1}$ be as in that theorem and for each tuple $1 \leq i_1, \dots, i_s \leq K$, $s \leq r$, and integer M , we define recursively

$$M^{(i_1, \dots, i_s)} := \left(\left(\left(M_{i_1}^{\gamma, F_{i_1}, d-1} \right)_{i_2}^{\gamma, F_{i_2}, d-1} \right) \dots \right)_{i_s}^{\gamma, F_{i_s}, d-1}.$$

Thus, $M^{(i_1)}$ is the integer $M_{i_1}^{\gamma, F_1, d-1}$ obtained in (1.3.2) by starting at M , $M^{(i_1, i_2)}$ is the integer $M_{i_2}^{\gamma, F_2, d-1}$ obtained by starting the sequence (1.3.2) at $M = M^{(i_1)}$, etc. In particular, notice that this sequence depends only on ε, F, d and M . Observe also that since each of the averages in (1.3.6) satisfies

$$\left\| \mathcal{A}_{N, N'}^{\mathbf{g}_m} \left[f_1, \dots, f_{j-1}, b_0^{(t)}, b_1^{(t)}, \dots, b_{j-1}^{(t)} \right] \right\|_{L^\infty(X)} \leq 1, \quad (1.3.7)$$

the sum on (1.3.6) is bounded by $\sum_{t=0}^{k-1} |\lambda_t| \leq C^*$.

We now proceed as follows. By the induction hypothesis we know that each of the reduced averages in (1.3.6) is bounded by γ for every pair $N, N' \in [M_*^{(i)}, F_1(M_*^{(i)})]$ and some $1 \leq i \leq K$, which depends on the particular average. By the pigeonhole principle and (1.3.7), this implies that we may find some $1 \leq i_1 \leq K$ such that the contribution to (1.3.6) of those averages which are not bounded by γ for every pair $N, N' \in [M_*^{(i_1)}, F_1(M_*^{(i_1)})]$ is at most $\left(\frac{K-1}{K}\right) C^*$. We now apply the induction hypothesis to these remaining averages with the function F_2 , the parameter γ and the starting point $M_*^{(i_1)}$. This way, for each of these remaining averages, we know that there exists some $1 \leq i \leq K$ such that the average is bounded by γ for every pair $N, N' \in [M_*^{(i_1, i)}, F_2(M_*^{(i_1, i)})]$. Since by construction of F_1 it is

$$[M_*^{(i_1, i)}, F_2(M_*^{(i_1, i)})] \subseteq [M_*^{(i_1)}, F_1(M_*^{(i_1)})],$$

we see that those averages which we bounded in the previous step remain bounded by γ on each of these new intervals. Thus, we may apply the pigeonhole principle as before to find some $1 \leq i_2 \leq K$ such that the contribution to (1.3.6) of those averages which are not bounded by γ for every pair $N, N' \in [M_*^{(i_1, i_2)}, F_2(M_*^{(i_1, i_2)})]$ is at most $\left(\frac{K-1}{K}\right)^2 C^*$.

Iterating the above process r times, we find a tuple $1 \leq i_1, \dots, i_r \leq K$ such that the set of reduced averages which are not bounded by γ for every pair

$$N, N' \in [M_*^{(i_1, \dots, i_r)}, F_r(M_*^{(i_1, \dots, i_r)})] = [M_*^{(i_1, \dots, i_r)}, F(M_*^{(i_1, \dots, i_r)})],$$

contributes at most

$$\left(\frac{K-1}{K}\right)^r C^* < \varepsilon/16,$$

to (1.3.6), upon choosing r sufficiently large in terms of ε and d . Since the sum over the remaining terms will be bounded by $\sum_{t=0}^{k-1} |\lambda_t| \gamma < \varepsilon/16$, we conclude that (1.3.6), and therefore

$$\left\| \mathcal{A}_{N, N'}^{\mathbf{g}} [f_1, \dots, f_{j-1}, f] \right\|_2 \leq \varepsilon/4,$$

for every $N, N' \in [M_*^{(i_1, \dots, i_r)}, F(M_*^{(i_1, \dots, i_r)})]$.

Notice that while the specific integer $M_*^{(i_1, \dots, i_r)}$ we have obtained depends on the set of functions f_1, \dots, f_{j-1}, f and the system \mathbf{g} , this integer belongs to the

sequence $M_*^{(j_1, \dots, j_r)}$, $1 \leq j_1, \dots, j_r \leq K$, which depends only on F, ε, d and M_* . The result follows from this observation with the sequence (1.3.5) given by the integers $M_*^{(j_1, \dots, j_r)}$, $1 \leq j_1, \dots, j_r \leq K$. \square

1.3.2 Proof of Theorem 1.13

We can now conclude the proof of Theorem 1.13. As it was done before, we fix X, G, F, ε, d and \mathbf{g} as in the statement of that theorem and assume without loss of generality that each reduction \mathbf{g}_m^* of \mathbf{g} is of complexity at most $d - 1$ and that the result is already proven for every $d' < d$. We will also write M_0 for the integer M to be chosen as the starting point of the sequence (1.3.2) in Theorem 1.13. Let δ, η be as specified at the beginning of §1.3.1 and write $C_i := C_i^{\delta, \eta}$ for the constants defined in (1.2.2). Given some positive integer L write Σ_L for the set of L -reducible functions and set

$$\Sigma_L^+ := \Sigma_L \cup B_2(\delta/C^*), \quad (1.3.8)$$

where we write $B_2(\delta/C^*)$ for the set of $f \in L^2(X)$ with $\|f\|_2 \leq \delta/C^*$. Consider on $L^2(X)$ the norms $\|\cdot\|_L := \|\cdot\|_{\Sigma_L^+}$ defined as in (1.2.1). It is easy to see that these norms are well defined and equivalent to $\|\cdot\|_{L^2(X)}$ (by the presence of the small L^2 ball and the fact that reducible functions are bounded by 1). Also, notice that $\Sigma_{L+1}^+ \subseteq \Sigma_L^+$ for every L which in turn implies (by Lemma 1.6) that $\|\cdot\|_{L+1}^* \leq \|\cdot\|_L^*$.

Given any integer M write $\psi(M) := F(M^*)$ where M^* is the integer obtained in Proposition 1.16 with $M_* = M$. Let $f_1, \dots, f_j \in L^\infty(X)$, $\|f_i\|_\infty \leq 1$, be given and consider for f_j a decomposition of the form provided in Proposition 1.11 with $(\|\cdot\|_L)_{L \in \mathbb{N}}$, ψ, δ, η as above and c equal to the constant c_1 in Lemma 1.15. This allows us to find a constant $1 \leq C_i \leq C^*$ and some integer M with $M_0 \leq M = O_{M_0, \varepsilon, F, d}(1)$, such that

$$f_j = \sum_{t=0}^{k-1} \lambda_t \sigma_t + u + v, \quad (1.3.9)$$

where $\sum_{t=0}^{k-1} |\lambda_t| \leq C_i$, each σ_t belongs to Σ_B^+ for some $B \geq \psi(M)$ (and therefore to $\Sigma_{\psi(M)}^+$), $\|u\|_A^* < \eta(C_i)$ for some $A < c_1 M$ and $\|v\|_2 < \delta$. We remark that this constant C_i is the one defined in (1.2.2) and that the integer M obtained belongs to the sequence given in Proposition 1.11, which does not depend on the family of norms $(\|\cdot\|_L)_{L \in \mathbb{N}}$ and in the present case is therefore independent of the particular system \mathbf{g} we have fixed (although it certainly depends on its complexity d , as well as on ε, F and M_0). Since $\|\sum_t^* \lambda_t \sigma_t\|_2 \leq \delta$, where the sum is restricted to those $\sigma_t \in B_2(\delta/C^*)$, we may assume that each σ_t in (1.3.9) actually belongs to $\Sigma_{\psi(M)}$, at the cost of softening our bound on v to $\|v\|_2 < 2\delta$.

We would like to use Lemma 1.15 to study the function u , but first we need to gain some control on its L^∞ norm. In order to do this, denote by $S \in \mathcal{X}$ the set of points on which the inequality $|v(s)| \leq C_i$ holds (in particular one has $\mu(S^c) < (2\delta/C_i)^2$)

and write $v' := u\mathbf{1}_{S^c} + v$. From the fact that $\|\sigma_j\|_{L^\infty(X)} \leq 1$ for every $\sigma_j \in \Sigma_{\psi(M)}$, (1.3.9) and the definition of S , one easily checks that $|u\mathbf{1}_{S^c}(x)| \leq 3|v(x)|$ a.e. and therefore $\|u\mathbf{1}_{S^c}\|_2 \leq 3\|v\|_2$. Hence, it follows that for every pair of integers N, N' ,

$$\begin{aligned} \|\mathcal{A}_{N,N'}[f_1, \dots, f_{j-1}, v']\|_2 &\leq \|\mathcal{A}_{N'}[f_1, \dots, f_{j-1}, v']\|_2 \\ &\quad + \|\mathcal{A}_N[f_1, \dots, f_{j-1}, v']\|_2 \\ &\leq 2(4\|v\|_2) \\ &< \varepsilon/3, \end{aligned} \tag{1.3.10}$$

where we are using Minkowski's inequality and the fact that $\|f_i\|_\infty \leq 1$ for every $1 \leq i \leq j-1$. Consider now $u\mathbf{1}_S$. Similarly as above, one sees that $\|u\mathbf{1}_S\|_{L^\infty(X)} \leq 3C_i$. Also, it follows from Lemma 1.6 that for every $\sigma \in \Sigma_A$ it is

$$\begin{aligned} |\langle u\mathbf{1}_S, \sigma \rangle| &\leq |\langle u, \sigma \rangle| + |\langle u\mathbf{1}_{S^c}, \sigma\mathbf{1}_{S^c} \rangle| \\ &\leq \|u\|_A^* + \|u\mathbf{1}_{S^c}\|_2 \|\sigma\mathbf{1}_{S^c}\|_2 \\ &< \eta(C_i) + 12\delta^2/C_i \\ &< 2\eta(C_i). \end{aligned}$$

We are now in a position to apply Lemma 1.15, which implies that for every pair $N, N' \geq M$

$$\begin{aligned} \|\mathcal{A}_{N,N'}[f_1, \dots, f_{j-1}, u\mathbf{1}_S]\|_2 &\leq \|\mathcal{A}_{N'}[f_1, \dots, f_{j-1}, u\mathbf{1}_S]\|_2 \\ &\quad + \|\mathcal{A}_N[f_1, \dots, f_{j-1}, u\mathbf{1}_S]\|_2 \\ &\leq \varepsilon/3. \end{aligned} \tag{1.3.11}$$

It only remains to analyze $\sum_{t=0}^{k-1} \lambda_t \sigma_t$. But we may now invoke Proposition 1.16 to conclude from our choice of ψ that

$$\left\| \mathcal{A}_{N,N'}[f_1, \dots, f_{j-1}, \sum_{t=0}^{k-1} \lambda_t \sigma_t] \right\|_{L^2(X)} < \varepsilon/3, \tag{1.3.12}$$

for every pair $M_i \leq N, N' \leq F(M_i)$ and some $M_i \in [M, \psi(M)]$ which belongs to the corresponding sequence (1.3.5). Theorem 1.13 then follows from (1.3.9), (1.3.10), (1.3.11), (1.3.12) and Minkowski's inequality.

1.4 The complexity of polynomial systems

In this section we will prove that every system of the form given in Theorem 1.4 has finite complexity, thereby finishing the proof of that theorem. In order to do this, we begin by reviewing some facts about polynomial sequences in nilpotent groups. For a detailed treatment of this topic, the reader is referred to the work of Leibman [37, 38].

For a G -sequence $g = (g(n))_{n \in \mathbb{Z}}$ taking values in a nilpotent group G and some integer m , we define the operator D_m which takes g to the G -sequence $(D_m g)(n) := g(n)g(n+m)^{-1}$. In particular, we have $\langle g|h \rangle_m(n) = (D_m g)(n)h(n+m)$, for every pair of G -sequences g, h and every positive integer m . We say that a G -sequence is *polynomial* if there exists some positive integer d such that for every choice of integers m_1, \dots, m_d , we have $D_{m_1} \dots D_{m_d} g = 1_G$, where we recall that 1_G stands for the constant sequence which equals the identity of G . It is known that if $(g(n))_{n \in \mathbb{Z}}$ is a sequence in a nilpotent group G which is of the form

$$g(n) = T_1^{p_1(n)} \dots T_k^{p_k(n)}, \quad (1.4.1)$$

where $T_1, \dots, T_k \in G$ and p_1, \dots, p_k is some set of integer valued polynomials, then g is a polynomial sequence. Indeed, each $T_i^{p_i(n)}$ is clearly a polynomial sequence and the product of polynomial sequences is polynomial by Lemma 1.17 below (the converse also holds, see for example [37]).

By a polynomial system we shall mean a system $\mathbf{g} = (g_1, \dots, g_j)$, where each g_i , $1 \leq i \leq j$, is a polynomial sequence. We define the size of such a system to be $|\mathbf{g}| = j$. To prove Theorem 1.4 it will suffice, by Theorem 1.13 and the fact that sequences of the form (1.4.1) are polynomial, to prove that every polynomial system has finite complexity.

In order to proceed, we will need to define the degree of a polynomial sequence. Unfortunately, the natural choice of taking the least positive integer d for which every d successive application of the above operators returns the identity is not appropriate for our purposes, since with this definition the set of polynomial sequences of degree $\leq d$ need not form a group. In order to amend this, we need to introduce some notation. Write $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ and $\mathbb{N}_* = \mathbb{N}_0 \cup \{-\infty\}$. We say a vector $\bar{d} = (d_1, \dots, d_k) \in \mathbb{N}_*^k$ is *superadditive* if $d_i \leq d_{i+1}$ for every $1 \leq i < k$ and $d_i + d_j \leq d_{i+j}$ for every pair i, j , where we are using the conventions $-\infty + t = -\infty$ for every $t \in \mathbb{N}_*$ and $-\infty < r$ for every $r \in \mathbb{N}_0$. Also, given a superadditive vector $\bar{d} = (d_1, \dots, d_k)$ and some nonnegative integer t , we write $\bar{d} - t := (d'_1, \dots, d'_k)$, where $d'_i = d_i - t$ if $t \leq d_i$ and $d'_i = -\infty$ otherwise. Notice that $\bar{d} - t$ so defined is also a superadditive vector.

Fix a nilpotent group G of nilpotency class s and let $G = G_1 \supset G_2 \supset \dots \supset G_s \supset G_{s+1} = \{1_G\}$ be its lower central series. As in [37, 38], we say a sequence $g = (g(n))_{n \in \mathbb{Z}}$ taking values in G is a polynomial sequence of (vector) degree $\leq (d_1, \dots, d_s)$ if $(D_{m_1} \dots D_{m_{d_k+1}} g)(n) \in G_{k+1}$ for every n , every $1 \leq k \leq s$ and every choice of $m_1, \dots, m_{d_k+1} \in \mathbb{Z}$. If $d_k = -\infty$ we take this to mean that g itself takes values in G_{k+1} . We will make use of the following results of Leibman.

Lemma 1.17 ([38, §3]). *Let $\bar{d} = (d_1, \dots, d_s)$ be a superadditive vector and let $t, t_1, t_2 \geq 0$ be nonnegative integers. Then we have the following properties.*

1. If g is a polynomial sequence of degree $\leq \bar{d} - t$, then $D_m g$ is a polynomial sequence of degree $\leq \bar{d} - (t + 1)$, for every $m \in \mathbb{Z}$.
2. The set of polynomial sequences of degree $\leq \bar{d} - t$ forms a group.
3. If g is a polynomial sequence of degree $\leq \bar{d} - t_1$ and h is a polynomial sequence of degree $\leq \bar{d} - t_2$, then $[g, h]$ is a polynomial sequence of degree $\leq \bar{d} - (t_1 + t_2)$, where $[g, h](n) := g^{-1}(n)h^{-1}(n)g(n)h(n)$.

Remark. The results of [38] concern the operators

$$(\tilde{D}_m g)(n) := g(n)^{-1}g(n+m) = (D_m g^{-1})(n).$$

Nevertheless, using Lemma 1.17 for these operators and a straightforward descending induction on t one can easily check that a G -sequence g has degree $\leq \bar{d} - t$ with respect to the operators \tilde{D}_m if and only if it has degree $\leq \bar{d} - t$ with respect to the operators D_m , from where we recover Lemma 1.17 as stated.

We say a polynomial system $\mathbf{g} = (g_1, \dots, g_j)$ has degree $\leq d$ if the degree of g_i is $\leq d$, for every $1 \leq i \leq j$. We will show that any system of degree $\leq \bar{d}$, for some superadditive vector $\bar{d} = (d_1, \dots, d_s)$, has finite complexity. Notice that this is enough to prove Theorem 1.4, since if a polynomial sequence g has degree $\leq (d_1, \dots, d_s)$, then it also has degree $\leq (d, 2d, \dots, sd)$, with $d = \max \{d_i : 1 \leq i \leq s\}$, and this last vector is clearly superadditive.

Given a polynomial system \mathbf{g} , we are concerned with the process that consists of passing from \mathbf{g} to an equivalent system \mathbf{g}' , then taking the m -reduction $(\mathbf{g}')_m^*$ of \mathbf{g}' for some m , passing to an equivalent system $((\mathbf{g}')_m^*)'$ and then taking the m' -reduction of this for some m' , etc. What we are free to choose in the above process is to which equivalent system we apply the reductions (but not the integer on which we subsequently reduce) and our objective is to show that there exists some constant C , depending on \mathbf{g} , such that for every sequence of positive integers m, m', m'', \dots we can go to the trivial system (1_G) by means of at most C repetitions of the above transformations. This clearly implies that the complexity of \mathbf{g} is at most C .

In order to simplify notation we will omit the reference to the specific sequence of integers on which we reduce. So for instance, we will generically refer to the reduction of a system $\mathbf{g} = (g_1, \dots, g_j)$ to be the system

$$\mathbf{g}^* = (g_1, \dots, g_{j-1}, \langle g_j | 1_G \rangle, \langle g_j | g_1 \rangle, \dots, \langle g_j | g_{j-1} \rangle).$$

Similarly, we have the identity

$$\langle g | h \rangle(n) = Dg(n)(Dh(n))^{-1}h(n),$$

provided, of course, that the omitted subindices are the same. We define a *step* to be the process of passing from a system \mathbf{g} to the reduction $(\mathbf{g}')^*$ of some system \mathbf{g}'

equivalent to \mathbf{g} . We will show that one can pass from a polynomial system \mathbf{g} to the trivial system in a number of steps which is bounded in terms of the size and degree of \mathbf{g} .

We define the complete reduction of a system \mathbf{g} to be the system

$$\mathbf{g}^{**} = (g_1, \dots, g_{j-1}, \langle g_j | g_1 \rangle, \dots, \langle g_j | g_{j-1} \rangle).$$

Thus, $\mathbf{g}^{**} = \mathbf{g}^* \setminus \{\langle g_j | 1_G \rangle\}$. We define a *complete step* in the same way as a step, but with the reduction replaced by the complete reduction. Complete steps are needed for a technical reason related to the inductive process to be applied. Precisely, in order to handle steps involving systems of degree $\leq \bar{d}$ we will need to assume some control on both steps and complete steps over systems of degree $\leq \bar{d} - 1$.

Theorem 1.4 follows from Theorem 1.13 and the following result.

Theorem 1.18. *Let \mathbf{g} be a polynomial system of size $|\mathbf{g}| \leq C_1$ and degree $\leq \bar{d}$, for some superadditive vector $\bar{d} = (d_1, \dots, d_s)$. Then,*

- *one can go from \mathbf{g} to the trivial system (1_G) in $O_{C_1, \bar{d}}(1)$ steps,*
- *one can go from \mathbf{g} to a system consisting of a single sequence of degree $\leq \bar{d}$ in $O_{C_1, \bar{d}}(1)$ complete steps,*

for every sequence of positive integers m, m', m'', \dots . In particular, \mathbf{g} has complexity $O_{C_1, \bar{d}}(1)$.

Proof. Let \bar{d} be as in the statement. We begin by noticing that the result is trivially true for systems of degree $\leq \bar{d} - (d_s + 1) = (-\infty, \dots, -\infty)$, since 1_G is the only sequence lying in G_{k+1} for every $1 \leq k \leq s$. We will proceed by induction. Since $\bar{d} - t$ is superadditive for every $t \geq 0$, it will suffice to prove that if Theorem 1.18 holds for systems of degree $\leq \bar{d} - 1$ then it also holds for systems of degree $\leq \bar{d}$.

Thus, let \mathbf{g} be as in the statement. We will first prove that we can go from \mathbf{g} to the trivial system in $O_{C_1, \bar{d}}(1)$ steps (and therefore, that \mathbf{g} has complexity $O_{C_1, \bar{d}}(1)$). In order to do this, observe that \mathbf{g} can be rewritten in the form

$$\mathbf{g} = \mathbf{h}_0 \cup \bigcup_{i=1}^l s_i \mathbf{h}_i, \tag{1.4.2}$$

for some polynomial sequences s_1, \dots, s_l of degree $\leq \bar{d}$, $l \leq C_1$, and some polynomial systems \mathbf{h}_i of degree $\leq \bar{d} - 1$ and size $\leq C_1$, with \mathbf{h}_0 possibly empty (for example, one could simply take $s_i = g_i$ and $h_i = (1_G)$ for every $1 \leq i \leq l$). Here, if $\mathbf{h} = (h_1, \dots, h_k)$, $s\mathbf{h}$ is the system (sh_1, \dots, sh_k) and the union of two systems $(h_1, \dots, h_k), (h'_1, \dots, h'_r)$ is understood to be the system $(h_1, \dots, h_k, h'_1, \dots, h'_r)$.

The idea will be to show that for systems of the form (1.4.2) one can perform steps in such a way that the resulting systems are also of the form (1.4.2) for the same set of sequences s_1, \dots, s_l . Furthermore, we will show that in finitely many steps we may actually discard the sequence s_l , therefore arriving at a system like (1.4.2) in which only the sequences s_1, \dots, s_{l-1} are present. Iterating this l times we shall then end up with a system of degree $\leq \bar{d} - 1$, from where one can proceed by induction.

In order to carry out this plan we begin by observing that if s_i, s_j are sequences of degree $\leq \bar{d}$ and h_i, h_j are sequences of degree $\leq \bar{d} - 1$, we have

$$\begin{aligned} \langle s_j h_j | s_i h_i \rangle &= D(s_j h_j) (D(s_i h_i))^{-1} s_i h_i \\ &= s_i D(s_j h_j) (D(s_i h_i))^{-1} [D(s_j h_j) (D(s_i h_i))^{-1}, s_i] h_i \\ &= s_i h^{j,i}, \end{aligned} \quad (1.4.3)$$

for some polynomial sequence $h^{j,i}$ which is seen to have degree $\leq \bar{d} - 1$ by Lemma 1.17. Furthermore, if $s_i = s_j = s$, it is easy to check that

$$\langle s h_j | s h_i \rangle = s \langle h_j | h_i \rangle.$$

It follows from these formulas that, provided $|\mathbf{h}_l| > 1$, the reduction \mathbf{g}^* of \mathbf{g} is equivalent to a system of the form

$$\mathbf{h}_0^{(1)} \cup \left(\bigcup_{i=1}^{l-1} s_i \mathbf{h}_i^{(1)} \right) \cup s_l \mathbf{h}_l^{**}, \quad (1.4.4)$$

for some systems $\mathbf{h}_0^{(1)}, \mathbf{h}_1^{(1)}, \dots, \mathbf{h}_{l-1}^{(1)}$ of degree $\leq \bar{d} - 1$ and size $|\mathbf{h}_0^{(1)}| \leq 2|\mathbf{h}_0| + 1$ and $|\mathbf{h}_i^{(1)}| \leq 2|\mathbf{h}_i|$ for every other i , and where we recall that \mathbf{h}_l^{**} refers to the complete reduction of \mathbf{h}_l . Explicitly, if $\mathbf{h}_i = (h_{i,1}, \dots, h_{i,j_i})$ for every $0 \leq i \leq l$, then

$$\mathbf{h}_0^{(1)} = (\langle s_l h_{l,j_l} | 1_G \rangle, h_{0,1}, \dots, h_{0,j_0}, \langle s_l h_{l,j_l} | h_{0,1} \rangle, \dots, \langle s_l h_{l,j_l} | h_{0,j_0} \rangle),$$

while $s_i \mathbf{h}_i^{(1)}$ equals

$$(s_i h_{i,1}, \dots, s_i h_{i,j_i}, \langle s_l h_{l,j_l} | s_i h_{i,1} \rangle, \dots, \langle s_l h_{l,j_l} | s_i h_{i,j_i} \rangle),$$

for every $1 \leq i \leq l - 1$. We see by (1.4.3) that this is of the desired form.

Observe now that if \mathbf{h} is equivalent to \mathbf{h}' then the system $s\mathbf{h}$ is also equivalent to $s\mathbf{h}'$. Since by the induction hypothesis we know that one can pass from \mathbf{h}_l to a system \mathbf{h} consisting of a single sequence of degree $\leq \bar{d} - 1$ in $O_{C_1, \bar{d}-1}(1)$ complete steps, it follows from the above observation and (1.4.4) that we may pass from \mathbf{g} to a system of the form

$$\mathbf{h}_0^{(2)} \cup \left(\bigcup_{i=1}^{l-1} s_i \mathbf{h}_i^{(2)} \right) \cup s_l \mathbf{h}, \quad (1.4.5)$$

in $O_{C_1, \bar{d}-1}(1)$ steps, where each system $\mathbf{h}_i^{(2)}$ has degree $\leq \bar{d} - 1$ and size $O_{C_1, \bar{d}}(1)$, and \mathbf{h} is a system consisting of a single sequence of degree $\leq \bar{d} - 1$. But then we see from (1.4.3) that the reduction of (1.4.5) will be of the form

$$\mathbf{h}_0^{(3)} \cup \left(\bigcup_{i=1}^{l-1} s_i \mathbf{h}_i^{(3)} \right),$$

with each $\mathbf{h}_i^{(3)}$ having degree $\leq \bar{d} - 1$ and size $O_{C_1, \bar{d}}(1)$. We have therefore succeeded in discarding the sequence s_l from our system. We can now repeat the same process as before with s_{l-1} in place of s_l . Since the size of $\mathbf{h}_{l-1}^{(3)}$ is $O_{C_1, \bar{d}}(1)$, we see that this new process finishes in $O_{C_1, \bar{d}}(1)$ steps, leaving us with a system of the form

$$\mathbf{h}_0^{(4)} \cup \left(\bigcup_{i=1}^{l-2} s_i \mathbf{h}_i^{(4)} \right).$$

Therefore, iterating the above process l times, we are finally left in $O_{C_1, \bar{d}}(1)$ steps with a system of degree $\leq \bar{d} - 1$ from where we may apply the induction hypothesis to obtain the trivial system in $O_{C_1, \bar{d}}(1)$ further steps, thereby completing the proof of the finite complexity of \mathbf{g} .

Now it only remains to show that one can pass from \mathbf{g} to a system consisting of a single sequence of degree $\leq \bar{d}$ in $O_{C_1, \bar{d}}(1)$ complete steps. But it is clear that the above reasoning to pass from \mathbf{g} to a system of degree $\leq \bar{d} - 1$ works in exactly the same way for complete steps, since the only things that may change are the systems $\mathbf{h}_0^{(1)}, \mathbf{h}_0^{(2)}, \mathbf{h}_0^{(3)}, \dots$, which nevertheless will always have degree $\leq \bar{d} - 1$ and whose size may only be smaller than in the previous case. Thus, the above reasoning allows us to pass to a system of degree $\leq \bar{d} - 1$ from where we may apply induction, as long as we are not left after any of the complete steps with a system which can be written in its entirety as $s_i \mathbf{h}$, for some s_i as above and \mathbf{h} of degree $\leq \bar{d} - 1$ and size $|\mathbf{h}| = 1$ (because if the whole system has size 1 the complete reduction is not defined). But since in such a case we are already done, this completes the proof of Theorem 1.18 and therefore of Theorem 1.4. \square

1.5 Further results

The next result is easily seen to follow from the methods discussed in the previous sections.

Theorem 1.19. *Let G be a nilpotent group of measure preserving transformations of a probability space (X, \mathcal{X}, μ) . Then, for every $T_1, \dots, T_l \in G$, every $f_1, \dots, f_r \in L^\infty(X)$, every set of polynomials $p_{i,j} : \mathbb{Z}^d \rightarrow \mathbb{Z}$, and every Følner sequence $\{\Phi_N\}_{N=1}^\infty$*

in \mathbb{Z}^d , the averages

$$\frac{1}{|\Phi_N|} \sum_{u \in \Phi_N} \prod_{j=1}^r \left(T_1^{p_{1,j}(u)} \dots T_l^{p_{l,j}(u)} \right) f_j \quad (1.5.1)$$

converge in $L^2(X, \mathcal{X}, \mu)$.

During the proof of Theorem 1.4 we used crucially the fact that the L^∞ norm is an algebra norm ($\|fg\|_\infty \leq \|f\|_\infty \|g\|_\infty$). While this is not true for the L^2 norm, if we are concerned with the study of a single function $f \in L^2(X)$, this issue is no longer present. Furthermore, in this case our polynomial systems will always have size 1, a fact that allows us to drop the hypothesis of nilpotency on our group G (because we no longer need the product of polynomial sequences to be polynomial). More generally, it is easy to see from these observations that our methods produce the following result, which was also conjectured by Bergelson and Leibman in [7].

Theorem 1.20. *Let G be a group of unitary operators on a Hilbert space \mathcal{H} . If $(g(n))_{n \in \mathbb{Z}}$ is a polynomial sequence in G , then*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N g(n)u$$

exists for every $u \in \mathcal{H}$.

This was established by Bergelson and Leibman [7] for nilpotent G . While our result drops this hypothesis, it should be noted that it is not presently known if there are polynomial sequences whose study does not essentially reduce to the nilpotent case.

1.6 Some examples of reductions

We now provide some concrete examples of how the process studied in Section §1.4 returns the trivial system for some polynomial systems. Given systems \mathbf{g} and \mathbf{h} we write $\mathbf{g} \sim \mathbf{h}$ to mean that both systems are equivalent and we write $\mathbf{g} \xrightarrow{m} \mathbf{h}$ to mean that \mathbf{h} is the m -reduction of \mathbf{g} . Given measure preserving transformations T, S, C of a probability space, we will use the convention of writing $T^{n^2} S^n C$ for the G -sequence g given by $g(n) = T^{n^2} S^n C$.

Before giving the examples, we note that an unpleasant feature of the process used in [53] is that in the simplest cases it is unnecessarily complicated (mainly, this happens when the group is abelian). A way to amend this is the trivial observation that the averages (1.3.1) for sequences $g_i(n) = h_i(n)C_i$, with C_i some set of transformations, equal the averages associated to the system $\mathbf{h} = (h_1, \dots, h_j)$ evaluated at the functions $C_i f_i$. Thus, we may extend the previous equivalence relation to include those pairs of systems which may be obtained from each other by adding or

removing a constant in the above manner. It is easy to check that the arguments of [53, §3] work equally well with this alternative notion of equivalence. We write $\mathbf{g} \sim^* \mathbf{h}$ if \mathbf{g} and \mathbf{h} are equivalent in this way and call this modification of the process 'cheating'. While there is no substantial gain by using this slight modification in the general case, many of the examples discussed below become much cleaner in this way. Nevertheless, we will also show in every case how the process is performed without cheating.

Example 1.21. As a trivial example, suppose our system is constant i.e. of the form (C_1, \dots, C_j) for some constant G -sequences C_1, \dots, C_j . Then its m -reduction is equivalent to

$$(1_G, C_1, \dots, C_{j-1}),$$

for every m , so in particular we get the trivial system after at most j steps. Of course, if one is allowed to cheat, one has

$$(C_1, \dots, C_j) \sim^* (1_G),$$

to begin with.

Example 1.22. Suppose we are given a linear system

$$(L_1^n C_1, \dots, L_j^n C_j),$$

consisting of commuting transformations $L_1, C_1, \dots, L_j, C_j$. Then, we see that the m -reduction of this system is given by

$$(L_1^n C_1, \dots, L_{j-1}^n C_{j-1}, L_j^{-m}, L_j^{-m} L_1^{n+m} C_1, \dots, L_j^{-m} L_{j-1}^{n+m} C_{j-1}) \quad (1.6.1) \\ \sim^* (L_1^n, \dots, L_{j-1}^n).$$

This highlights the advantage of cheating. Indeed, this would allow us to go to the trivial system in j steps, while without cheating we would require more than $2 \uparrow \uparrow j$ steps. We now see how the latter is accomplished. Our objective is to eliminate L_{j-1}^n from the reduction (1.6.1) (cf. the general strategy given in [53, §4]). In general, if we are given a system of the form

$$(L_1^n C_{1,1}, \dots, L_1^n C_{1,i_1}, \dots, L_k^n C_{k,1}, \dots, L_k^n C_{k,i_k}),$$

with the transformations generating an abelian group, its m -reduction will be equivalent to

$$(L_k^{-m}, L_1^n C_{1,1}, \dots, L_1^n C_{1,i_1}, L_1^n C_{1,1} L_k^{-m} L_1^m, \dots, \\ L_1^n C_{1,i_1} L_k^{-m} L_1^m, \dots, L_{k-1}^n C_{k-1,1}, \dots, L_{k-1}^n C_{k-1,i_{k-1}}, \\ L_{k-1}^n C_{k-1,1} L_k^{-m} L_{k-1}^m, \dots, L_{k-1}^n C_{k-1,i_{k-1}} L_k^{-m} L_{k-1}^m, L_k^n C_{k,1}, \dots, L_k^n C_{k,i_k-1}).$$

In particular, L_i^n appears twice as many times as before for every $1 \leq i < k$, while L_k^n appears one time less, therefore disappearing after i_k steps. Notice also that at each step we get twice plus one as many constant sequences as before. Applying these observations we see that the system $(L_1^n C_1, \dots, L_j^n C_j)$ reduces to one consisting only of constant sequences after at most $a(j)$ steps, with $a : \mathbb{N} \rightarrow \mathbb{N}$ the function recursively defined by $a(1) = 1$ and $a(n+1) = a(n) + 2^{a(n)}$. We may then proceed as in Example 1.21.

Example 1.23. Consider now a system of the form

$$(S_1^n C_1, \dots, S_{j-1}^n C_{j-1}, T^{n^2} S_j^n C_j),$$

for commuting T, S_i, C_i . The m -reduction is given by

$$(S_1^n C_1, \dots, S_{j-1}^n C_{j-1}, T^{-2mn-m^2} S_j^{-m}, \\ T^{-2mn-m^2} S_j^{-m} S_1^{n+m} C_1, \dots, T^{-2mn-m^2} S_j^{-m} S_{j-1}^{n+m} C_{j-1}),$$

which is a linear system and therefore reduces to the trivial system by the procedure discussed in Example 1.22.

Example 1.24. If we are given a system of size 1 consisting of a polynomial G -sequence g then it is obvious that the number of steps required to reach the trivial system is the total degree of g . Notice that this is true even when the group G is not assumed to be nilpotent.

Example 1.25. Consider $(T^{n^2}, T^{n^2} S^n)$ for commuting T and S . The m -reduction is given by

$$(T^{n^2}, T^{-2mn-m^2} S^{-m}, T^{n^2} S^{-m}) \\ \xrightarrow{l} (T^{n^2}, T^{-2mn-m^2} S^{-m}, T^{-2ln-l^2}, T^{n^2}, T^{-2ln-l^2-2m(n+l)-m^2} S^{-m}),$$

and this is equivalent to a system of the form studied in Example 1.23.

Example 1.26. Consider (T^n, S^n) with T and S generating a nilpotent group. The m_1 -reduction is given by

$$(T^n, S^{-m_1}, S^{-m_1} T^{m_1} T^n).$$

Write $C := S^{-m_1} T^{m_1}$, $C^{(1)} := S^{-m_1}$. Then the m_2 -reduction of this is given by

$$(T^n, C^{(1)}, CT^{-m_2} C^{-1}, CT^{-m_2} C^{-1} T^{m_2} T^n, CT^{-m_2} C^{-1} C^{(1)}) \\ \sim (C^{(1)}, C^{(2)}, C^{(3)}, T^n, [C^{-1}, T^{m_2}] T^n).$$

for some constant G -sequences $C^{(2)}, C^{(3)}$ which depend on m_1, m_2 . By the same reasoning we see that after reducing at m_3, \dots, m_l (and passing to equivalent systems) we get the system

$$(C^{(1)}, C^{(2)}, \dots, C^{(c(l))}, T^n, [[[[C^{-1}, T^{m_2}]^{-1}, T^{m_3}]^{-1}, \dots]^{-1}, T^{m_l}] T^n),$$

for some constant G -sequences $C^{(i)}$, $1 \leq i \leq c(l)$, with $c : \mathbb{N} \rightarrow \mathbb{N}$ the increasing function defined recursively by $c(1) = 1$ and $c(n+1) = 2c(n) + 1$. Clearly, this is equivalent to

$$(C^{(0)}, C^{(1)}, \dots, C^{(c(l))}, T^n),$$

for some l of size at most $s+1$, with s the nilpotency class of the group. Since any reduction of this last system will be a constant system of size $c(l+1)$ it follows that our original system (S^n, T^n) reduces to the trivial one in at most $s+2+c(s+1)$ steps. Of course, $s+2$ steps would have sufficed if we were allowed to cheat.

Example 1.27. Our last example is the system (T^{n^2}, S^{n^2}) for commuting T and S . We have

$$\begin{aligned} (T^{n^2}, S^{n^2}) &\xrightarrow{m_1} (T^{n^2}, S^{-2nm_1-m_1^2}, S^{-2nm_1-m_1^2}T^{n^2}T^{2nm_1+m_1^2}) \\ &\sim (S^{-2nm_1-m_1^2}, T^{n^2}, S^{-2nm_1-m_1^2}T^{n^2}T^{2nm_1+m_1^2}) \\ &\xrightarrow{m_2} (S^{-2nm_1-m_1^2}, T^{n^2}, T^{-2nm_2-m_2^2-2m_1m_2}S^{2m_1m_2}, \\ &\quad T^{-2nm_2-m_2^2-2m_1m_2}S^{-2nm_1-m_1^2}, T^{n^2}S^{2m_1m_2}T^{-2m_1m_2}) \\ &\xrightarrow{m_3} (S^{-2nm_1-m_1^2}, T^{n^2}, T^{-2nm_2-m_2^2-2m_1m_2}S^{2m_1m_2}, \\ &\quad T^{-2nm_2-m_2^2-2m_1m_2}S^{-2nm_1-m_1^2}, \\ &\quad T^{-2nm_3-m_3^2}, T^{-2nm_3-m_3^2}S^{-2nm_1-2m_1m_3-m_1^2}, \\ &\quad T^{n^2}, T^{-2nm_3-m_3^2-2nm_2-2m_2m_3-m_2^2-2m_1m_2}S^{2m_1m_2}, \\ &\quad T^{-2nm_3-m_3^2-2nm_2-2m_2m_3-m_2^2-2m_1m_2}S^{-2nm_1-2m_1m_3-m_1^2}) \end{aligned}$$

and this last system is equivalent to one of the form studied in Example 1.23.

Chapter 2

The inverse sieve problem

2.1 Introduction

In this chapter we will study a strong manifestation of algebraic rigidity, related with the distribution of sets in residue classes. The study of such 'local' distributional properties is one of the main topics of interest in analytic number theory. Examples abound, but folkloric ones include Dirichlet's theorem, which tells us that the primes are uniformly distributed along primitive residue classes, and the open problem of determining how large may the least quadratic non-residue be.

Another example worth mentioning is the Bombieri-Vinogradov theorem [9]. Write $\psi(x; a, q)$ for the number of primes congruent to $a \pmod{q}$ below x and φ for Euler's totient function, that is,

$$\varphi(n) = n \prod_{p|n} \left(1 - \frac{1}{p}\right).$$

We say that the primes have level of distribution θ if, for every $A > 0$ and $\varepsilon > 0$, we have the estimate

$$\sum_{q \leq x^{\theta - \varepsilon}} \max_{(a, q) = 1} \left| \psi(x; a, q) - \frac{x}{\varphi(q)} \right| \ll_{A, \varepsilon} \frac{x}{(\log x)^A},$$

This is saying that the primes are very well behaved mod q *on average*. The Bombieri-Vinogradov theorem states that the primes have level of distribution $1/2$. This is a remarkable result that equals what the Generalized Riemann Hypothesis can say in this situation.

It is conjectured that the primes have level of distribution 1, but this seems to be far out of reach of present technology. A spectacular result of Goldston, Pintz and Yildirim [27] shows that, if there exists some $\varepsilon > 0$ such that the primes have

level of distribution $1/2 + \varepsilon$, then there must exist some absolute constant $C = C(\varepsilon)$ and infinitely many primes p , such that $p + C$ is also prime.

The above example serves to illustrate two points. First of all, besides being interesting in itself, there are strong applications to be drawn from knowing that a set is well behaved in residue classes. This is particularly true for the application of sieve methods (see [17]). The second point, which will be the most relevant for our study, is that obvious obstructions notwithstanding, the main goal of the problems mentioned above is to show that the sets of interest are, in fact, very well distributed in residue classes.

2.1.1 Sieves and algebraic structure

From the point of view of discrete inverse theorems, the last point mentioned in the previous paragraph raises an obvious question: may it be the case that every set that is very badly distributed in residue class must be of a specific form? Since we would expect a random set to be fairly well distributed, the main question becomes whether a set occupying very few residue classes for many primes p has to have some specific structure. The remarkable observation that this might indeed be the case is due to Croot and Elsholtz [14] and Helfgott and Venkatesh [30]. Writing $[N]$ for the set of integers $\{0, \dots, N\}$ their observation can be resumed in the following principle:

Inverse Sieve Problem. *Suppose a set $S \subseteq [N]^d$ occupies very few residue classes mod p for many primes p . Then, either S is small, or it possesses some strong algebraic structure.*

An important example to keep in mind here is that of the squares. Indeed, let $S = \{1, 4, 9, 16, \dots\} \subseteq [N]$ be the set of squares. It is well known that this set occupies only $(p + 1)/2$ residue classes mod p for every prime p . That is, the set of squares is very badly distributed in residue classes. However, it is clear that this set is also a 'strongly algebraic' set. The guess is then that every large subset of $[N]$ occupying so few residue classes must be of this form, i.e. lie in the image of a polynomial (see Conjecture 2.16).

There is a good reason why such inverse sieve results are of much interest in number theory. One of the main features of sieve theory is the uniformity of its results, which is a consequence of the fact that sieves only take into account the cardinality of the classes occupied by the set. However, a clear drawback of this is that the bounds thus obtained are limited to what happens in extremal cases. By stating that such extremal sets must have a very specific structure, inverse results should allow one to retain the uniformity of the sieve while providing much stronger bounds. The reader may consult the book of Kowalski [35, §2.5] for further discussion of the potential applications of this phenomenon and [21] for applications of similar classifications in arithmetic combinatorics.

In the present chapter we give a satisfactory answer to the inverse sieve problem for every $d \geq 2$. In order to discuss our results suppose we are given a big integer set $S \subseteq [N]^d$ occupying $O(p^{d-1})$ residue classes in $(\mathbb{Z}/p\mathbb{Z})^d$ for many primes p . What does this imply about S ? By the Lang-Weil inequality, we know that this condition is satisfied by the set of integer points of a proper algebraic variety of small degree and one would expect a partial converse to also hold. That is, that any big set $S \subseteq [N]^d$ occupying only that many residue classes for every prime p should essentially be contained inside the solution set of a polynomial of low degree. When $d = 1$ this follows from Gallagher's larger sieve [26] (not to be confused with the conjecture discussed in §2.6.3). The case $d = 2$ was proven by Helfgott and Venkatesh in [30], by applying the Bombieri-Pila determinant method [10] to obtain a two-dimensional generalization of the larger sieve. Although their methods are only capable of handling the case $d \leq 2$, they conjectured that such an inverse theorem should in fact hold for every dimension d . In this chapter we introduce a different approach and use it to answer this question by giving the following best possible result.

Theorem 2.1. *Let $0 \leq k < d$ be integers and let $\varepsilon, \alpha, \eta > 0$ be positive real numbers. Then, there exists a constant C depending only on the above parameters, such that for any set $S \subseteq [N]^d$ occupying less than αp^k residue classes for every prime p at least one of the following holds:*

- (*S is small*) $|S| \ll_{d,k,\varepsilon,\alpha} N^{k-1+\varepsilon}$,
- (*S is strongly algebraic*) There exists a nonzero polynomial $f \in \mathbb{Z}[x_1, \dots, x_d]$ of degree at most C and coefficients bounded by N^C vanishing at more than $(1 - \eta)|S|$ points of S .

Theorem 2.1 is sharp. Indeed, the reader may consult Section §2.5 for examples of sets of size $|S| \gg N^{k-1}$ occupying less than p^k residue classes for every prime p but possessing no algebraic structure. On the other hand, we only need to require from S that it occupies few residue classes for sufficiently many small primes (see Theorem 2.5). More generally, we will show in Theorem 2.14 that assuming some necessary regularity conditions, every set of size $\gg N^\varepsilon$ occupying few residue classes for many primes p must satisfy condition (ii). In Section §2.6.2 we shall give an easy application of this generalization to the characterization of functions preserving some structure when reduced to prime moduli.

Taking $d = 2$ in Theorem 2.1 we recover the result of [30]. Actually, the methods of Helfgott and Venkatesh are capable of handling the case $k = 1$ of Theorem 2.1, that is, when S is assumed to occupy only $O(p)$ residue classes. However, the approach fails as soon as the set occupies more than $p \log p$ classes. The reason for this is that their method, as well as the larger sieve itself, is in essence a counting argument (see §2.3.1) and therefore needs the *number* of classes occupied by S to be

small, while the high dimensional setting requires us to take advantage of the local *density* of S being small. This type of obstacle is not specific to the problem at hand, but arises whenever one tries to extend these kind of sieves to higher dimensional settings (see [34, Remark 3] for some discussion). So while we do make use of the larger sieve, in order to establish Theorem 2.1 we need to introduce an approach that overcomes this difficulty by taking advantage of the structure of the set and which we believe to be applicable in more general situations.

2.1.2 Organization

The rest of this chapter is organized as follows. After setting up some notation, in Section §2.2 we state and discuss Proposition 2.3, which is the main ingredient of our study, and use it to deduce Theorem 2.1. Roughly speaking, this proposition says that every set satisfying hypothesis similar to those of Theorem 2.1 admits a subset of size $O(r^k)$ such that if a polynomial identity of degree r holds at this set then it must also hold at a positive proportion of the points of S . Then, in Section §2.3, we review some facts about the larger sieve and apply them to obtain a key uniformization lemma. Using this, the proof of Proposition 2.3 is carried out in Section §2.4. Finally, in §2.5 we construct several examples showing that our results are sharp, while in §2.6 we discuss further consequences of our methods as well as the remaining case ($d = 1$) of the inverse sieve problem.

2.2 A conditional proof of Theorem 1.1

2.2.1 General notation

We now fix some notation. By $O_{c_1, \dots, c_k}(X)$ we shall mean a quantity which is bounded by $C_{c_1, \dots, c_k}X$ where C_{c_1, \dots, c_k} is some finite positive constant depending on c_1, \dots, c_k . Also, we shall write $Y \ll_{c_1, \dots, c_k} X$ to mean $|Y| = O_{c_1, \dots, c_k}(X)$. However, since we will generally be concerned with the study of a set S satisfying the hypothesis of Proposition 2.3 for some parameters d, h, κ and ε as in the statement of that proposition, we will free up some notation by assuming that all implied constants in the O, \ll notation *always* depend on these parameters even though this may not be explicitly stated. So for instance $Y \ll_\eta X$ stands for $Y \ll_{\eta, d, h, \kappa, \varepsilon} X$. Throughout the rest of this chapter we will let the letter c denote a small positive constant whose exact value may vary at each occurrence.

Given a statement $\phi(x)$ with respect to an element $x \in [N]^d$ we will write $\mathbf{1}_{\phi(x)}$ for the function which equals 1 if $\phi(x)$ is true and 0 otherwise. Also, we shall write $\pi_i : \mathbb{Z}^d \rightarrow \mathbb{Z}$, $1 \leq i \leq d$, for the projection to the i th coordinate.

The letter p will always refer to a prime number. We write \mathcal{P} for the set of primes and given any magnitude Q , we denote $\mathcal{P}(Q)$ the set of primes $p \leq Q$. Since

we will usually need to consider the weight $\frac{\log p}{p}$ over \mathcal{P} , for a finite subset $P \subseteq \mathcal{P}$ we write $w(P) := \sum_{p \in P} \frac{\log p}{p}$. We shall use the estimates

$$w(\mathcal{P}(Q)) = \log Q + O(1),$$

and

$$\sum_{p \in \mathcal{P}(Q)} \log p \sim Q,$$

without explicit mention.

2.2.2 Characteristic sets

The purpose of this section is to state Proposition 2.3, which is the key ingredient of our argument, and use it to derive Theorem 2.1. What this proposition essentially says, is that for any ill-distributed set S as in the statement of Theorem 2.1, one may find a very small “characteristic” subset $A \subseteq S$ such that if a small polynomial vanishes at A then it also vanishes at a positive proportion of S . The task of proving Theorem 2.1 is thus reduced to that of finding a polynomial which vanishes at A , and this will always be possible since A is small.

Before proceeding, we need to define exactly what we mean for a polynomial to be small. Given a parameter N and some integer $d > 0$ by an r -polynomial, for a positive integer r , we shall mean any polynomial f with integer coefficients satisfying $|f(n)| < N^{3r}$ for every $n \in [N]^d$. The exponent $3r$ is chosen in order to guarantee that if N is sufficiently large in terms of r and d , then a polynomial $f \in \mathbb{Z}[x_1, \dots, x_d]$ of degree at most r , with coefficients bounded in absolute value by N^r , is an r -polynomial. This leads us to the following definition.

Definition 2.2. Let $0 < \delta \leq 1$ be a positive real number and $r > 0$ some integer. We say a subset A of a set S is (r, δ) -characteristic for S if we can find some subset $A \subseteq B \subseteq S$ of size $|B| \geq \delta|S|$ such that whenever an r -polynomial vanishes at A , then it also vanishes at B .

We can now state Proposition 2.3 which says that ill-distributed sets always admit small characteristic subsets.

Proposition 2.3. *Let $d, h \geq 1$ be arbitrary integers and $\varepsilon > 0$ some positive real number. Set $Q = N^{\frac{\varepsilon}{2d}}$ and let $P \subseteq \mathcal{P}(Q)$ satisfy $w(P) \geq \kappa \log Q$ for some $\kappa > 0$. Also, let r be an arbitrary positive integer. Suppose $S \subseteq [N]^d$ is a set of size $|S| \gg N^{d-h-1+\varepsilon}$ occupying at most αp^{d-h} residue classes mod p for every prime $p \in P$ and some $\alpha > 0$. Then, if N is sufficiently large, there exists a set $A \subseteq S$ of size $|A| = O(r^{d-h})$ which is (r, δ) -characteristic for S , for some $\delta > 0$ which depends on $d, h, \kappa, \varepsilon$ but is independent of S, N or r .*

Remarks. The exact value of Q in the above statement is irrelevant and may be replaced by any small power of N . The reason why we have made the change of variables $h := d - k$ with respect to Theorem 2.1 is that in the arguments to follow we shall always set the quantity h to be fixed and induct on d . We believe it is simpler to introduce this change of notation at an early stage.

To see why such a result might be expected consider some polynomial f vanishing at an integer point x . Since polynomials descend to congruence classes, this means that for any other integer y satisfying $y \equiv x \pmod{p}$ for a prime p , we will have $p|f(y)$. Thus, if we are given a set S which occupies very few residue classes, one may then hope to find a small subset A such that given some $y \in S$ there are a lot of primes p for which $y \equiv x \pmod{p}$ for some $x \in A$. It would then follow that if a polynomial vanishes at A then there would be many primes p dividing $f(y)$. If furthermore f is small, then this can only hold if $f(y) = 0$. Notice that this is similar to the general idea of the larger sieve, where one uses the fact that S occupies few residue classes mod p to conclude the existence of too many pairs of elements of S occupying the same class, and contrasts this with the fact that no fixed pair of distinct bounded integers can occupy the same residue class for many primes (see §2.3.1).

On the other hand, the size hypothesis on S is necessary. For instance, one may construct small (logarithmic size) sets $S \subseteq [N]$ as in [30, §4.3] which occupy few residue classes for large moduli just because they are small, but which however have at most one element in each residue class, making the above argument unviable in this situation. Furthermore, it is clear that a similar pathology occurs in higher dimensions, by considering for instance the product set $S \times [N]$. For the general construction of this type of sets and to see that in fact one cannot take $\varepsilon = 0$ in Proposition 2.3 the reader is referred to §2.5.

In order to deduce Theorem 2.1 from Proposition 2.3 we will need to find a polynomial which vanishes at a specific set of points. This will be accomplished in a standard way by means of Siegel's lemma.

Lemma 2.4 (Siegel). *Suppose we are given a system of m linear equations*

$$\sum_{j=1}^n a_{ij}\beta_j = 0 \quad \forall 1 \leq i \leq m,$$

in n unknowns $(\beta_1, \dots, \beta_n)$, $n > m$, where the coefficients (a_{ij}) are integers not all equal to 0 and bounded in magnitude by some constant C . Then, the above system has a non trivial integer solution $(\beta_1, \dots, \beta_n)$ with

$$|\beta_j| \leq 1 + (Cn)^{m/(n-m)},$$

for all $1 \leq j \leq n$.

Proof of Theorem 2.1 assuming Proposition 2.3. Let the hypothesis be as in the statement of Theorem 2.1 and write $h := d - k$. Assume condition (i) fails, so that $|S| \gg N^{d-h-1+\varepsilon}$. We claim that for any given integer r there exists a set $A \subseteq S$ of size $|A| = O_\eta(r^{d-h})$ which is $(r, 1 - \eta)$ -characteristic for S , provided N is sufficiently large. To see this we begin by noticing that Proposition 2.3 implies the existence of some $\delta \gg 1$ such that for every subset $S' \subseteq S$ with $|S'| \geq \eta|S|$ there exists a set $A' \subseteq S'$ of size $|A'| = O_\eta(r^{d-h})$ which is (r, δ) -characteristic for S' . From now on we fix this value of δ . Let A_0 be such a characteristic subset for S and let B_0 consist of those elements of S which vanish at every r -polynomial that vanishes at A_0 , so in particular $|B_0| \geq \delta|S|$. If $\delta \geq 1 - \eta$ we are done, otherwise we have that $S_1 := S \setminus B_0$ satisfies $|S_1| \geq \eta|S|$ and therefore contains a characteristic subset $A_1 \subseteq S_1$ as above. If we now let B_1 denote those points of S_1 vanishing at every r -polynomial that vanishes at A_1 we see that either we get the claim with $A = A_0 \cup A_1$ or the set $S_2 := S_1 \setminus B_1$ satisfies

$$\eta|S| \leq |S_2| \leq (1 - \delta)^2|S|.$$

After iterating this process j times we see that if the set $A = \bigcup_{i=0}^{j-1} A_i$ is not $(r, 1 - \eta)$ -characteristic for S then we can find some $S_j \subseteq S$ with

$$\eta|S| \leq S_j \leq (1 - \delta)^j|S|.$$

Since this last possibility cannot hold for some large $j = O_\eta(1)$, the claim follows.

Now it only remains to find some r -polynomial f which vanishes at A and which is of the form given in Theorem 2.1. Notice that this is plausible since the size of A is $\ll_\eta r^{d-h}$ while an r -polynomial has $\sim r^d$ degrees of freedom. We now make this rigorous by means of Siegel's lemma. Thus, we may assume $d|r$ and consider the system of $|A|$ linear equations in $\binom{r}{d} + 1$ unknowns given by

$$\sum_{\mathbf{i}=\{i_1, \dots, i_d\} \leq r/d} \beta_{\mathbf{i}} a^{\mathbf{i}} = 0 \quad \forall a \in A, \quad (2.2.1)$$

where $\mathbf{i} \leq l$ stands for $i_j \leq l$ for all $1 \leq j \leq d$ and where we use the multi-index notation $a^{\mathbf{i}} = a_1^{i_1} \dots a_d^{i_d}$ for $a = (a_1, \dots, a_d)$. Notice that $|a^{\mathbf{i}}| \leq N^r$ for every \mathbf{i} and that a solution $(\beta_{\mathbf{i}})$ of (2.2.1) corresponds to the coefficients of a polynomial vanishing at A . If we now choose $r = O_\eta(1)$ large enough so that $\binom{r}{d} + 1 > 3|A|$ it follows by Siegel's lemma that there exists an integer solution $(\beta_{\mathbf{i}})$ to (2.2.1) with

$$|\beta_{\mathbf{i}}| \ll_r N^{r/2} \leq N^r,$$

provided N is sufficiently large. We thus see that the polynomial

$$f := \sum_{\mathbf{i} \leq r/d} \beta_{\mathbf{i}} x^{\mathbf{i}},$$

is of the desired form (assuming again that N is sufficiently large) and, taking $C = r$, this concludes the proof of Theorem 2.1. \square

Notice that we have actually proved the following slight strengthening of Theorem 2.1 in which the set S is only required to be badly distributed in a dense subset of the primes.

Theorem 2.5. *Let $0 \leq k < d$ be integers and let $\varepsilon, \eta > 0$ be positive real numbers. Set $Q = N^{\frac{\varepsilon}{2d}}$ and let $P \subseteq \mathcal{P}(Q)$ satisfy $w(P) \geq \kappa \log Q$ for some $\kappa > 0$. Suppose $S \subseteq [N]^d$ is a set of size $|S| \gg N^{k-1+\varepsilon}$ occupying at most αp^k residue classes mod p for every prime $p \in P$ and some $\alpha > 0$. Then there exists a nonzero polynomial $f \in \mathbb{Z}[x_1, \dots, x_d]$ of degree $O_\eta(1)$ and coefficients bounded by $N^{O_\eta(1)}$ which vanishes at more than $(1 - \eta)|S|$ points of S .*

Remark. Since we have already mentioned that the exact value of Q in Proposition 2.3 is irrelevant, it follows that Theorem 2.5 also holds with Q any small power of N .

2.3 Applying the larger sieve in high dimensions

2.3.1 A review of the larger sieve

We will now quickly review some facts about Gallagher's larger sieve and use them to prove two easy lemmas which we shall need later. For further discussion of the larger sieve and its consequences the reader may consult [12, Section 2.2] and of course Gallagher's original paper [26].

Before proceeding we need to state some further notation that will be used in this and the next sections. When studying a set $S \subseteq [N]^d$ we will denote by $[S]_p$ the set of residue classes mod p occupied by S . Given such a set S , we shall be largely concerned with how many elements of S belong to a given residue class, so it is important for us to have a specific notation for this subset. Thus, given a residue class $\mathbf{a} = (a_1, \dots, a_d) \pmod{p}$ we write $S(\mathbf{a}; p)$ to refer to those elements of S which are congruent to $\mathbf{a} \pmod{p}$. Moreover, we shall sometimes consider some $a \in \mathbb{Z}/p\mathbb{Z}$ and write $S(a; p)$ for those elements of S having their first coordinate congruent to $a \pmod{p}$. Since we will always use the bold font \mathbf{a} to denote a vector residue class and since where this class lives shall be clear from the context altogether, we believe the similarity of both notations will not cause any confusion. Finally, if p is fixed, we may simply write $S(\mathbf{a})$ and $S(a)$ for the above sets.

Fix now a set S and consider some parameter Q . The main idea of the larger sieve is to count in two different ways the number of distinct pairs $x, y \in S$ and primes $p \leq Q$ such that $x \equiv y \pmod{p}$. Given two such integers $x, y \in [N]$ it is clear that those primes for which they are congruent are exactly those dividing $|x - y| \leq N$ and therefore

$$\sum_{p \leq Q} \sum_{\substack{x, y \in S \\ x \neq y}} \mathbf{1}_{x \equiv y \pmod{p}} \log p \leq |S|^2 \log N. \quad (2.3.1)$$

On the other hand, we have that the left hand side of (2.3.1) equals

$$\sum_{p \leq Q} \sum_{a \pmod{p}} |S(a; p)|^2 \log p - |S| \sum_{p \leq Q} \log p. \quad (2.3.2)$$

Notice that the above argument also works if $S \subseteq [N]^d$ since if p is a prime for which $x \equiv y \pmod{p}$ then p must divide $|\pi_1(x) - \pi_1(y)|$ which is bounded by N .

As an example, we have the following result due to Gallagher [26]. Suppose we are given a set $S \subseteq [N]$ occupying at most αp residue classes, then the Cauchy-Schwarz inequality implies

$$\sum_{a \pmod{p}} |S(a; p)|^2 \geq \frac{1}{\alpha p} |S|^2.$$

Combining this with (2.3.1) and (2.3.2) we obtain

$$\frac{1}{\alpha} \log Q + O\left(\frac{|Q|}{|S|}\right) \leq \log N + O(1).$$

Taking $Q = |S|$ we conclude that $|S| \ll_{\alpha} N^{\alpha}$.

For the purposes of our work, we need to apply Gallagher's sieve in a slightly more general context. Precisely, we will use the following lemma which is also classical.

Lemma 2.6. *Let $X \subseteq [N]$ be some set of integers and set $Q = N^{\gamma}$ for some $\gamma > 0$. Let $c_1, c_2 > 0$ be positive real numbers. Suppose there is a set of primes $P \subseteq \mathcal{P}(Q)$ with $w(P) \geq c_1 \log Q$ such that for every $p \in P$ there are at least $c_2 |X|$ elements of X lying in at most αp residue classes for some $\alpha > 0$ independent of p . Then, if α is sufficiently small in terms of c_1, c_2 and γ , it must be $|X| < Q$.*

Proof. Again, we count the number of pairs $x, y \in X$ and $p \in P$ with $x \equiv y \pmod{p}$. On one hand, we have as before that

$$\sum_{p \in P} \sum_{\substack{x, y \in X \\ x \neq y}} \mathbf{1}_{x \equiv y \pmod{p}} \log p \leq |X|^2 \log N. \quad (2.3.3)$$

On the other hand, using the Cauchy-Schwarz inequality we see that our hypothesis on X implies

$$\sum_{a \pmod{p}} |X(a; p)|^2 \geq \frac{1}{\alpha p} (c_2 |X|)^2,$$

from where it follows that the left hand side of (2.3.3) is at least

$$\frac{c_1 c_2^2}{\alpha} |X|^2 \log Q + O(|Q| |X|).$$

It is then clear that if α is sufficiently small, then the only way for (2.3.3) to hold is to have $|X| < Q$. \square

Finally, we prove the following easy consequence of the larger sieve which already handles the case $d = h$ of Proposition 2.3.

Lemma 2.7. *Let $Q = N^\gamma$ for some $\gamma > 0$ and let $P \subseteq \mathcal{P}(Q)$ be some set of primes with $w(P) \geq c_1 \log Q$ for some $c_1 > 0$. Let $S \subseteq [N]^d$ occupy less than c_2 residue classes mod p for every prime $p \in P$ and some constant c_2 . Then $|S| = O_{c_1, c_2, \gamma}(1)$.*

Proof. Gallagher's sieve implies in this case

$$\log N \geq \left(\frac{1}{c_2} - \frac{1}{|S|} \right) \sum_{p \in P} \log p \gg \left(\frac{1}{c_2} - \frac{1}{|S|} \right) N^{\gamma c_1},$$

and clearly, for sufficiently large N , this can only hold if $|S| \leq c_2$. \square

2.3.2 Genericity

Our strategy to prove Proposition 2.3 will be to partition S into many lower dimensional subsets and apply induction. However, the main obstacle we encounter in doing so (and which is not merely a technical issue, as can be seen from the examples in §2.5) is the possibility that the resulting subsets are rather independent from each other, in the sense that they do not share many residue classes. If this happens, then the fact that a small polynomial vanishes at one of this subsets will not give us much information about the behavior of this polynomial in the other subsets. However, in order for this to happen it would be necessary for these subsets to occupy very few residue classes and this would imply the existence of too many elements in each subset occupying the same residue class. While with our hypothesis one cannot guarantee that this never happens, the goal of this section is to show that this indeed does not happen on average, which will be sufficient for our arguments.

We begin with the following definition.

Definition 2.8 (Genericity). Given a real number $B > 0$ and some integer $l > 0$ we say that a set $S \subseteq [N]^d$ is (B, l) -generic mod p if

$$\frac{|S(\mathbf{a}; p)|}{|S|} < \frac{B}{p^l},$$

for every residue class $\mathbf{a} \pmod{p}$.

Given a set of primes $P \subseteq \mathcal{P}(Q)$ we shall write $P' \leftrightarrow P$ to mean a subset $P' \subseteq P$ with $w(P') \gg w(P)$. Recall that by our conventions in §2.2.1 the implied constants depend on the parameters $d, h, \varepsilon, \kappa$ of Proposition 2.3. The rest of this section is devoted to the proof of the following lemma.

Lemma 2.9. *Let $d, h \geq 1$ be arbitrary integers and let $\varepsilon > 0$ be some positive real number. Set $Q = N^{\frac{\varepsilon}{2d}}$ and let $P \subseteq \mathcal{P}(Q)$ satisfy $w(P) \geq \kappa \log Q$ for some $\kappa > 0$. Suppose $S \subseteq [N]^d$ is a set of size $|S| \gg N^{d-h-1+\varepsilon}$ occupying at most αp^{d-h} residue classes mod p for every prime $p \in P$ and some $\alpha > 0$. Then there exists $B = O(1)$ and a set of primes $P' \hookrightarrow P$ such that for each $p \in P'$ there is some subset $\mathcal{G}_p(S) \subseteq S$, $|\mathcal{G}_p(S)| \gg |S|$, which is $(B, d-h)$ -generic mod p .*

Remarks. Here again the exact value of Q is not important as long as it is a small power of N . Also, as in the previous statements, all the hypothesis are necessary because of the examples in §2.5.

Proof. From now on fix an integer $h \geq 1$. If $d \leq h$ the result is trivial, so we may proceed by induction on d . Thus, let $d \geq h+1$ be some integer and assume the result holds for every smaller dimension.

Take S and P as in the statement and recall that $\pi_i(S)$ is the projection of S to the i th coordinate. We claim that for some $1 \leq i \leq d$ there exists a set $S' \subseteq S$ with $|S'| \geq |S|/2^d$ such that every $A \subseteq S'$ with $|A| \geq |S'|/2$ satisfies $|\pi_i(A)| \geq Q$. Indeed, if the claim fails with $S' = S$ and $i = 1$ we may find some subset $S_1 \subseteq S$ with $|S_1| \geq |S|/2$ and $|\pi_1(S_1)| < Q$. Then, if the claim fails again with $S' = S_1$ and $i = 2$, we get some $S_2 \subseteq S_1$ with

$$|S_2| \geq |S_1|/2 \geq |S|/4,$$

and

$$|\pi_1(S_2)|, |\pi_2(S_2)| < Q.$$

Iterating this d times either we get the claim or end up with a set $S_d \subseteq S$ satisfying

$$|S| \leq 2^d |S_d| \leq 2^d |\pi_1(S_d)| \dots |\pi_d(S_d)| < 2^d Q^d.$$

By our choice of Q this is clearly absurd for sufficiently large N and therefore the claim follows.

Since it suffices to prove the lemma for such a subset S' we may assume without loss of generality that $S' = S$ and permuting the coordinates if necessary we may also assume $i = 1$. Hence, we have that

$$|\pi_1(A)| \geq Q \text{ for every } A \subseteq S \text{ with } |A| \geq |S|/2. \quad (2.3.4)$$

We wish to construct a dense subset of S which is in an adequate position to apply the induction hypothesis. Since we will be working with the first coordinate, given some $a \in \mathbb{Z}/p\mathbb{Z}$, we will write $S(a; p)$ to refer to those elements of S having their first coordinate congruent to $a \pmod{p}$. Let B_1 be some large constant to be specified later. Since $|[S]_p| \leq \alpha p^{d-h}$, it is clear that there can be at most $\alpha p/B_1$

residue classes $a \in [\pi_1(S)]_p \subseteq \mathbb{Z}/p\mathbb{Z}$ for which $||S(a; p)||_p \geq B_1 p^{d-h-1}$. We denote by $\mathcal{E}_1(p)$ this exceptional set. Also, we write

$$\mathcal{E}_2(p) := \left\{ a \in [\pi_1(S)]_p : |S(a; p)| \geq \frac{B_1}{\alpha p} |S| \right\}.$$

From the obvious fact that

$$\sum_{a \in \mathbb{Z}/p\mathbb{Z}} |S(a; p)| = |S|,$$

it follows that $|\mathcal{E}_2(p)| \leq \alpha p/B_1$ and therefore $|\mathcal{E}(p)| \leq 2\alpha p/B_1$, where $\mathcal{E}(p) := \mathcal{E}_1(p) \cup \mathcal{E}_2(p)$. By means of the larger sieve we may now deduce that not too many integers in $[N]$ can lie in $\mathcal{E}(p)$ for many $p \in P$. Indeed, consider the set X which consists of all elements $x \in [N]$ for which

$$\sum_{p \in P} \mathbf{1}_{x(\bmod p) \in \mathcal{E}(p)} \frac{\log p}{p} \geq \frac{1}{2} w(P).$$

By the pigeonhole principle, one may then find a set of primes $P_1 \subseteq P$ with $w(P_1) \geq \frac{1}{4} w(P)$ and such that

$$\left| \bigcup_{a \in \mathcal{E}(p)} X(a; p) \right| \geq \frac{1}{4} |X|,$$

for every $p \in P_1$. It then follows from Lemma 2.6 that upon choosing B_1 sufficiently large, we can ensure that $|X| < Q$.

By (2.3.4), we deduce that

$$|S \setminus \pi_1^{-1}(X)| \geq \frac{1}{2} |S|.$$

We may therefore find a subset $S' \subseteq S$ with $|S'| \geq \frac{1}{4} |S|$ which does not intersect $\pi_1^{-1}(X)$ and such that $S'_x := \pi_1^{-1}(x) \cap S'$ satisfies $|S'_x| \gg N^{d-h-2+\varepsilon}$ for every $x \in \pi_1(S')$. Every such x lies outside of X and therefore has associated a set of primes $P_x \hookrightarrow P$ for which $x(\bmod p) \notin \mathcal{E}(p)$. Since $\mathcal{E}_1(p) \subseteq \mathcal{E}(p)$, we may apply the induction hypothesis to S'_x for every x to see that there exists sets of primes $P'_x \hookrightarrow P_x$ and constants $c, B_2 > 0$ independent of x , such that for each $p \in P'_x$ there is a $(B_2, d-h-1)$ -generic mod p set $\mathcal{G}_p(S'_x) \subseteq S'_x$ containing at least $c|S'_x|$ elements.

Since the sets P'_x constructed above satisfy $P'_x \hookrightarrow P$, with the implied constant independent of x , we may apply again the pigeonhole principle to locate some set of primes $P' \hookrightarrow P$ and some constant $c > 0$, such that for each $p \in P'$ there are at least $c|S'|$ elements $s \in S'$ for which $p \in P'_{\pi_1(s)}$. It thus follows that if for a prime $p \in P'$ we consider the set

$$\mathcal{G}_p(S) := \bigcup_{x: p \in P'_x} \mathcal{G}_p(S'_x),$$

then

$$|\mathcal{G}_p(S)| \gg |S'| \gg |S|,$$

and $\mathcal{G}_p(S) \cap \pi_1^{-1}(x) = \mathcal{G}_p(S'_x)$ is a $(B_2, d - h - 1)$ -generic set for every $x \in \pi_1(\mathcal{G}_p(S))$. Also, we see that there are at most

$$\frac{B_1}{\alpha p} |S| \ll \frac{B_1}{\alpha p} |\mathcal{G}_p(S)|,$$

elements of $\mathcal{G}_p(S)$ having the same first coordinate mod p since by construction it does not lie in $\mathcal{E}_2(p)$. It thus follows that $\mathcal{G}_p(S)$ is a B -generic set for some large B depending on B_1 and B_2 but independent of p and this concludes the proof of Lemma 2.9. \square

2.4 The proof of Proposition 2.2

In this section we give a proof of Proposition 2.3. As we did in the proof of Lemma 2.9 we will fix an integer h and induct on d . Since for $d \leq h$ the result is either trivial or follows from Lemma 2.7 we may assume $d \geq h + 1$ and that the result holds for all smaller dimensions.

Before we proceed, we give a brief discussion of the strategy of the proof. By our size hypothesis on S , we know that there must exist some coordinate i such that both the projection $\pi_i(S)$ and the corresponding sections $\pi_i^{-1}(x) \cap S$ are big (at least on average), and therefore generically distributed in the residue classes they occupy (by Lemma 2.9). Combining these two facts, one can deduce the existence of $m \gg r$ sections $\pi_i^{-1}(x) \cap S$ of S such that the probability of some $s \in S$ of being congruent mod p to some element of these sections is roughly m/p for many p (see Lemma 2.10). This in turn implies that if an r -polynomial f vanishes at these sections, $f(s)$ is expected to be divisible by many primes, which by the boundedness of f would imply that $f(s) = 0$. Thus, it only remains to find a set that is characteristic for these m sections, but by the induction hypothesis each section admits a characteristic subset and the result then follows by taking the union of these.

2.4.1 Summary of notation

To help the reader, we provide below a summary of previously introduced notation that will be needed during the proof.

- $|A|$ - the cardinality of a set A ,
- $w(P) := \sum_{p \in P} \frac{\log p}{p}$.
- $P' \hookrightarrow P$ - a subset of primes $P' \subseteq P$ with $w(P') \gg w(P)$,

- $\pi_i(A)$ - the projection of A to the i th coordinate,
- $A_x := \pi_1^{-1}(x) \cap A$ for a set $A \subseteq [N]^d$ (but P_x will have a different meaning for P a set of primes),
- $[S]_p$ - the set of residue classes occupied by $S \bmod p$,
- $S(\mathbf{a}; p)$ - those elements of S congruent to $\mathbf{a} \bmod p$,
- $S(a; p)$ - those elements of S with first coordinate congruent to $a \bmod p$,

2.4.2 Building good sections

We now turn to the details. We are thus given a set S and some positive integer r . Our first step will be to find generic sets inside the sections of S for many primes p . Proceeding as in the beginning of the proof of Lemma 2.9 we may assume that

$$|\pi_1(A)| \geq Q \text{ for every } A \subseteq S \text{ with } |A| \geq |S|/2. \quad (2.4.1)$$

This allows us, at the cost of passing to a subset of half density if necessary, to get the bound

$$|S_x| \leq 2|S|/Q \text{ for every } x \in [N], \quad (2.4.2)$$

where $S_x := \pi_1^{-1}(x) \cap S$. Finally we may also assume, again by passing to a subset of half density if necessary, that $|S_x| \gg N^{d-h-2+\varepsilon}$ for every $x \in \pi_1(S)$.

Let B be some large constant. For every prime p we denote by $\mathcal{E}(p)$ the set of residue classes $a \in \mathbb{Z}/p\mathbb{Z}$ for which $|[S(a; p)]_p| \geq Bp^{d-h-1}$ (recall that $S(a; p)$ stands for those elements of S having their first coordinate congruent to $a \pmod{p}$ and thus $[S(a; p)]_p$ consists of those residue classes in $[S]_p \subseteq (\mathbb{Z}/p\mathbb{Z})^d$ having a as a first coordinate). Since $|\mathcal{E}(p)| \leq \alpha p/B$, applying Lemma 2.6 as in the proof Lemma 2.9, we conclude by (2.4.1) that if B is chosen sufficiently large, we may find some $S' \subseteq S$, $|S'| \gg |S|$, such that for each $x \in \pi_1(S')$ we have $P_x \hookrightarrow P$, with the implied constant independent of x , and where

$$P_x := \{p \in P : x \pmod{p} \notin \mathcal{E}(p)\}.$$

This places us in a position in which we can apply the induction hypothesis to each section S'_x of S' to find some $\delta_0 \gg 1$ independent of x such that each S'_x admits a (r, δ_0) -characteristic subset of size $O(r^{d-h-1})$. In particular, we see that at the cost of passing to a subset of S' of density δ_0 if necessary, we may assume that inside each S'_x we can find a set of size $O(r^{d-h-1})$ which is $(r, 1)$ -characteristic for the whole section. Notice that since we are refining the sections, we still get a bound of the form $|S'_x| \gg N^{d-h-2+\varepsilon}$ for every $x \in \pi_1(S')$. Thus, we may also apply Lemma 2.9 to every such S'_x obtaining sets of primes $P'_x \hookrightarrow P_x$ such that for every $p \in P'_x$ we can find a $(C, d-h-1)$ -generic subset $\mathcal{G}_p(S'_x) \subseteq S'_x$, $|\mathcal{G}_p(S'_x)| \gg |S'_x|$, where C and the

implied constants are independent of p and x . In particular, we may find some set of primes $P' \hookrightarrow P$ such that for each $p \in P'$ the set

$$\mathcal{G}_p(S) := \bigcup_{x:p \in P'_x} \mathcal{G}_p(S_x),$$

satisfies

$$|\mathcal{G}_p(S)| \gg |S'| \gg |S|,$$

and each nonempty section $(\mathcal{G}_p(S))_x$ of $\mathcal{G}_p(S)$ is a $(C, d - h - 1)$ -generic set.

From now on we write $\mathcal{G}_p := \mathcal{G}_p(S)$. The next lemma is crucial as it allows us to find sections of S containing the residue class of many elements of S for many primes p .

Lemma 2.10. *There exists a set $\mathcal{B} \subseteq S'$, $|\mathcal{B}| \gg |S|$, such that for every non empty section \mathcal{B}_x of \mathcal{B} there is a set of primes*

$$P_x \hookrightarrow P' \hookrightarrow P,$$

with

$$\left| \left\{ s \in S' : [s]_p \in [\mathcal{B}_x]_p \right\} \right| \geq \frac{c|S|}{p},$$

for every $p \in P_x$, where $c > 0$ does not depend on x or p .

Proof. We begin by fixing a prime $p \in P'$ and considering some residue class $a \in [\pi_1(\mathcal{G}_p)]_p$. Since p is fixed we will simply write $\mathcal{G}_p(a)$ to denote those elements of \mathcal{G}_p with first coordinate congruent to $a \pmod{p}$. Also, given a class $\mathbf{b} \in (\mathbb{Z}/p\mathbb{Z})^d$ we write $\mathcal{G}_p(\mathbf{b})$ for those elements of \mathcal{G}_p congruent to $\mathbf{b} \pmod{p}$. By the pigeonhole principle and the fact that by construction of P' it is $|\mathcal{G}_p(a)| \leq Bp^{d-h-1}$ it follows that we may find some $\mathbf{b}_1 \in [\mathcal{G}_p(a)]_p \subseteq (\mathbb{Z}/p\mathbb{Z})^d$ with

$$|\mathcal{G}_p(\mathbf{b}_1)| \geq |\mathcal{G}_p(a)| / (Bp^{d-h-1}).$$

Consider now the set $\mathcal{B}_1 \subseteq \mathcal{G}_p(a)$ defined by

$$\mathcal{B}_1 := \bigcup_{s:[s]_p=\mathbf{b}_1} (\mathcal{G}_p)_{\pi_1(s)}, \quad (2.4.3)$$

that is, \mathcal{B}_1 is the union of those sections $(\mathcal{G}_p)_x$ in \mathcal{G}_p containing a representative of \mathbf{b}_1 .

Since each $(\mathcal{G}_p)_x$ is a $(C, d - h - 1)$ -generic set, we have that

$$|(\mathcal{G}_p)_x| \geq \frac{p^{d-h-1}}{C} |(\mathcal{G}_p)_x(\mathbf{b}_1)|$$

and therefore

$$|\mathcal{B}_1| \geq \frac{p^{d-h-1}}{C} |\mathcal{G}_p(\mathbf{b}_1)| \geq \frac{1}{BC} |\mathcal{G}_p(a)|. \quad (2.4.4)$$

Notice now that since $|\mathcal{G}_p(a)| \geq |\mathcal{B}_1|$ and $|\mathcal{G}_p(a)_p| \leq Bp^{d-h-1}$, by the first inequality of (2.4.4) and the pigeonhole principle we may find another residue class $\mathbf{b}_2 \in [\mathcal{G}_p(a)]_p$ with

$$\begin{aligned} |\mathcal{G}_p(\mathbf{b}_2)| &\geq \frac{1}{Bp^{d-h-1}} |\mathcal{G}_p(a) \setminus \mathcal{G}_p(\mathbf{b}_1)| \\ &\geq \frac{1}{Bp^{d-h-1}} \left(1 - \frac{C}{p^{d-h-1}}\right) |\mathcal{G}_p(a)|, \end{aligned}$$

which is at least $|\mathcal{G}_p(a)|/(2Bp^{d-h-1})$ if $p^{d-h-1} > 2C$. In such a case, if we now define \mathcal{B}_2 as in (2.4.3), but this time with respect to \mathbf{b}_2 , the same reasoning that gives (2.4.4) implies

$$|\mathcal{B}_2| \geq \frac{1}{2BC} |\mathcal{G}_p(a)|.$$

Iterating this process we end up with a sequence $\mathbf{b} = \{\mathbf{b}_1, \dots, \mathbf{b}_q\}$ of residue classes, $q = \lceil \frac{p^{d-h-1}}{2C} \rceil$, satisfying

$$\begin{aligned} |\mathcal{G}_p(\mathbf{b}_j)| &\geq \frac{1}{Bp^{d-h-1}} \left| \mathcal{G}_p(a) \setminus \bigcup_{i=1}^{j-1} \mathcal{G}_p(\mathbf{b}_i) \right| \\ &\geq \frac{1}{Bp^{d-h-1}} \left(1 - \frac{(j-1)C}{p^{d-h-1}}\right) |\mathcal{G}_p(a)| \\ &\geq \frac{|\mathcal{G}_p(a)|}{2Bp^{d-h-1}}, \end{aligned}$$

and $|\mathcal{B}_j| \geq \frac{1}{2BC} |\mathcal{G}_p(a)|$. In particular, we have that

$$\sum_{j=1}^q |\mathcal{B}_j| \geq \frac{q}{2BC} |\mathcal{G}_p(a)|. \quad (2.4.5)$$

Now, we consider the set

$$\mathcal{B}[a] := \left\{ s \in \mathcal{G}_p(a) : \sum_{j=1}^q \mathbf{1}_{s \in \mathcal{B}_j} \geq \frac{q}{4BC} \right\}.$$

Notice that $\mathcal{B}[a]_x := \mathcal{B}[a] \cap \pi_1^{-1}(x)$ equals $(\mathcal{G}_p)_x$ whenever this intersection is not empty. Also, (2.4.5) implies

$$|\mathcal{B}[a]| \geq \frac{1}{4BC} |\mathcal{G}_p(a)|. \quad (2.4.6)$$

We see that $\mathcal{B}[a]$ is very close to what we want, since if we take any nonempty section $\mathcal{B}[a]_x$ of this set, then there are at least $|\mathcal{G}_p(a)|/(4BC)^2$ elements $s \in \mathcal{G}(a)$ such that $s \equiv y \pmod{p}$ for some $y \in \mathcal{B}[a]_x$.

We now let $\mathcal{R} \subseteq [\pi_1(S)]_p$ consist of those residue classes $a \in \mathbb{Z}/p\mathbb{Z}$ with $|\mathcal{G}_p(a)| \geq \frac{1}{2p}|\mathcal{G}_p|$ and write

$$\mathcal{B}[p] := \{s \in S' : S'_{\pi_1(s)} \cap \mathcal{B}[a] \neq \emptyset \text{ for some } a \in \mathcal{R}\}.$$

In other words, $\mathcal{B}[p]$ consists of those sections of S' intersecting $\bigcup_{a \in \mathcal{R}} \mathcal{B}[a]$. In particular, since $\mathcal{B}[p]$ contains the disjoint union $\bigcup_{a \in \mathcal{R}} \mathcal{B}[a]$, we see from (2.4.6) and the definition of \mathcal{R} that

$$|\mathcal{B}[p]| \geq \frac{1}{8BC}|\mathcal{G}_p| \geq c|S|,$$

for some constant c independent of p .

Recall now that $w(P') \geq c \log Q$. For an element $s \in S'$ write P'_s for the set of primes $p \in P'$ for which $s \in \mathcal{B}[p]$. It follows from the above paragraph that for an appropriate choice of c the set

$$\mathcal{B} := \{s \in S' : w(P'_s) \geq c \log Q\}, \quad (2.4.7)$$

satisfies $|\mathcal{B}| \geq c|S|$. It is easy to check that \mathcal{B} is of the desired form. \square

2.4.3 Construction of the characteristic set

To conclude the proof of Proposition 2.3 we will show that if an r -polynomial vanishes at the sections \mathcal{B}_x for $\gg_r 1$ distinct values of x , then it must also vanish at a positive proportion of S . To this end, we choose m distinct sections of S' having nontrivial intersection with \mathcal{B} , where $m = O_r(1)$ is to be specified later. Notice that by (2.4.2) and Lemma 2.10 this is always possible provided N is sufficiently large. Call

$$\mathcal{L} := S'_{x_1} \cup \dots \cup S'_{x_m},$$

the union of these sections. Let $P_{\mathcal{L}}$ consist of those primes p for which there exists a pair of sections $S'_{x_i} \neq S'_{x_j}$ in \mathcal{L} with $[S'_{x_i}]_p \cap [S'_{x_j}]_p \neq \emptyset$. Given such a pair of sections the fact that $[S'_{x_i}]_p \cap [S'_{x_j}]_p \neq \emptyset$ implies in particular that $x_i \equiv x_j \pmod{p}$. Since $x_i \neq x_j$ this implies that the sum of $\log p$ over such primes is bounded by $\log N$. Thus, we see that

$$\sum_{p \in P_{\mathcal{L}}} \log p \leq \binom{m}{2} \log N, \quad (2.4.8)$$

and this implies that $w(P_{\mathcal{L}}) \ll_r \log \log N$.

We now consider on S' the function

$$\psi_{\mathcal{L}}(s) := \sum_{p \leq Q} \mathbf{1}_{\exists x \in \mathcal{L} : s \equiv x \pmod{p}} \log p.$$

Thus, $\psi_{\mathcal{L}}(s)$ measures the extent to which the residue classes occupied by s have a representative in \mathcal{L} . If we write P_i to denote the set of primes in Lemma 2.10 corresponding to the section $S'_{x_i} \cap \mathcal{B}$ of \mathcal{B} , it follows from this lemma and (2.4.8) that

$$\begin{aligned} \sum_{s \in S'} \psi_{\mathcal{L}}(s) &\geq \sum_{i=1}^m \sum_{p \in P_i \setminus P_{\mathcal{L}}} \sum_{s \in S'} \mathbf{1}_{\exists x \in S'_{x_i} : s \equiv x \pmod{p}} \log p \\ &\geq \sum_{i=1}^m \sum_{p \in P_i \setminus P_{\mathcal{L}}} \frac{c|S|}{p} \log p \\ &\geq m|S| (c \log Q + O_r(\log \log N)) \\ &\geq c_0 m|S| \log Q, \end{aligned}$$

for some $c_0 > 0$ and sufficiently large N .

Set $\delta = \frac{\varepsilon c_0}{4d}$ and suppose there are at most $\delta|S|$ elements $s \in S'$ with $\psi_{\mathcal{L}}(s) \geq 3r \log N$. Since $\psi_{\mathcal{L}}(s) \leq m \log N$ for every $s \notin \mathcal{L}$ we conclude that

$$c_0 m|S| \log Q \leq |\mathcal{L}|2Q + |S|3r \log N + \delta|S|m \log N,$$

where we used that

$$\sum_{p \leq Q} \log p \leq 2Q.$$

for large Q . Hence, by (2.4.2) we derive that

$$m \left(\frac{\varepsilon c_0}{2d} - \delta - \frac{4}{\log N} \right) \leq 3r.$$

Taking $m = 7r/\delta$ we get a contradiction for sufficiently large N . We may therefore assume that the set

$$A := \{s \in S' : \psi_{\mathcal{L}}(s) \geq 3r \log N\},$$

has size $|A| \geq \delta|S|$ for the above choices of m and δ .

We will now show that if an r -polynomial vanishes at \mathcal{L} , then it also vanishes at A . Indeed, let f be such a polynomial and let $x \in A$ be arbitrary. By definition, we have $|f(x)| < N^{3r}$. On the other hand, if p is a prime for which there exists some $y \in \mathcal{L}$ with $x \equiv y \pmod{p}$, then the fact that $f(y) = 0$ implies that $p|f(x)$. But by definition of A the product of all such p is at least N^{3r} so we see that the only way for this to hold is to have $f(x) = 0$, which proves our claim.

By the induction hypothesis and our construction of S' we know that for each $S'_{x_i} \in \mathcal{L}$ we may find a $(r, 1)$ -characteristic set of size $O(r^{d-h-1})$. Taking the union of these m sets we have thus found a set of size $O(r^{d-h})$ which is (r, δ) -characteristic for S , with δ as above. This concludes the proof of Proposition 2.3.

2.5 Ill-distributed sets with no algebraic structure

In this section we provide some examples of high dimensional ill-distributed sets possessing no algebraic structure. In particular, we show that the assertion of Theorem 2.1 fails when $\varepsilon = 0$. To begin with, we use a slight modification of the construction given in [30, §4.3] to see that, given any $0 < \eta < 1$, one may construct a subset of $[N]$ of size $\gg (\log N)^\eta$ which occupies at most p^η residue classes for every prime p and which possesses no algebraic structure. Indeed, if N is sufficiently large, we may find some integer Q with $Q < \log N < 2Q$ such that the product of all primes $p \leq Q$, say R , satisfies $N^{1/4} < R < N$ (this, of course, is very crude). For each prime $p \leq Q$ choose $\lfloor p^\eta \rfloor$ residue classes. Then, by the Chinese remainder theorem, there are $\sim R^\eta$ elements below R belonging to a selected class for every $p \leq Q$. Choose $\lfloor (\log N)^\eta / 2 \rfloor$ of these elements and call this set X . Notice that for all primes $p > Q$ we have $p^\eta > |X|$ and therefore X occupies at most p^η residue classes for these primes p . Since by construction it also occupies that many classes for all primes $p \leq Q$, we get the claim.

We now proceed to give some examples of ill-distributed sets with no algebraic structure. The first one already shows that Theorem 2.1 is best possible.

Example 2.11. This follows readily from the above construction. Fix some pair of positive integers d, h with $d \geq h + 1$ and consider $h + 1$ different sets X_1, \dots, X_{h+1} constructed as in the previous paragraph with $\eta = 1/(h + 1)$. If we define the set

$$S := \{(x_1, \dots, x_d) \in [N]^d : x_i \in X_i \forall 1 \leq i \leq h + 1\},$$

then we have that $|S| \gg N^{d-h-1} \log N$ while $|[S]_p| \leq p^{d-h}$ for every prime p , from where it follows that we cannot take $\varepsilon = 0$ in Theorem 2.1.

Example 2.12. One can generalize the above example by “perturbing” arbitrary algebraic sets. We show a simple instance of this. Let $d = 3$ and consider two polynomials $f, g \in \mathbb{Z}[x]$. Let X and Y be sets of size $\gg (\log N)^{1/2}$ occupying at most $p^{1/2}$ residue classes for every prime p . Then, we see that

$$\{(x, f(x) \cdot X, g(x) \cdot Y) : x \in [N]\}$$

is a big set of integer points occupying at most p^2 residue classes.

Finally, we show that not all possible counterexamples are perturbations of strongly algebraic sets.

Example 2.13. Fix some small $\varepsilon > 0$. By the Chinese remainder theorem one can construct a set $X \subseteq [N]$ of size $|X| \sim N^{1-\varepsilon}$ occupying only one residue class for every prime $p \leq \varepsilon \log N$. Take $K = \lfloor (\varepsilon \log N)^{1/3} \rfloor$ and let $f_1, \dots, f_K, g_1, \dots, g_K$ be a

family of polynomials. Also, let $X_1, \dots, X_K, Y_1, \dots, Y_K$ be arbitrary sets of size at most $(\varepsilon \log N)^{1/3}$. Then

$$\bigcup_{i=1}^K \{(x, f_i(x) \cdot X_i, g_i(x) \cdot Y_i) : x \in X\}$$

is a big set of integer points occupying at most p^2 residue classes for every prime p . Notice that this construction is of a different nature than the one given in Example 2.12, since the union of that many algebraic sets may not retain any algebraic structure itself.

It follows from the above examples that strange things can happen if one allows the set to possess too many very small sections. However, we shall show in Theorem 2.14 below that the methods presented in this chapter do indeed work as long as one avoids this type of situation.

2.6 Further results and conjectures

2.6.1 A generalization of Theorem 1.1

We now state the most general result which follows at once from our methods. Let $0 \leq k < d$ be integers and let $\varepsilon > 0$ be some positive real number. We say a set $S \subseteq [N]^d$ is $(1, \varepsilon)$ -regular if $|S| \geq N^\varepsilon$. Recursively, we say $S \subseteq [N]^d$ is (k, ε) -regular if there exists some $1 \leq i \leq d$ such that for every $x \in [N]$

1. $|\pi_i^{-1}(x) \cap S| \leq |S|/N^\varepsilon$,
2. $\pi_i^{-1}(x) \cap S$ is either empty or $(k-1, \varepsilon)$ -regular.

The first condition allows us to recover (2.4.2), while the second one enables us to use Lemma 2.9 and the induction hypothesis as it was done in the main argument. As a consequence, one recovers the conclusions preceding Lemma 2.10 and from here the proof of (the analogous of) Proposition 2.3 proceeds without further modifications. One can thus deduce that any (k, ε) -regular S set occupying $\ll p^k$ residue classes admits a bounded polynomial vanishing at a positive proportion of S . Furthermore, since it is easy to see that any subset $S' \subseteq S$ of a (k, ε) -regular set with $|S'| \geq \eta|S|$ admits a $(k, \varepsilon/2)$ -regular subset of half density (provided N is sufficiently large in terms of η) we can in fact deduce the following stronger result.

Theorem 2.14. *Let $0 \leq k < d$ be integers and let $\varepsilon, \eta, \alpha$ be positive real numbers. Then there exists $C = O_{\varepsilon, \eta, \alpha, k, d}(1)$ such that for every (k, ε) -regular set $S \subseteq [N]^d$ occupying less than αp^k residue classes for every prime p , there exists a nonzero polynomial $f \in \mathbb{Z}[x_1, \dots, x_d]$ of degree at most C and coefficients bounded by N^C vanishing at more than $(1 - \eta)|S|$ points of S .*

One would expect a reasonable set to be well approximated by a bounded union of (k, ε) -regular subsets, in which case it is clear that the same conclusion holds. For example, it was implicitly shown in the proof of Theorem 2.1 that given any $\eta > 0$ and any $S \subseteq [N]^d$ with $|S| \gg N^{k-1+\varepsilon}$ there exists some $S' \subseteq S$ with $|S'| \geq (1-\eta)|S|$ which is the union of a bounded number (in terms of η and d) of $(k, \varepsilon/4d)$ -regular subsets.

Finally, it is important to note that one cannot hope to do much better than Theorem 2.14 in this generality, since the regularity conditions are necessary in order to avoid those constructions emerging from the Chinese Remainder Theorem as in §2.5.

2.6.2 Approximate reduction

We shall give a quick application of Theorem 2.14 to the study of functions preserving some structure when reduced modulo a prime, that is, functions f for which knowing the class of $x \pmod{p}$ gives us information about the class of $f(x) \pmod{p}$. Thus, given a positive integer K , we say a function $f : [N]^k \rightarrow [N^r]^t$ has K -approximate reduction if

$$\left| [f([N]^k(\mathbf{a}))]_p \right| \leq K,$$

for every $\mathbf{a} \in (\mathbb{Z}/p\mathbb{Z})^k$ and every prime p . That is, f is said to have K -approximate reduction if once the residue class of x is fixed, there are at most K possibilities for the residue class of $f(x)$. When $K = 1$ this implies the very strong property of *recurrence mod p* and using this, it was shown by Hall [28] and Ruzsa [41] (see also [43, §XV.41]) that for large N the only functions having 1-approximate reduction are polynomials (notice that we are assuming our functions to have polynomial growth, which is in fact a necessary condition [28]). Since the graph of a function $f : [N]^k \rightarrow [N^r]^t$ is always a $(k, 1/2r)$ -regular set, it follows from Theorem 2.14 that this is indeed a very robust phenomenon:

Corollary 2.15. *Suppose $f : [N]^k \rightarrow [N^r]^t$ has K -approximate reduction and let $\Gamma(f)$ be the graph of f . Then there exists $C = O_{k,r,t,K}(1)$ and a polynomial $P \in \mathbb{Z}[x_1, \dots, x_d]$ of degree at most C and coefficients bounded by N^C , such that P vanishes at more than $(1-\eta)|\Gamma(f)|$ points of $\Gamma(f)$.*

2.6.3 The inverse sieve problem in one dimension

We conclude by mentioning a very strong version of the inverse sieve problem which is conjectured to hold for sets $S \subseteq [N]$ (see [14, Problem 7.2] and [30]).

Conjecture 2.16. *Suppose that $S \subseteq [N]$ is some set of integers of size $|S| \geq N^\varepsilon$ occupying less than αp residue classes for some $0 < \alpha < 1$ and every prime p . Then S has a large intersection with the image of a nonlinear integer polynomial of degree bounded in terms of α and ε .*

As a more precise instance of this, they conjecture for example that if a set S has size $|S| \geq N^{0.49}$ say, and occupies less than $2p/3$ residue classes mod p for every prime p , then most of S must be contained in a set of the form $\{an^2 + bn + c : n \in \mathbb{Z}\}$. This can be seen as an inverse conjecture for the large sieve [22, 40].

Conjecture 2.16 seems to be hard. In the particular case in which the residue classes occupied by S lie outside some interval of length $(p-1)/2$, nontrivial estimates were obtained by Green, and also by Bourgain, with the bound $|S| \ll (\log \log N)^c$ being recently achieved by Green and Harper [23]. In general, as noted by Helfgott and Venkatesh [30, §4.2], Conjecture 2.16 implies that there are $\ll_\varepsilon N^\varepsilon$ points on an irrational curve within a square of side N , which is itself a well known open problem.

Bibliografía

- [1] J. Avigad, P. Gerhardy, H. Towsner, *Local stability of ergodic averages*, Trans. Amer. Math. Soc. **362** (2010), 261-288.
- [2] T. Austin, *On the norm convergence of non-conventional ergodic averages*, Ergod. Th. and Dynam. Sys. **30** (2010), 321-338.
- [3] T. Austin, *Pleasant extensions retaining algebraic structure, I*, preprint available at [arXiv:0905.0518](https://arxiv.org/abs/0905.0518).
- [4] T. Austin, *Pleasant extensions retaining algebraic structure, II*, preprint available at [arXiv:0910.0907](https://arxiv.org/abs/0910.0907).
- [5] V. Bergelson, *Weakly mixing PET*, Ergod. Th. and Dynam. Sys. **7** (1987), 337-349.
- [6] E. Breuillard, B. Green, T. Tao, *The structure of approximate groups*, to appear in Publ. Math. IHES.
- [7] V. Bergelson, A. Leibman, *A nilpotent Roth theorem*, Invent. Math. **147** (2) (2002), 429-470.
- [8] V. Bergelson, T. Tao, T. Ziegler, *An inverse theorem for the uniformity seminorms associated with the action of \mathbb{F}_p^∞* , Geom. Funct. Anal. **19** (2010), no. 6, 1539–1596.
- [9] E. Bombieri, *Le Grand Crible dans la Théorie Analytique des Nombres*, Astérisque, vol. 18 (Société mathématique de France, Paris, 1974).
- [10] E. Bombieri, J. Pila, *The number of integral points on arcs and ovals*, Duke Math. J. **59** (1989), 337-357.
- [11] Q. Chu, N. Frantzikinakis, B. Host, *Ergodic averages of commuting transformations with distinct degree polynomial iterates*, Proc. London Math. Soc. **102** (5) (2011), 801-842.

- [12] A. C. Cojocaru, M. R. Murty, *Introduction to Sieve Methods and Their Applications*. London Mathematical Society Student Texts, Cambridge University Press, 2005.
- [13] J. P. Conze, E. Lesigne, *Théorèmes ergodiques pour des mesures diagonales*, Bull. Soc. Math. France **112** (2) (1984), 143-175.
- [14] E. Croot, V. F. Lev, *Open problems in additive combinatorics*, in Additive Combinatorics, CRM Proc. Lecture Notes **43**, 207-233, Amer. Math. Soc., Providence, RI, 2007.
- [15] N. Frantzikinakis, B. Kra, *Convergence of multiple ergodic averages for some commuting transformations*, Ergod. Th. and Dynam. Sys. **25** (2005), 799-809.
- [16] G. R. Freiman, *Foundations of a structural theory of set addition*, Kazan Gos. Ped. Inst., Kazan, 1966.
- [17] J. Friedlander, H. Iwaniec, *Opera de cribro*, American Mathematical Society Colloquium Publications, 57. American Mathematical Society, Providence, RI, 2010.
- [18] H. Furstenberg, *Ergodic behavior of diagonal measures and a theorem of Szemerédi on arithmetic progressions*, J. d'Analyse Math. **31** (1977), 204-256.
- [19] H. Furstenberg, B. Weiss, *A mean ergodic theorem for $\frac{1}{N} \sum_{n=1}^N f(T^n x) g(T^{n^2} x)$* . Convergence in Ergodic Theory and Probability. Walter de Gruyter, Berlin, New York (1996), 193-227.
- [20] W. T. Gowers, *Decompositions, approximate structure, transference, and the Hahn-Banach theorem*, Bull. London Math. Soc. **42** (2010), 573-606.
- [21] B. J. Green, *Approximate groups and their applications: work of Bourgain, Gamburd, Helfgott and Sarnak*, Current Events Bulletin of the AMS, 2010.
- [22] B. J. Green, *On a variant of the large sieve*, preprint, [arXiv:0807.5037](https://arxiv.org/abs/0807.5037).
- [23] B. J. Green, *Private communication*.
- [24] B. J. Green, I. Z. Ruzsa, *Freiman's theorem in an arbitrary abelian group*, J. Lond. Math. Soc. (2) **75** (2007), no. 1, 163-175.
- [25] B. J. Green, T. Tao, T. Ziegler, *An inverse theorem for the Gowers U^{s+1} norm*, Ann. of Math. **176** (2012), no. 2, 1231-1372.
- [26] P. X. Gallagher, *A larger sieve*, Acta Arith. **18** (1971), 77-81.
- [27] D. A. Goldston, J. Pintz, C. Y. Yildirim, *Primes in tuples I*, Ann. of Math. **170** (2009), no. 2, 819-962.

- [28] R. R. Hall, *On pseudopolynomials*, *Mathematika* **18** (1971), 71-77.
- [29] H. A. Helfgott, *Growth and generation in $SL_2(\mathbb{Z}/p\mathbb{Z})$* , *Ann. of Math. (2)* **167** (2008), no. 2, 601-623.
- [30] H. A. Helfgott, A. Venkatesh, *How small must ill-distributed sets be?*, *Analytic number theory. Essays in honour of Klaus Roth*. Cambridge University Press, 2009, 224-234.
- [31] B. Host, *Ergodic seminorms for commuting transformations and applications*, *Studia Math.* **195** (2009), 31-49.
- [32] B. Host, B. Kra, *Nonconventional ergodic averages and nilmanifolds*, *Ann. of Math. (2)* **161** (1) (2005), 397-488.
- [33] B. Host, B. Kra, *Convergence of polynomial ergodic averages*, *Isr. J. Math.* **149** (2005), 1-19.
- [34] E. Kowalski, *The ubiquity of surjective reduction in random groups*, unpublished notes, <http://www.math.ethz.ch/~kowalski/notes-unpublished.html>.
- [35] E. Kowalski, *The large sieve and its applications: arithmetic geometry, random walks and discrete groups*. Cambridge Tracts in Math. 175, Cambridge University Press, 2008.
- [36] B. Kra, *The Green-Tao theorem on arithmetic progressions in the primes: an ergodic point of view*, *Bull. Amer. Math. Soc.* **43** (2006), no. 1, 3-23.
- [37] A. Leibman, *Polynomial sequences in groups*, *Journal of Algebra* **201** (1998), 189-206.
- [38] A. Leibman, *Polynomial mappings of groups*, *Isr. J. Math* **129** (2002), 29-60.
- [39] A. Leibman, *Convergence of multiple ergodic averages along polynomials of several variables*, *Isr. J. Math.* **146** (2005), 303-316.
- [40] H. L. Montgomery, *A note on the large sieve*, *J. London Math. Soc.* **43** (1968), 93-98.
- [41] I. Z. Ruzsa, *On congruence preserving functions*, *Mat. Lapok.* **22** (1971), 125-134.
- [42] I. Z. Ruzsa, *Generalized arithmetic progressions and sumsets*, *Acta Math. Hungar.* **65** (1994), 379-388.
- [43] J. Sándor, D. S. Mitrinovic, B. Crstici, *Handbook of number theory I*, Springer, 2006.

- [44] E. Szemerédi, *On sets of integers containing no k elements in arithmetic progression*, Acta Arith. **27** (1975), 299-345.
- [45] T. Tao, *What is good mathematics?*, Bull. Amer. Math. Soc. **44** (2007), no. 4, 623-634.
- [46] T. Tao, *Norm convergence of multiple ergodic averages for commuting transformations*, Ergod. Th. and Dynam. Sys. **28** (2008), 657-688.
- [47] T. Tao, *Freiman's theorem for solvable groups*, Contrib. Discrete Math. **5** (2010), no. 2, 137-184.
- [48] T. Tao, V. Vu, *Additive combinatorics*, Cambridge Studies in Advanced Mathematics, 105. Cambridge University Press, Cambridge, 2006.
- [49] T. Tao, V. Vu, *Inverse Littlewood-Offord theorems and the condition number of random matrices*, Ann. of Math. (2) **169** (2009), 595-632.
- [50] T. Tao, T. Ziegler, *The inverse conjecture for the gowers norm over finite fields in low characteristic*, preprint, [arXiv:1101.1469](https://arxiv.org/abs/1101.1469)
- [51] H. Towsner, *Convergence of diagonal ergodic averages*, Ergod. Th. and Dynam. Sys. **29** (2009), 1309-1326.
- [52] M. Walsh, *The inverse sieve problem in high dimensions*, Duke Math. J. **161** (2012), no. 10, 2001-2022.
- [53] M. Walsh, *Norm convergence of nilpotent ergodic averages*, Ann. of Math. **175** (2012), no. 3, 1667-1688.
- [54] Q. Zhang, *On the convergence of the averages $(1/N) \sum_{n=1}^N f_1(R^n x) f_2(S^n x) f_3(T^n x)$* , Monatsh. Math. **122** (3) (1996), 275-300.
- [55] T. Ziegler, *Universal characteristic factors and Furstenberg averages*, J. Amer. Math. Soc. **20** (2007), 53-97.