TRANSACTIONS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 370, Number 11, November 2018, Pages 8135–8153
http://dx.doi.org/10.1090/tran/7432
Article electronically published on August 9, 2018

LIMITS OF FUNCTIONS ON GROUPS

BALÁZS SZEGEDY

ABSTRACT. Our goal is to develop a limit approach for a class of problems in additive combinatorics that is analogous to the limit theory of dense graph sequences. We introduce a metric, convergence and limit objects for functions on discrete groups and use it to study limits of measurable functions on compact abelian groups. As an application we find exact minimizers for densities of linear configurations of complexity 1.

1. Introduction

The so-called graph limit theory (see [10], [11], [2], [9]) gives an analytic approach to a large class of problems in graph theory. A very active field of applications is extremal graph theory where, roughly speaking, the goal is to find the maximal (or minimal) possible value of a graph parameter in a given family of graphs and to study the structure of graphs attaining the extremal value. A classical example is Mantel's theorem which implies that a triangle free graph H on $2n$ vertices maximizes the number of edges if H is the complete bipartite graph with equal color classes. Another example is given by the Chung-Graham-Wilson theorem [3]. If we wish to minimize the density of the four cycles in a graph H with edge density $1/2$, then H has to be sufficiently quasi-random. However, the perfect minimum of the problem (that is $1/16$) cannot be attained by any finite graph but one can get arbitrarily close to it. Such problems justify graph limit theory where in an appropriate completion of the set of graphs the optimum can always be attained if the extremal problem satisfies a certain continuity property. Furthermore, one can use variational principles at the exact maximum or minimum bringing the tools of differential calculus into graph theory.

Extremal graph (and hypergraph) theory has a close connection to additive combinatorics. It is well known that the triangle removal lemma by Szemerédi and Ruzsa implies the qualitative version of Roth's theorem on three term arithmetic progressions. The proof relies on an encoding of an integer sequence (or a subset in an abelian group) by a graph that is rather similar to a Cayley graph. Such representations of additive problems in graph theory hint at a limit theory for subsets in abelian groups that is closely connected to graph limit theory. This new limit theory, that is actually a limit theory for functions on abelian groups, was initiated by the author in [15], [16], and [13] in a rather general form.

Received by the editors February 27, 2015, and, in revised form, March 15, 2017.

2010 *Mathematics Subject Classification.* Primary 05D99; Secondary 43A99.

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement n°617747. The research was partially supported by the MTA Rényi Institute Lendület Limits of Structures Research Group.

Motivated by Szemerédi's theorem on arithmetic progressions Gowers initiated a theory of higher order Fourier analysis in [6] (for a textbook on the topic see [18]). He introduced a sequence of norms $\|.\|_{U_k}$ (called uniformity norms) for functions on finite abelian groups. Roughly speaking, in k-th order Fourier analysis functions with small U_{k+1} norm are considered to be "random like". Seperation of noise and structure is a central topic in higher order Fourier analysis. The larger k is, the more functions are considered to be structured and their description gets increasingly harder. Correspondingly, there is a hierarchy of increasingly fine limit notions related to k-th order Fourier analysis as k goes to infinity and the limit objects get increasingly complex. The focus of this paper is the linear case $k = 1$ that was called "harmonic analytic limit" in [15]. This case is interesting in its own right, covers numerous important questions, and is illustrative for the more general limit concept.

We introduce metric, convergence and limit objects for subsets in abelian groups. More generally, since subsets can be represented by their characteristic functions, we study the convergence of functions on abelian groups. This extends the range of possible applications of our approach to problems outside additive combinatorics.

In the first part of the paper we study a metric $\hat{d}$ and a related convergence notion for l^2 functions on discrete (not necessarily commutative) groups. It is important that the metric $\hat{d}$ allows us to compare two functions defined on different groups. In chapter 3 we introduce a distance d for measurable functions $f \in L^2(A_1), g \in L^2(A_2)$ defined on compact abelian groups A_1, A_2 such that $d(f,g) := \hat{d}(\hat{f},\hat{g})$ where $\hat{f}$ and $\hat{g}$ denote the Fourier transforms of f and g. In additive combinatorics, we can use the distance d to compare subsets in finite abelian groups in the following way. If $S_1 \subseteq A_1$ and $S_2 \subseteq A_2$ are subsets in finite abelian groups A_1 and A_2, then their distance is $d(1_{S_1}, 1_{S_2})$. This allows us to talk about convergent sequences of subsets in a sequence of abelian groups.

A crucial property of the metric d (see Theorem 3) is that it puts a compact topology on the set of all pairs (f, A) where A is a compact abelian group and f is a measurable function on A with values in a fixed compact convex set $K \subset \mathbb{C}$. As a consequence we have that any sequence of subsets $\{S_i \subseteq A_i\}_{i=1}^{\infty}$ in finite abelian groups A_i has a convergent subsequence with limit object which is a measurable function of the form $f : A \to [0,1]$ where A is some compact abelian group. This result is analogous to graph limit theory where graph sequences always have convergent subsequences with limit object which is a symmetric measurable function of the form $W : [0,1]^2 \to [0,1]$.

The success of a limit theory depends on how many interesting parameters are continuous with respect to the convergence notion. The parameters that are most interesting in additive combinatorics are densities of linear configurations. A linear configuration is given by a finite set of linear forms, i.e., homogeneous linear multivariate polynomials over $\mathbb{Z}$. For example, a 3-term arithmetic progression is given by the linear forms $a,\ a+b,\ a+2b$. If f is a bounded measurable function on a compact abelian group A, then we can compute the density of 3-term arithmetic progressions in f as the expected value $\mathbb{E}_{a,b\in A}(f(a)f(a+b)f(a+2b))$ according to the normalized Haar measure on A. This density concept can be generalized to an arbitrary linear configuration $\mathcal{L} = \{L_1, L_2, \ldots, L_k\}$ and the density of $\mathcal{L}$ in f is denoted by $t(\mathcal{L}, f)$ (see formula (1) and the following sentence). Gowers and Wolf introduced a complexity notion [7] for linear configurations called *true complexity* (see Definition 4.1 in this paper). A useful upper bound for the true complexity is the so-called Cauchy-Schwarz complexity developed by Green and Tao in [8].

We prove the following fact (for precise formulation see Theorem 5).

Theorem 1. *If $\mathcal{L}$ has true complexity at most 1, then the density function of $\mathcal{L}$ is continuous in the metric d.*

Examples for linear configurations of complexity 1 include the 3-term arithmetic progression (this was shown in [8]), the parallelogram $a,\ a+b,\ a+c,\ a+b+c$, and the system $\mathcal{L}_H := \{x_i + x_j \ :\ (i,j) \in E(H)\}$ where H is an arbitrary finite graph on $\{1, 2, \dots, n\}$. The last example gives a close connection with graph limit theory. The density of $\mathcal{L}_H$ in $f \in L^\infty(A)$ is equal to the density of the graph H in the symmetric kernel $W : A \times A \to \mathbb{C}$ defined by $W(x,y) = f(x+y)$. Note that if f has values in $[0,1]$, then W is a graphon in the graph limit language. We will elaborate on this connection in chapter 10.

Let $\mathcal{L}$ be an arbitrary linear configuration. For $0 \leq \delta \leq 1$ and $n \in \mathbb{N}$ let $\rho(\delta, n, \mathcal{L})$ denote the minimal possible density of $\mathcal{L}$ in subsets of $\mathbb{Z}_n$ of size at least δn. Let $\rho(\delta, \mathcal{L}) := \liminf_{p\to\infty} \rho(\delta, p, \mathcal{L})$ where p runs through the prime numbers. A result by Candela and Sisask [1] implies that the lim inf can be replaced by lim in the definition of $\rho(\delta, \mathcal{L})$. Note that the qualitative version of Roth's theorem is equivalent with the fact that $\rho(\delta, \mathcal{L}) > 0$ if $\delta > 0$ and $\mathcal{L} = \{a,\ a+b,\ a+2b\}$.

Theorem 2. *Let $\mathcal{L}$ be a linear configuration of true complexity at most 1. For every $0 \leq \delta \leq 1$ we have that*

$$\rho(\delta, \mathcal{L}) = \min_{A,f}(t(\mathcal{L}, f)),$$

where f runs through all measurable functions of the form $f : A \to [0,1]$ with $\mathbb{E}(f) = \delta$ on compact abelian groups A with torsion-free Pontrjagin dual groups.

We emphasize that in Theorem 2 we obtain $\rho(\delta, \mathcal{L})$ as an actual minimum and thus there is some function $f_{\delta,\mathcal{L}}$ realizing the value $\rho(\delta)$. If for example $\mathcal{L} = \{a,\ a+b,\ a+2b\}$, then it is not hard to deduce the qualitative version of Roth's theorem from Theorem 2 using Lebesgue's density theorem. We sketch the proof at the end of chapter 10. It would be very interesting to find the explicit form of a minimizer $f_{\delta,\mathcal{L}}$ for every δ or even to obtain any information on $f_{\delta,\mathcal{L}}$ such as on which abelian group it is defined.

It is important to mention that our convergence notion behaves quite differently from usual convergence notions in functional analysis. There is an example for a convergent sequence of functions, all of them defined on the circle (complex unit circle with multiplication or equivalently the quotient group $\mathbb{R}/\mathbb{Z}$), but the limit object exists only on the torus (see the example at the end of chapter 3).

In the proofs we will extensively use ultralimit methods. Ultralimit methods in graph and hypergraph regularization and limit theory were first introduced in [4]. There are two different reasons to use these methods. One is that they seem to help to get rid of a great deal of technical difficulties and provide cleaner proofs for most of our statements. The other reason is that they point to an interesting connection between ergodic theory and our limit theory. The ultraproduct $\mathbf{A}$ of compact abelian groups $\{A_i\}_{i=1}^\infty$ behaves as a measure preserving system. Our limit concept can easily be explained through a factor $\mathcal{F}(\mathbf{A})$ of $\mathbf{A}$ which is a variant of the so-called Kronecker factor.

2. A limit notion for functions on discrete groups

For an arbitrary group G we denote by $l^2(G)$ the Hilbert space of all functions $f : G \to \mathbb{C}$ such that $\|f\|_2^2 = \sum_{g\in G} |f(g)|^2 < \infty$. If $f \in l^2(G)$ and $\epsilon \geq 0$, then we denote by $\mathrm{supp}_\epsilon(f)$ the set $\{g : g \in G, |f(g)| > \epsilon\}$. In particular, $\mathrm{supp}(f) := \mathrm{supp}_0(f)$ is the support of f. Note that if $\epsilon > 0$, then $|\mathrm{supp}_\epsilon(f)| \leq \|f\|_2^2/\epsilon^2$ and $\mathrm{supp}(f) = \bigcup_{n=1}^\infty \mathrm{supp}_{1/n}(f)$ is a countable (potentially finite) set. We denote by $\langle f\rangle$ the subgroup of G generated by $\mathrm{supp}(f)$. It is clear that $\langle f\rangle$ is a countable (potentially finite) group.

Two functions $f_1 \in l^2(G_1)$ and $f_2 \in l^2(G_2)$ are called isomorphic if there is a group isomorphism $\alpha : \langle f_1\rangle \to \langle f_2\rangle$ such that $f_1 = f_2 \circ \alpha$. Let us denote by $\mathcal{M}$ the isomorphism classes of l^2 functions on groups. Our goal is to define a metric space structure on $\mathcal{M}$. We will need the next definition.

Definition 2.1. Let G_1 and G_2 be groups. A partial isomorphism of weight n is a bijection $\phi : S_1 \to S_2$ between two subsets $S_1 \subseteq G_1, S_2 \subseteq G_2$ such that $g_1^{\alpha_1} g_2^{\alpha_2} \dots g_n^{\alpha_n} = 1$ holds if and only if $\phi(g_1)^{\alpha_1}\phi(g_2)^{\alpha_2}\dots\phi(g_n)^{\alpha_n} = 1$ for every sequence $g_i \in S_1, \alpha_i \in \{-1, 0, 1\}$ with $1 \leq i \leq n$.

Definition 2.2. Let $f_1 \in l^2(G_1)$ and $f_2 \in l^2(G_2)$. An ϵ-isomorphism between f_1 and f_2 is a partial isomorphism $\phi : S_1 \to S_2$ of weight $\lceil 1/\epsilon\rceil$ between sets with $\mathrm{supp}_\epsilon(f_1) \subseteq S_1 \subseteq G_1$ and $\mathrm{supp}_\epsilon(f_2) \subseteq S_2 \subseteq G_2$ such that $|f_1(g) - f_2(\phi(g))| \leq \epsilon$ holds for every $g \in S_1$. We define $\hat{d}(f_1, f_2)$ as the infimum of all ϵ's such that there is an ϵ-isomorphism between f_1 and f_2.

Note that both partial isomorphism and ϵ-isomorphism are symmetric notions in the sense that if ϕ is a partial isomorphism (resp., ϵ-isomorphism), then ϕ^{-1} is also a partial isomorphism (resp., ϵ-isomorphism).

Proposition 2.1. *The function $\hat{d}$ is a metric on $\mathcal{M}$.*

Proof. First we show that $\hat{d}(f_1, f_2) = 0$ if and only if f_1 and f_2 are isomorphic. If f_1 is isomorphic to f_2, then it is clear that $d(f_1, f_2) = 0$. For the other direction assume w.l.o.g. that $\|f_2\|_2 \leq \|f_1\|_2$. Let $\alpha_n : S_{1,n} \to S_{2,n}$ be a $1/n$-isomorphism between f_1 to f_2 for every n. Clearly, for every element $g \in \mathrm{supp}(f_1)$ there are finitely many possible elements in the sequence $\{\alpha_n(g)\}_{n=1}^\infty$ since $\lim_{n\to\infty} f_2(\alpha_n(g)) = f_1(g)$ and there are finitely many elements h in G_2 on which $|f_2(h)| > |f_1(g)|/2$. Using that the support of f_1 is countable we obtain that there is a subsequence $\{\beta_n\}$ of $\{\alpha_n\}$ such that the sequences $\{\beta_n(g)\}$ stabilize (become constant) after finitely many steps for every g with $|f_1(g)| > 0$. This defines a map $\beta = \lim \beta_n$ from $\mathrm{supp}(f_1)$ to $\mathrm{supp}(f_2)$. It is clear that β extends to an injective homomorphism from $\langle f_1\rangle$ to $\langle f_2\rangle$ and it satisfies $f_2(\beta(g)) = f_1(g)$ for every $g \in \langle f_1\rangle$. Using $\|f_2\|_2 \leq \|f_1\|_2$ it follows that every element in $\mathrm{supp}(f_2)$ is in the image of β and so β is a value preserving isomorphism between $\langle f_1\rangle$ and $\langle f_2\rangle$.

It remains to check the triangle inequality for the metric d. Assume that $\alpha : S_1 \to S_2$ is an ϵ-isomorphism between f_1 and f_2 and assume that $\beta : S_2' \to S_3$ is an ϵ'-isomorphism between f_2 and f_3. Without loss of generality we can assume (by reversing arrows if necessary) that $\epsilon' \geq \epsilon$. We have the following inclusions:

$$\beta^{-1}(\mathrm{supp}_{\epsilon'+\epsilon}(f_3)) \subseteq \beta^{-1}(\mathrm{supp}_{\epsilon'}(f_3)) \subseteq \beta^{-1}(S_3) = S_2',$$
$$\beta^{-1}(\mathrm{supp}_{\epsilon'+\epsilon}(f_3)) \subseteq \mathrm{supp}_{\epsilon}(f_2) \subseteq S_2,$$
$$\alpha(\mathrm{supp}_{\epsilon'+\epsilon}(f_1)) \subseteq \mathrm{supp}_{\epsilon'}(f_2) \subseteq S_2 \cap S_2'.$$

Let $T_2 = \beta^{-1}(\mathrm{supp}_{\epsilon'+\epsilon}(f_3)) \cup \mathrm{supp}_{\epsilon'}(f_2)$ (observe that $T_2 \subseteq S_2 \cap S_2'$) and let $T_1 = \alpha^{-1}(T_2)$, $T_3 = \beta(T_2)$. We have that $\mathrm{supp}_{\epsilon'+\epsilon}(f_1) \subseteq T_1$ and $\mathrm{supp}_{\epsilon'+\epsilon}(f_3) \subseteq T_3$. Let $\gamma : T_1 \to T_3$ be the restriction of $\beta \circ \alpha$ to T_1. Using $T_2 \subseteq S_2 \cap S_2'$ we get that γ is a bijection. To complete the proof of the triangle inequality we show that γ is an $(\epsilon' + \epsilon)$-isomorphism. We have that γ is a bijection and that $|f_1(g_1) - f_3(\gamma(g_1))| \leq \epsilon' + \epsilon$ holds for every $g \in T_1$. It remains to check that γ is a partial isomorphism of weight $\lceil 1/(\epsilon'+\epsilon)\rceil$. This follows from the fact that the composition of a partial isomorphism of weight n and a partial isomorphism of weight m is a partial isomorphism of weight $\min(n, m)$. However, the minimum of $\lceil 1/\epsilon\rceil$ and $\lceil 1/\epsilon'\rceil$ is at least $\lceil 1/(\epsilon' + \epsilon)\rceil$. $\square$

Lemma 2.1. *Assume that a sequence $\{f_i\}_{i=1}^{\infty}$ of l^2 functions on abelian groups converges in $\hat{d}$ to $f \in l^2(G)$. Then $\langle f\rangle$ is also abelian.*

Proof. Let $g_1, g_2 \in \mathrm{supp}(f)$ be two elements. Let $\epsilon = \min(|f(g_1)|/2, |f(g_2)|/2, 1/4)$. Then by convergence of f_i there is an index i such that there is an ϵ-isomorphism ϕ between f and f_i. Since $g_1, g_2 \in \mathrm{supp}_\epsilon f$ we have that ϕ is defined on g_1, g_2 and $\phi(g_1)\phi(g_2)\phi(g_1)^{-1}\phi(g_2)^{-1} = 1$ implies that $g_1g_2g_1^{-1}g_2^{-1} = 1$ because $\epsilon < 1/4$. $\square$

For every real number $a > 0$ let $\mathcal{M}_a$ denote the subset of $\mathcal{M}$ consisting of equivalence classes of functions $f \in l^2(G)$ with $\|f\|_2 \leq a$.

Proposition 2.2. *The metric space $(\mathcal{M}_a, \hat{d})$ is compact for every $a > 0$.*

For the proof of Proposition 2.2 we will need the next lemma. Let F_r denote the free group in r generators.

Lemma 2.2. *Assume that $\{G_n\}_{n=1}^{\infty}$ is a sequence of groups and for every n we have a sequence of elements $\{g_{n,i}\}_{i=1}^{\infty}$ in G_n. Then there is a sequence of elements $\{g_i\}_{i=1}^{\infty}$ in some group G and a set $S \subseteq \mathbb{N}$ such that for every $r \in \mathbb{N}$ and word $w \in F_r$ there is a natural number N_w such that if $k \in S$ and $k > N_w$, then $w(g_{k,1}, g_{k,2}, \ldots, g_{k,r}) = 1$ if and only if $w(g_1, g_2, \ldots, g_r) = 1$.*

Proof. Let $\{w_i\}_{i=1}^{\infty}$ be an arbitrary ordering of the words in $\bigcup_{r=1}^{\infty} F_r$ with $w_i \in F_{r_i}$. We construct a sequence of infinite subsets $S_i \subseteq \mathbb{N}$ in a recursive way. Assume that $S_0 = \mathbb{N}$. If S_{i-1} is already constructed, then we construct S_i in a way that S_i is an infinite subset in S_{i-1} and either $w_i(g_{s,1}, g_{s,2}, \ldots, g_{s,r_i}) = 1$ holds for every $s \in S_i$ or $w_i(g_{s,1}, g_{s,2}, \ldots, g_{s,r_i}) \neq 1$ holds for every $s \in S_i$. This can be clearly achieved since S_{i-1} is infinite and thus at least one of the two options holds infinitely many times for indices inside S_{i-1}. We then choose a sequence $\{s_i\}_{i=1}^{\infty}$ such that $s_i \in S_i$ and $s_i < s_j$ hold for every pair $i < j$. We obtain for $\{s_i\}_{i=1}^{\infty}$ that for every $r \in \mathbb{N}$ and word $w \in F_r$ either $w(g_{s_i,1}, g_{s_i,2}, \ldots, g_{s_i,r}) = 1$ holds with finitely many exceptions or $w_r(g_{s_i,1}, g_{s_i,2}, \ldots, g_{s_i,r}) \neq 1$ holds with finitely many exceptions. Let W denote the collection of words for which the first case holds. Let G be the group with generators $\{g_i\}_{i=1}^{\infty}$ and relations $\{w(g_1, g_2, \ldots, g_r) = 1 | r \in \mathbb{N}, w \in F_r \cap W\}$. It is clear from the construction of W that every relation that G satisfies in its generators is already listed in W. This follows from the fact that if a word w is not in W, then for an arbitrary finite subset W' in W there is a witness among the groups G_{s_i} in which w does not hold but all words in W' hold. Now we have that $S = \{s_i\}_{i=1}^{\infty}$ and G with $\{g_i\}_{i=1}^{\infty}$ satisfies the lemma. □

Proof of Proposition 2.2. Let $\{f_n : G_n \to \mathbb{C}\}_{n=1}^{\infty}$ be a sequence of functions of l^2 norm at most a. For every n let $\{g_{n,i}\}_{i=1}^{\infty}$ be an ordering of the elements in $\text{supp}(f_n)$ in such a way that $|f_n(g_{n,i})| \geq |f_n(g_{n,j})|$ whenever $i < j$. If $\text{supp}(f_n)$ is finite, then, to make the list infinite, we add additional elements from outside $\text{supp}(f_n)$ to the list. If the group G_n is finite, then we enlarge G_n to an infinite group containing G_n (say $G_n \times \mathbb{Z}$) such that f_n takes the value 0 on the new group elements and then we can make the list infinite with elements from outside G_n.

Let $S \subseteq \mathbb{N}$, G and $\{g_i\}_{i=1}^{\infty}$ be chosen for the sequences $\{g_{n,i}\}_{i=1}^{\infty}$ according to Lemma 2.2. Let $S' \subseteq S$ be an infinite subset of S such that

$$a_i := \lim_{n\to\infty, n\in S'} f_n(g_{n,i})$$

exists for every $i \in \mathbb{N}$. Now we define the function $f : G \to \mathbb{C}$ such that $f(g_i) = a_i$ inside the set $\{g_i\}_{i=1}^{\infty}$ and $f(g) = 0$ for the rest of the elements. It is clear that f is well defined since $g_{n,i} \neq g_{n,j}$ holds for every n if $i \neq j$ and thus $g_i \neq g_j$. It is clear that $\|f\|_2 \leq \liminf_{n\in S'} \|f_n\|_2$ and thus $\|f\|_2 \leq a$.

To create an ϵ-isomorphism between f and f_n (if $n \in S'$ is large enough) we consider the sets $T_n = \{g_{n,i} : i \leq a^2/\epsilon^2\}$ and the set $T = \{g_i : i \leq a^2/\epsilon^2\}$. Let $\alpha_n : T_n \to T$ be the bijection defined by $\alpha_n(g_{n,i}) = g_i$. It is clear that $\text{supp}_\epsilon(f_n) \subseteq T_n$ holds for every n and that $\text{supp}_\epsilon(f) \subseteq S$. The construction guarantees that $|f_n(g) - f(\alpha_n(g)| \leq \epsilon$ holds if $n \in S'$ is large enough. Furthermore, the property given by Lemma 2.2 shows that α_n is a partial isomorphism of weight m for an arbitrary $m \in \mathbb{N}$ if $n \in S'$ is large enough. This completes the proof. □

3. Convergence notions on compact abelian groups

In this chapter we deal with compact abelian groups, Haar measure, Fourier transform, and Pontrjagin duality. The tools that we use are covered in the textbook [17]. Compact abelian groups in this paper will be assumed to be second countable and thus the Pontrjagin dual group is always countable. For a compact abelian group G we denote by $L^2(G)$ the Hilbert space of Borel measurable complex valued functions f on G with $\int |f|^2 \, d\mu \leq \infty$ where μ is the normalized Haar measure.

If $H \subseteq G$ is a closed subgroup of G, then we have that $\tau_H : G \to G/H$ is continuous and Haar measure preserving. Let $L^2_H(G)$ denote the Hilbert subspace $\tau_H \circ L^2(G/H)$ in $L^2(G)$. We have that $L^2_{H_1}(G) \cap L^2_{H_2}(G) = L^2_{\langle H_1,H_2\rangle}(G)$. It follows that for $f \in L^2(G)$ there is a unique largest closed subgroup $H(G,f)$ such that $f \in L^2_{H(G,f)}(G)$. In other words $H(G,f)$ is the largest closed subgroup of G such that there is a unique function $f' \in L^2(G/H(G,f))$ with $f = \tau_{H(G,f)} \circ f'$. It is clear that $H(G/H(G,f),f')$ is trivial. We say that the function $f' \in L^2(G/H(G,f))$ is the *economic representation* of $f \in L^2(G)$.

The economic representation can also be described through Fourier transforms. Let $\hat{G}$ denote the Pontrjagin dual of G and let $\hat{f} \in l^2(\hat{G})$ denote the Fourier transform of f. For a closed subgroup $H \subseteq G$ we have a natural embedding of $\widehat{G/H}$ into $\hat{G}$. We have that $f \in L^2_H(G)$ if and only if $\widehat{G/H}$ contains the support of $\hat{f}$ inside $\hat{G}$. It follows that the economic representation f' of f is the Fourier transform of the restriction of f to the group $\langle\text{supp}(\hat{f})\rangle$ generated by the support of $\hat{f}$. In particular,

f' is defined on the dual group of $\langle \mathrm{supp}(\hat{f})\rangle$ which is the factor group of G with the subgroup H is the intersection of the kernels of the characters in $\langle \mathrm{supp}(\hat{f})\rangle$.

Let $f_1 \in L^2(G_1)$ and $f_2 \in L^2(G_2)$ be functions on the compact abelian groups G_1 and G_2 with economic representations $(f_1', G_1/H(G_1, f_1))$ and $(f_2', G_2/H(G_2, f_2))$. We say that f_1 and f_2 are *isomorphic* if and only if there is a continuous isomorphism $\phi : G_1/H(G_1, f_1) \to G_2/H(G_2, f_2)$ such that $f_1' = \phi \circ f_2'$. It is clear that this notion of isomorphism is an equivalence relation. Using the above dual description of economic representations we have that f_1 and f_2 is isomorphic if and only if $\hat{f}_1$ is isomorphic to $\hat{f}_2$ in the sense of chapter 2.

Note that f_1, f_2 are isomorphic if and only if there is a third function $f_3 \in L^2(G_3)$ and continuous epimorphisms $\alpha_i : G_i \to G_3$ for $i = 1, 2$ such that $f_3(\alpha_i(g)) = f_i(g)$ holds for almost every g with respect to the Haar measure in G_i. This follows from the fact that the economic representations of f_1 and f_2 must factor through α_1 and α_2.

Let $\mathcal{H}$ denote the set of isomorphism classes of Borel measurable L^2 functions on compact abelian groups. We introduce the distance d on $\mathcal{H}$ by $d(f_1, f_2) := \hat{d}(\hat{f}_1, \hat{f}_2)$. The metric d induces a convergence notion on $\mathcal{H}$. If we say $\{f_i\}_{i=1}^{\infty}$ is convergent, then we mean convergence in d if not stated explicitly in which other meaning it is convergent. Let $\mathcal{H}_a$ denote the set of functions in $\mathcal{H}$ with L^2-norm at most a. Using the fact that Fourier transform preserves the L^2-norm we have by Lemma 2.1 and Proposition 2.2 the following statement.

Proposition 3.1. *$(\mathcal{H}_a, d)$ is a compact metric space for every $a > 0$.*

For a set $K \subseteq \mathbb{C}$ let $\mathcal{H}(K)$ denote the set of functions in $\mathcal{H}$ which take values in K. We will prove the next theorem.

Theorem 3. *If $K \subseteq \mathbb{C}$ is a compact convex set, then $(\mathcal{H}(K), d)$ is a compact metric space.*

Corollary 3.1. *If $\{f_i\}_{i=1}^{\infty}$ is a sequence of $\{0, 1\}$ valued functions in $\mathcal{H}$ converging to f in the metric d, then the values of f are in the interval $[0, 1]$.*

Theorem 3 is somewhat surprising. The metric d is given in terms of Fourier transforms, however, it is not trivial to relate the set of values of a function to the properties of its Fourier transform. The condition that K is convex turns out to be necessary in Theorem 3. Corollary 3.1 is useful when we study limits of sets in abelian groups by the limits of their characteristic functions. We give the proof of Theorem 3 in a later chapter.

In general if $\{f_i\}_{i=1}^{\infty}$ converges to f (in the sense of this paper) it is not necessarily true that $\{\|f_i\|_2\}_{i=1}^{\infty}$ converges to $\|f\|_2$. We only have that $\limsup_{i\to\infty} \|f_i\|_2 \geq \|f\|_2$. This motivates the next definition. We say that a sequence $\{f_i\}_{i=1}^{\infty}$ in $\mathcal{H}$ is *tightly convergent* if it converges in d and the limit f satisfies $\lim_{i\to\infty} \|f_i\|_2 = \|f\|_2$. Tight convergence can be metrized by the distance

$$d'(f_1, f_2) := d(f_1, f_2) + |\|f_1\|_2 - \|f_2\|_2|.$$

Convergence in d' is stronger than convergence in d and it has stronger consequences. To formulate our result we need the following notation. For a measurable function f on a compact abelian group A we denote by μ_f the probability distribution of $f(x)$ where x is chosen randomly from A according to the Haar measure. The measure μ_f is a Borel probability distribution on $\mathbb{C}$.

Theorem 4. *Let $\{f_i\}_{i=1}^{\infty}$ be a sequence of uniformly bounded functions in $\mathcal{H}$ converging to f in d'. Then μ_{f_i} converges to μ_f in the weak topology of measures.*

Note that the above theorem is not true for convergence in d. A trivial example for a tightly convergent sequence is an L^2-convergent sequence of functions on a fixed compact abelian group A. However, there are more interesting examples. We finish this chapter with an example which shows that a sequence of L^2 functions on the circle group $\mathbb{R}/\mathbb{Z}$ can have a limit (even in d') which cannot be defined on the circle group. The limit object exists on the torus. Let $f_n(x) = e^{2i\pi x} + e^{2in\pi x}$ be defined on $\mathbb{R}/\mathbb{Z}$ for $n \in \mathbb{N}$. It is easy to see that f_n is convergent and the limit is the function $f(x,y) = e^{2i\pi x} + e^{2i\pi y}$ on the torus $\mathbb{R}/\mathbb{Z} \times \mathbb{R}/\mathbb{Z}$. Note that the sequence f_n is tightly convergent since $\|f_n\|_2 = \|f\|_2 = \sqrt{2}$.

4. Densities of linear configurations in functions on abelian groups

In this chapter we state our main theorem regarding the convergence of the densities of linear configurations of complexity 1. We will follow the language introduced by Gowers and Wolf in [7]. Recall from the introduction that a linear form is a homogeneous linear multivariate polynomial with coefficients in $\mathbb{Z}$. If $L = \lambda_1 x_1 + \lambda_2 x_2 + \cdots + \lambda_n x_n$ is a linear form, then we can evaluate it in an arbitrary abelian group A by giving values from A to the variables x_i and thus it becomes a function of the form $L : A^n \to A$. A system $L_1, L_2, \ldots, L_k$ of linear forms determines a type of linear configuration. An example for a linear configuration is the 3-term arithmetic progression which is encoded by the linear forms $x_1,\ x_1 + x_2,\ x_1 + 2x_2$. Assume that A is a compact abelian group and $\mathcal{F} = \{f_i\}_{i=1}^{k}$ is a system of bounded measurable functions in $L^\infty(A)$. Assume furthermore that $\mathcal{L} = \{L_1, L_2, \ldots, L_k\}$ is a system of linear forms in $\mathbb{Z}(x_1, x_2, \ldots, x_n)$. Then it is usual to define the density of the configuration $\mathcal{L}$ in $\mathcal{F}$ by the formula

$$t(\mathcal{L}, \mathcal{F}) := \mathbb{E}_{x_1,x_2,\ldots,x_n \in A} \prod_{i=1}^{k} f_i(L_i(x_1, x_2, \ldots, x_n)). \tag{1}$$

If $f_i = f$ for every $1 \le i \le k$ in the function system $\mathcal{F}$, then we use the notation $t(\mathcal{L}, f)$ for $t(\mathcal{L}, \mathcal{F})$.

In this chapter we address the following type of problem.

Assume that $\mathcal{L} = \{L_1, L_2, \ldots, L_k\}$ is a linear configuration and $\mathcal{A}$ is a class of compact abelian groups. Under what conditions on $\mathcal{L}$ and $\mathcal{A}$ is the function $f \mapsto t(\mathcal{L}, f)$ continuous in the metric d when functions are assumed to be uniformly bounded measurable functions on groups in $\mathcal{A}$?

The role of the class $\mathcal{A}$ is to exclude certain degeneracies that occur for number theoretic reasons. For example, the linear form $2x$ becomes degenerated on the elementary abelian group $(\mathbb{Z}/2\mathbb{Z})^m$. We will need the following definition introduced by Gowers and Wolf in a slightly different form in [7].

Definition 4.1. Let $\mathcal{L} = \{L_1, L_2, \ldots, L_k\}$ be a linear configuration. The **true complexity** of $\mathcal{L}$ in a class $\mathcal{A}$ of abelian groups is the smallest number $m \in \mathbb{N}$ with the following property. For every $\epsilon > 0$ there exists $\delta > 0$ such that if $A \in \mathcal{A}$ is any abelian group and $\mathcal{F} = \{f_i\}_{i=1}^{k}$ is a system of measurable functions with $|f_i| \le 1$ and $\|f_j\|_{U_{m+1}} \le \delta$ for some j, then $t(\mathcal{L}, \mathcal{F}) \le \epsilon$.

In the above definition $\|.\|_{U_{m+1}}$ denotes Gowers's $m+1$-th uniformity norm. Our main theorem states the following.

Theorem 5. *Let $a > 0$. Let $\mathcal{L}$ be a linear configuration and let $\mathcal{A}$ be a family of compact abelian groups such that $\mathcal{L}$ has true complexity at most 1 in $\mathcal{A}$. Then $f \to t(\mathcal{L}, f)$ is continuous with respect to the metric d for measurable functions $f \in L^\infty(A)$ with $A \in \mathcal{A}$ and $|f| \leq a$.*

5. Ultraproducts and ultralimits

Let ω be a non-principal ultrafilter on the natural numbers. Let $\{X_i\}_{i=1}^\infty$ be a sequence of sets. For two elements $x = (x_1, x_2, \dots)$ and $y = (y_1, y_2, \dots)$ in the product $\prod_{i=1}^\infty X_i$ we say that $x \sim_\omega y$ if $\{i : x_i = y_i\} \in \omega$. It is well known that $\sim_\omega$ is an equivalence relation. The set $\prod_\omega X_i := \left(\prod_{i=1}^\infty X_i\right)/\sim_\omega$ is called the *ultraproduct* of the sets X_i.

Let T be a compact Polish space and let $\{t_i\}_{i=1}^\infty$ be a sequence in T. The ultralimit $\lim_\omega t_i$ is the unique point t in T with the property that for every open set U containing t the set $\{i : t_i \in U\}$ is in ω. (This definition implies that $\lim_\omega t_i$ is always an element of the closure of the set $\{t_i : i \in \mathbb{N}\}$.) Let $\{f_i : X_i \to T\}_{i=1}^\infty$ be a sequence of functions. We define $f = \lim_\omega f_i$ as the function on $\prod_\omega X_i$ whose value on the equivalence class of $(x_1, x_2, \dots)$ is $\lim_\omega f_i(x_i)$.

Let $\{X_i, \mu_i\}_{i=1}^\infty$ be pairs where X_i is a compact Polish space and μ_i is a probability measure on the Borel sets of X_i. We denote by $\mathbf{X}$ the ultraproduct space $\prod_\omega X_i$. The space $\mathbf{X}$ has the following structures on it.

Strongly open sets: We call a subset of $\mathbf{X}$ strongly open if it is the ultraproduct of open sets $\{S_i \subset X_i\}_{i=1}^\infty$.

Open sets: We say that $S \subset \mathbf{X}$ is open if it is a countable union of strongly open sets. Open sets on $\mathbf{X}$ form a σ-topology. This is similar to a topology but it has the weaker axiom that only countable unions of open sets are required to be open.

Lemma 5.1. *$\mathbf{X}$ with the above σ-topology is countably compact. This means that if $\mathbf{X}$ is covered by countably many open sets, then there is a finite subsystem which covers $\mathbf{X}$.*

Proof. Since every open set is a countable union of strongly open sets it is clearly enough to prove the statement for covering systems of $\mathbf{X}$ with strongly open sets. Let $\{O_i\}_{i=1}^\infty$ be such a system. Now each O_i is the ultraproduct of open sets $\{W_{k,i} \subseteq X_k\}_{k=1}^\infty$. Let $W_k := \bigcup_i W_{k,i}$. We have that $\prod_\omega W_k \supseteq \bigcup_i O_i = \mathbf{X}$ and thus $\prod_\omega W_k = \mathbf{X}$. It follows that $K := \{k : W_k = X_k\}$ is in ω. For each $k \in K$ let $f(k)$ denote the largest natural number such that $\bigcup_{i=1}^{f(k)} W_{k,i} \neq X_k$ (if $W_{k,1} = X_k$, then $f(k)$ is defined to be 0). By compactness of X_k we have that $f(k)$ is finite. Let us construct a sequence $\{x_k\}_{k=1}^\infty$ such that $x_k \in X_k \setminus (\bigcup_{i=1}^{f(k)} W_{k,i})$ if $k \in K$ and x_k is arbitrary if $k \in \mathbb{N} \setminus K$. The equivalence class of $(x_1, x_2, \dots) \in \mathbf{X}$ is covered by some element O_t from the covering system. It follows that the set $K' := \{k : x_k \in W_{k,t}\}$ is in ω. This means that $f(k) < t$ holds for every $k \in K \cap K'$ and thus $\bigcup_{i=1}^t W_{k,i} = X_k$ holds for $k \in K' \cap K$. Since $K \cap K' \in \omega$ we have that $\bigcup_{i=1}^t O_i = \mathbf{X}$. □

Borel sets and measurable sets: A subset of $\mathbf{X}$ is called Borel if it is in the σ-algebra generated by strongly open sets. A subset of $\mathbf{X}$ is called measurable if it is in the the completion of the Borel σ-algebra.

Ultralimit measure: If $S \subseteq \mathbf{X}$ is a strongly open set of the form $S = \prod_\omega S_i$, then we define $\mu(S)$ as $\lim_\omega \mu_i(S_i)$. It is a classical fact that μ extends as a probability measure to the σ-algebra of all measurable sets on $\mathbf{X}$. If $\mathbf{X}$ is the ultraproduct of finite sets, then the statement can be found in [4]. (See Proposition 2.2.) The proof of the general case is not much different. A good exposition of the subject is Evan Warner's Ph.D. thesis [19] where the statement is discussed in its full generality.

Ultralimit functions: Let T be a compact Hausdorff topological space. Let $\{f_i : X_i \to T\}_{i=1}^{\infty}$ be a sequence of Borel measurable functions. We call functions of the form $f = \lim_\omega f_i$ ultralimit functions. It is easy to see that ultralimit functions can always be modified on a 0 measure set such that they become measurable in the Borel σ-algebra on $\mathbf{X}$. This means that ultralimit functions are automatically measurable in the completion of the Borel σ-algebra.

Measurable functions: It is an important fact (see proposition 5.1 in [4] and proposition 3.8 in [19]) that every bounded measurable function on $\mathbf{X}$ is almost everywhere equal to some ultralimit function $f = \lim_\omega f_i$.

Continuity: A function $f : \mathbf{X} \to T$ from $\mathbf{X}$ to a topological space T is called continuous if $f^{-1}(U)$ is open in $\mathbf{X}$ for every open set in T. It follows from Lemma 5.1 that the image $f(\mathbf{X})$ of a continuous function $f : \mathbf{X} \to T$ in T is countably compact with respect to the restriction of the topology of T to $f(\mathbf{X})$. If T is metrizable, then also $f(\mathbf{X})$ is metrizable and thus countably compactness implies compactness. We will need the next lemma.

Lemma 5.2. *A continuous function $f : \mathbf{X} \to \mathbb{R}^n$ is the ultralimit of uniformly bounded continuous functions $f_i : X_i \to \mathbb{R}^n$.*

Proof. Observe first that it is enough to prove the statement for functions of the form $f : \mathbf{X} \to \mathbb{R}$ and the general statement follows by coordinate wise application. We have that $f(\mathbf{X})$ is a compact subset in $\mathbb{R}$ and thus $f(\mathbf{X}) \in (-a, a)$ for some $a \in \mathbb{R}^+$. Let $\epsilon > 0$ be fixed and let $U_i = (-a-\epsilon+i\epsilon/2, -a+i\epsilon/2)$ for $i = 1, 2, \ldots, t = \lceil 4a/\epsilon \rceil$. It is clear that the intervals U_i cover $(-a, a)$. For each $1 \le i \le t$ we have that $f^{-1}(U_i)$ is the union of countably many strongly open sets $\{Q_{i,j}\}_{j=1}^{\infty}$. By $\bigcup_{i,j} Q_{i,j} = \bigcup_i f^{-1}(U_i) = \mathbf{X}$ and Lemma 5.1 we have that there is a finite subsystem $\{S_k\}_{k=1}^r$ of $\{Q_{i,j}\}_{i,j}$ which covers $\mathbf{X}$. Let us choose points $\{x_k \in S_k\}_{k=1}^r$. Let $\{S_{k,j} \in X_j\}_{j=1}^{\infty}$ be sequences of open sets for every $1 \le k \le r$ such that $\prod_\omega S_{k,j} = S_k$ holds. Using that $\{S_k\}_{k=1}^r$ covers $\mathbf{X}$ we have that $T := \{j : \bigcup_k S_{k,j} = X_j\}$ is in ω. For every $j \in T$ we can choose a partition of unity $\{\rho_{k,j} : X_j \to [0,1]\}_{k=1}^r$ subordinated to the covering $\{S_{k,j}\}_{k=1}^r$. The functions $\rho_{k,j}$ are continuous and their sum is the constant 1 function. Furthermore $\rho_{k,j}$ is supported on $S_{k,j}$. If $j \in \mathbb{N} \setminus T$ we define $\rho_{k,j}$ to be 0. Now let $f_j := \sum_{k=1}^r \rho_{k,j} f(x_k)$ for $j \in \mathbb{N}$. Let $f' := \lim_\omega f_j$ and $\rho_k := \lim_\omega \rho_{k,j}$. By the additivity of ultralimits we have that $f' = \sum_{k=1}^r \rho_k f(x_k)$ and that $\sum_{k=1}^r \rho_k = 1_{\mathbf{X}}$. Now let $x \in \mathbf{X}$ be arbitrary. We have that whenever a set $f^{-1}(U_i)$ contains x, then $|f(x) - f(x_i)| \le \epsilon$ because $f(U_i)$ has diameter at most ϵ. since $f'(x)$ is a convex combination of the values $\{f(x_i) : x \in U_i\}$ we have that $|f'(x) - f(x)| \le \epsilon$ holds everywhere.

Now for an arbitrary $\epsilon > 0$ we produced a sequence of continuous function $\{f_j\}_{i=1}^{\infty}$ such that $|\lim_\omega f_j - f|_{\sup} \le \epsilon$ and that $\|f_j\|_{\sup} \le \|f\|_{\sup}$ holds for every $j \in \mathbb{N}$. Now we produce sequences for $m \in \mathbb{N}$ recursively. If $m = 1$, then let $\{f_j^1\}_{j=1}^{\infty}$ be a sequence which satisfies the above conditions with $\epsilon = 1/2$ for f. In general if $\{f_j^{m-1}\}_{j=1}^{\infty}$ is already produced, then we produce a new sequence $\{f_j^m\}_{j=1}^{\infty}$ with $\epsilon = 1/2^m$ for the function $f - \sum_{l=1}^{m-1} \lim_\omega f_j^m$. Note that if $m > 1$, then we have that

$\|f_j^m\|_{\sup} \le 1/2^{m-1}$ holds for every j. It follows that $g_j := \sum_{m=1}^{\infty} f_j^m$ is continuous. It is also clear that $\lim_\omega g_j = f$. □

Lemma 5.3. *Let T be a compact Polish space. Then the ultralimit of continuous functions $\{f_i : X_i \to T\}$ is continuous.*

Proof. Let $f = \lim_\omega f_i$. Let U be an open set in T. We can choose a countable family of open sets $\{W_j\}_{j=1}^{\infty}$ such that $U = \bigcup_j W_j$ and $\bar{W}_j \subseteq U$. We claim that $f^{-1}(U) = \bigcup_j \prod_\omega f_i^{-1}(W_j)$. Let j be fixed and assume that $x \in \prod_\omega f_i^{-1}(W_j)$. It follows from the basic properties of ultralimits that $f(x) \in \bar{W}_j \subseteq U$. This implies that $f^{-1}(U)$ contains $\prod_\omega f_i^{-1}(W_j)$ for every j. To see the other containment of the claim let $x \in f^{-1}(U)$. We have that there is j such that $f(x) \in U_j$. Assume that x is the equivalence class of $(x_1, x_2, \dots)$. We have by the properties of ultralimits that $\{i : f_i(x_i) \in U_j\}$ has to be in ω and thus $x \in \prod_\omega f_i^{-1}(W_j)$. □

6. The Fourier σ-algebra

If A is a compact abelian group, then linear characters are continuous homomorphisms of the form $\chi : A \to \mathcal{C}$ where $\mathcal{C}$ is the complex unit circle with multiplication as the group operation. Note that on compact abelian groups we typically use $+$ as the group operation. However, if we think of $\mathcal{C}$ as a subset of $\mathbb{C}$, then we are forced to use multiplicativ notation. On the other hand, if we think of $\mathbb{C}$ as the group $\mathbb{R}/\mathbb{Z}$, then we are basically forced to use additive notation.

Linear characters are forming the Fourier basis in $L^2(A)$. In particular, linear characters generate the whole Borel σ-algebra on A. Assume now that $\mathbf{A} = \prod_\omega A_i$ is the ultraproduct of compact abelian groups. Linear characters of $\mathbf{A}$ can be similarly defined as for compact abelian groups. In this case we require them to be continuous in the σ-topology on $\mathbf{A}$.

Proposition 6.1. *A function $\chi \in L^\infty(\mathbf{A})$ is a linear character if and only if $\chi = \lim_\omega \chi_i$ for some sequence $\{\chi_i \in L^\infty(A_i)\}_{i=1}^{\infty}$ of linear characters.*

The proof of the proposition relies on a rigidity result saying that almost linear characters on compact groups can be corrected to proper characters. For a function f we denote by f^* the pointwise complex conjugate of f.

Lemma 6.1. *For every $\epsilon > 0$ there is $\delta > 0$ such that if $f : A \to \mathbb{C}$ is a continuous function on a compact abelian group A with the property that $|f(x+a)f^*(x) - f(y+a)f^*(y)| \le \delta$, $||f(x)| - 1| \le \delta$ for every $x, y, a \in A$ and $|f(0) - 1| \le \delta$, then there is a character χ of A such that $|\chi(x) - f(x)| \le \epsilon$ holds for every $x \in A$.*

Proof. As a tool we introduce group theoretic expected values of random variables taking values in $\mathcal{C}$. Let l denote the arc length metric on the circle group $\mathcal{C} \simeq \mathbb{R}/\mathbb{Z}$ normalized by the total length 2π. It is clear that the metric l is topologically equivalent with the complex metric $|x - y|$ on $\mathcal{C}$. Assume that a random variable X takes its values in an arc of the circle group of length $1/3$. Then there is a lift Y of X to $\mathbb{R}$ such that $Y + \mathbb{Z} = X$ and Y takes its values in an interval of length $1/3$. The lift Y with this property is unique up to an integer shift. Then we define $\mathbb{E}(X) \in \mathbb{R}/\mathbb{Z}$ as $\mathbb{E}(Y) + \mathbb{Z}$. Switching to multiplicative notation in $\mathcal{C}$ this expected value satisfies $\mathbb{E}(X_1X_2) = \mathbb{E}(X_1)\mathbb{E}(X_2)$ where X_1, X_2 take values in an arc of length $1/6$.

Let us define $f_2(x) = f(x)/|f(x)|$. If $\delta < 1$, then $f(x) \neq 0$ on A and thus f_2 is defined on A. If $\delta > 0$ is small enough, then for every fixed t the function $x \mapsto f(x+t)f^*(x)$ takes values in an arc of length at most $1/6$. For every $t \in A$ let $g(t) = \mathbb{E}_x(f(x+t)f^*(x))$ where $\mathbb{E}$ is the group theoretic expected value. If δ is small enough, then $|g(t) - f(t)| \leq \epsilon$ holds for every $t \in A$ because $|f(x+t)f^*(x) - f(t)f^*(0)| \leq \delta$ and $f(0)$ is close to 1. Using our multiplicativity property of $\mathbb{E}$ we have for every pair $a, b \in A$ that

$$\begin{aligned} g(a+b)g^*(b) &= \mathbb{E}_x(f(x+a+b)f^*(x)f^*(x+b)f(x)) \\ &= \mathbb{E}_x(f(x+a+b)f^*(x+b)) = \mathbb{E}_x((x+a)f^*(x)) = g(a). \end{aligned}$$

This implies that g is a linear character of A. □

Now we are ready to prove Proposition 6.1

Proof. The continuity of χ guarantees that $\chi = \lim_\omega f_i$ for some sequence of continuous functions f_i on A_i (see Lemma 5.2). The fact that χ is a character implies that there is a sequence δ_i such that f_i satisfies the conditions of Lemma 6.1 with δ_i for every i and $\lim_\omega \delta_i = 0$. It follows by Lemma 6.1 that there is a sequence of linear characters χ_i on A_i such that $\lim_\omega \max(|\chi_i - f_i|) = 0$. Thus we have that $\lim_\omega \chi_i = \lim_\omega f_i = \chi$. □

Proposition 6.1 implies that the set of linear characters of $\mathbf{A}$ (also as a group) is equal to $\prod_\omega \hat{A}_i$. We denote this set by $\hat{\mathbf{A}}$. If $f \in L^2(\mathbf{A})$, then the Fourier transform of f on $\mathbf{A}$ is the function $\hat{f} \in l^2(\hat{\mathbf{A}})$ defined by $\hat{f}(\chi) = (f, \chi)$. If $f = \lim_\omega f_i$, then we have that $\hat{f} = \lim_\omega \hat{f}_i$.

It was observed in [14] that linear characters of $\mathbf{A}$ no longer span $L^2(\mathbf{A})$. This shows that in general we only have $\|\hat{f}\|_2 \leq \|f\|_2$ instead of equality. Furthermore, the σ-algebra $\mathcal{F}(\mathbf{A})$ generated by linear characters on $\mathbf{A}$ is smaller than the whole ultraproduct σ-algebra on $\mathbf{A}$. (The only exception is the case when $\mathbf{A}$ is a finite group. This can happen if the groups A_i are finite and there is a uniform bound on their size.)

We call $\mathcal{F}(\mathbf{A})$ the **Fourier σ-algebra** on $\mathbf{A}$. The fact that the Fourier σ-algebra is not the complete σ-algebra on $\mathbf{A}$ gives rise to the interesting operation $f \mapsto \mathbb{E}(f|\mathcal{F}(\mathbf{A}))$ that isolates the "Fourier part" of a function $f \in L^2(\mathbf{A})$. Using that linear characters of $\mathbf{A}$ are closed with respect to multiplication we obtain that linear characters are forming a basis in $L^2(\mathcal{F}(\mathbf{A}))$. This implies that if $f \in L^2(\mathbf{A})$, then $\hat{f} = \hat{g}$ where $g = \mathbb{E}(f|\mathcal{F}(\mathbf{A}))$. Thus we have that $\|\hat{f}\|_2 = \|\hat{g}\|_2 = \|\mathbb{E}(f|\mathcal{F}(\mathbf{A}))\|_2$. In particular, $\|f\|_2 = \|\hat{f}\|_2$ holds if and only if f is measurable in $\mathcal{F}(\mathbf{A})$.

The Fourier σ-algebra has an elegant description in terms of the second Gowers norm U_2. Recall that the U_2-norm [5], [6] of a function $f \in L^\infty(A)$ on a compact abelian group A is defined by

$$\|f\|_{U_2} = \Big(\mathbb{E}_{x,a,b \in A} f(x)f^*(x+a)f^*(x+b)f(x+a+b)\Big)^{1/4}. \tag{2}$$

The next lemma gives a description of the U_2-norm in terms of Fourier analysis.

Lemma 6.2. *If $f \in L^\infty(A)$, then $\|f\|_{U_2} = \|\hat{f}\|_4$ and thus $\|\hat{f}\|_\infty \leq \|f\|_{U_2} \leq (\|f\|_2\|\hat{f}\|_\infty)^{1/2}$.*

One can define $\|f\|_{U_2}$ by the formula (2) for functions on ultraproduct groups. With this definition we have that $\|f\|_{U_2} = \lim_\omega \|f_i\|_{U_2}$ whenever $f = \lim_\omega f_i$. The main difference from the compact case is that $\|.\|_{U_2}$ is no longer a norm for functions in $L^\infty(\mathbf{A})$. It is only a semi-norm. However, the next lemma shows that $\|.\|_{U_2}$ is a norm when restricted to $L^\infty(\mathcal{F}(\mathbf{A}))$ and that $\mathcal{F}(\mathbf{A})$ is the largest σ-algebra with this property.

Lemma 6.3. *If $g \in L^\infty(\mathbf{A})$, then $\|g\|_{U_2} = 0$ if and only if g is orthogonal to $L^2(\mathcal{F}(\mathbf{A}))$. A function $f \in L^\infty(\mathbf{A})$ is measurable in $\mathcal{F}(\mathbf{A})$ if and only if f is orthogonal to every function $g \in L^\infty(\mathbf{A})$ with $\|g\|_{U_2} = 0$. In particular, we have that $\|.\|_{U_2}$ is a norm on $L^\infty(\mathcal{F}(\mathbf{A}))$.*

Proof. We can assume that $g = \lim_\omega g_i$ for some sequence of functions $\{g_i \in L^\infty(A_i)\}_{i=1}^\infty$ such that $\|g_i\|_\infty \leq \|g\|_\infty$ holds for every i. Assume first that $\|g\|_{U_2} = 0$. Let $\chi = \lim_\omega \chi_i$ be an ultralimit of linear characters. Using Lemma 6.2 we have that $|(g_i, \chi_i)| \leq \|\hat{g}_i\|_\infty \leq \|g_i\|_{U_2}$ and thus

$$|(g, \chi)| = \lim_\omega |(g_i, \chi_i)| \leq \lim_\omega \|g_i\|_{U_2} = \|g\|_{U_2} = 0.$$

It follows that g is orthogonal to the space $L^2(\mathcal{F}(\mathbf{A}))$ spanned by linear characters of $\mathbf{A}$. For the other direction assume that $g \neq 0$ is orthogonal $L^2(\mathcal{F}(\mathbf{A}))$. For every i we choose a linear character χ_i on A_i such that $|(g_i, \chi_i)| = \|\hat{g}_i\|_\infty$. We have by Lemma 6.2 and by $\|g_i\|_2 \leq \|g_i\|_\infty \leq \|g\|_\infty$ that $|(g_i, \chi_i)| \geq \|g_i\|_{U_2}^2 \|g\|_\infty^{-1}$. Then we have for $\chi = \lim_\omega \chi_i$ that $0 = |(g, \chi)| \geq (\lim_\omega \|g_i\|_{U_2}^2)\|g\|_\infty^{-1}$. It follows that $\|g\|_{U_2} = 0$.

Now we prove the second part of the statement. If $f \in L^\infty(\mathcal{F}(A))$, then by the first part of the statement f has to be orthogonal to every $g \in L^\infty(\mathbf{A})$ with $\|g\|_{U_2} = 0$. For the other direction assume that $f \in L^\infty(\mathbf{A})$ is orthogonal to every $g \in L^\infty(\mathbf{A})$ with $\|g\|_{U_2} = 0$. Let $h := f - \mathbb{E}(f|\mathcal{F}(\mathbf{A})) \in L^\infty(\mathbf{A})$. Note that since $\mathbb{E}$ is an orthogonal projection it follows that $(f, h) = \|h\|_2^2$. From

$$\mathbb{E}(h|\mathcal{F}(\mathbf{A})) = \mathbb{E}(f - \mathbb{E}(f|\mathcal{F}(\mathbf{A}))|\mathcal{F}(\mathbf{A})) = \mathbb{E}(f|\mathcal{F}(\mathbf{A})) - \mathbb{E}(f|\mathcal{F}(\mathbf{A})) = 0$$

we have that h is orthogonal to the whole space $L^2(\mathcal{F}(\mathbf{A}))$ and so by the first statement it follows that $\|h\|_{U_2} = 0$. It implies by our assumption on f that $(f, h) = 0$ and thus $\|h\|_2^2 = 0$. Now we have that $h = 0$ and $f = \mathbb{E}(f|\mathcal{F}(\mathbf{A}))$ is measurable in $\mathcal{F}(\mathbf{A})$. □

Let $\hat{\mathcal{Q}} : L^2(\mathbf{A}) \to \mathcal{M}$ be such that $\hat{\mathcal{Q}}(f)$ is the isomorphism class of $\hat{f}$ in $\mathcal{M}$. Let furthermore $\mathcal{Q}(f)$ denote the isomorphism class in $\mathcal{H}$ representing the Fourier transform of $\hat{\mathcal{Q}}(f)$. Note that $\mathcal{Q}(f) = \mathcal{Q}(\mathbb{E}(f|\mathcal{F}(\mathbf{A})))$. We have that $\mathcal{Q}(f)$ can be represented as a measurable function on some second countable compact abelian group with $\|\mathcal{Q}(f)\|_2 \leq \|f\|_2$ which in some sense imitates f. However, it is not even clear from this definition that if f is a bounded function, then $\mathcal{Q}(f)$ is also bounded. The next theorem provides a structure theorem for functions in $L^\infty(\mathcal{F}(\mathbf{A}))$ and describes $\mathcal{Q}(f)$.

Theorem 6. *A function $f \in L^\infty(\mathbf{A})$ is measurable in $\mathcal{F}(\mathbf{A})$ if and only if there is a continuous, surjective, measure preserving homomorphism $\phi : \mathbf{A} \to A$ to some second countable compact abelian group A and a function $h \in L^\infty(A)$ such that $f = h \circ \phi$ (up to 0 measure change). Furthermore, $d(h, \mathcal{Q}(f)) = 0$ implying that the isomorphism class of h is $\mathcal{Q}(f)$.*

Proof. Assume first that $f = h \circ \phi$ for some homomorphism ϕ and function h as in the statement. Let $h = \sum_{i=1}^\infty \lambda_i \chi_i$ be the Fourier decomposition of h converging in $L^2(A)$ where χ_i is a sequence of linear characters of A. We have that $\chi_i \circ \phi$ is a linear character of $\mathbf{A}$ for every i. The measure preserving property of ϕ implies that $f = \sum_{i=1}^{\infty} \lambda_i(\chi_i \circ \phi)$ and thus f is measurable in $\mathcal{F}(\mathbf{A})$.

For the other direction assume that $f \in L^\infty(\mathcal{F}(\mathbf{A}))$. Then $f = \sum_{i=1}^{\infty} a_i \chi_i$ for some (distinct) linear characters $\{\chi_i\}_{i=1}^{\infty}$ of $\mathbf{A}$ where the convergence is in $L^2(\mathbf{A})$ and $\|f\|_2^2 = \sum_{i=1}^{\infty} |a_i|^2$. Let us consider the homomorphism $\phi : \mathbf{A} \to \mathcal{C}^\infty$ such that the i-th coordinate of $\phi(x) = \chi_i(x)$. (Recall that $\mathcal{C}$ is the group $\mathbb{R}/\mathbb{Z}$ or equivalently the complex unit circle with respect to multiplication. The group $\mathcal{C}^\infty$ is a compact abelian group.) Using the continuity of ϕ we have that the image A of ϕ is a closed subgroup in $\mathcal{C}^\infty$. Let ν denote the Borel measure on A satisfying $\nu(S) = \mu(\phi^{-1}(S))$ where μ is the ultralimit measure on $\mathbf{A}$. The fact that ϕ is a homomorphism implies that ν is a group invariant Borel probability measure on A and thus ν is equal to the normalized Haar measure. In other words, ϕ is measure preserving with respect to the Haar measure on A.

Let us denote by α_i the i-th coordinate function on A. It is clear that $\{\alpha_i\}_{i=1}^{\infty}$ is a system of linear characters of A. Since ϕ is surjective it induces an injective homomorphism $\hat{\phi} : \hat{A} \to \hat{\mathbf{A}}$ defined by $\hat{\phi}(\chi) = \chi \circ \phi$ with the property that $\hat{\phi}(\alpha_i) = \chi_i$ holds for every i. We have that $h = \sum_{i=1}^{\infty} a_i \alpha_i$ (which is defined up to a 0 measure set on A) is convergent in L^2 and has the property that $f = h \circ \phi$ (up to a 0 measure set). The fact that $\hat{\phi}$ is an injective homomorphism implies that $\hat{d}(\hat{h}, \hat{f}) = 0$ and thus $d(h, \mathcal{Q}(f)) = 0$. □

If $\mathcal{L}$ is a system of linear forms and $f \in L^\infty(\mathbf{A})$, then we can define $t(\mathcal{L}, f)$ by the formula (1) using the ultralimit measure on $\mathbf{A}$.

Proposition 6.2. *Let $f \in L^\infty(\mathcal{F}(\mathbf{A}))$ and let $\mathcal{L}$ be a system of linear forms. Then $t(\mathcal{L}, f) = t(\mathcal{L}, \mathcal{Q}(f))$. Furthermore, if $\mathcal{L}$ has complexity 1 in a family $\mathcal{A}$ of compact abelian groups, $\mathbf{A}$ is an ultraproduct of groups in $\mathcal{A}$ and $f \in L^\infty(\mathbf{A})$, then $t(\mathcal{L}, f) = t(\mathcal{L}, \mathcal{Q}(f))$.*

Proof. For the first part we use Theorem 6. We get that $f = h \circ \phi$ for some measure preserving homomorphsim $\phi : \mathbf{A} \to A$. It follows that $t(\mathcal{L}, f) = t(\mathcal{L}, h) = t(\mathcal{L}, \mathcal{Q}(f))$. For the second part let $f = \lim_\omega f_i$ and $g = \mathbb{E}(f|\mathcal{F}(\mathbf{A})) = \lim_\omega g_i$ for some functions with $\|f_i\|_\infty \leq \|f\|_\infty$ and $\|g_i\|_\infty \leq \|g\|_\infty$. We have that

$$\lim_\omega \|f_i - g_i\|_{U_2} = \|f - g\|_{U_2} = 0.$$

Then using that $\mathcal{L}$ has complexity 1 we obtain $t(\mathcal{L}, \mathcal{Q}(f)) = t(\mathcal{L}, \mathcal{Q}(g)) = t(\mathcal{L}, g) = \lim_\omega t(\mathcal{L}, g_i) = \lim_\omega t(\mathcal{L}, f_i) = t(\mathcal{L}, f)$. □

7. The ultraproduct descriptions of $\hat{d}$ and d convergence

We give a simple and useful description of $\hat{d}$ convergence using ultrafilters. The price that we pay for the simplicity is that we don't get an explicit metric on $\mathcal{M}$, we only get the concept of convergence.

Theorem 7. *Let $a > 0$. Assume that $\{f_i\}_{i=1}^{\infty}$ is a sequence in $\mathcal{M}_a$ that converges to f in $\hat{d}$; then f is isomorphic to $\lim_\omega f_i$ for every (non-principal) ultrafilter ω. Consequently a sequence $\{f_i\}_{i=1}^{\infty}$ in $\mathcal{M}_a$ is convergent in $\hat{d}$ if and only if the isomorphism class of $\lim_\omega f_i$-limit doesn't depend on the choice of the ultrafilter ω.*

Proof. For every i let $\alpha_i : T_i \to S_i$ be an ϵ_i-isomorphism between f_i and f with $T_i \subseteq G_i, S_i \subseteq G$ such that $\lim_{i\to\infty} \epsilon_i = 0$. Assume that $\{h_i\}_{i=1}^{\infty}$ represents an element h in $\prod_\omega G_i$ that is in $\mathrm{supp}(g)$ where $g = \lim_\omega f_i$. We have for some set $S \in \omega$ that $|f_i(h_i)| > |g(h)|/2$ and $\epsilon_i \leq |g(h)|/4$ for $i \in S$. It follows that $\alpha_i(h_i) \in \mathrm{supp}_{|g(h)|/4}(f)$ holds for every $i \in S$. Since $\mathrm{supp}_{|g(h)|/4}$ is finite we have that $\lim_\omega \alpha_i(h_i)$ exists and it is an element in G that we denote by $\beta(h)$. The map $\beta : \mathrm{supp}(g) \to \mathrm{supp}(f)$ is a partial isomorphism of arbitrary high weight and so it extends to an isomorphism from $\langle g\rangle$ to $\langle f\rangle$. It is clear that β is also an isomorphism between f and g. $\square$

Corollary 7.1. *Let $a > 0$. Assume that $\{f_i\}_{i=1}^{\infty}$ is a sequence of functions with $f_i \in L^\infty(A_i)$ and $\|f_i\|_\infty \leq a$ for some sequence $\{A_i\}_{i=1}^{\infty}$ of compact abelian groups. If $\{f_i\}_{i=1}^{\infty}$ converges to $f \in \mathcal{H}_a$ in the metric d, then $f = \mathcal{Q}(\lim_\omega f_i)$ for an arbitrary (non-principal) ultrafilter ω.*

Proof. Since the Fourier transform of $f' = \lim_\omega f_i$ is the ultralimit of the Fourier transforms of f_i we have by Theorem 7 that $\hat{d}(\hat{f}', \hat{f}) = 0$. It follows that $\mathcal{Q}(f') = f$. $\square$

Corollary 7.2. *Let $a > 0$. Assume that $\{f_i\}_{i=1}^{\infty}$ is a convergent sequence of functions with $f_i \in L^\infty(A_i)$ and $\|f_i\|_\infty \leq a$ for some sequence $\{A_i\}_{i=1}^{\infty}$ of compact abelian groups. Then the limit f of $\{f_i\}_{i=1}^{\infty}$ can be represented as a function on some compact abelian group A such that the dual group of A is a subgroup in $\prod_\omega \hat{A}_i$.*

Proof. We have by Corollary 7.1 that $f = \mathcal{Q}(\lim_\omega f_i)$. This means that $\hat{f}$ has an injective embedding into $\hat{\mathbf{A}}$ where $\mathbf{A} = \prod_\omega A_i$. By $\hat{\mathbf{A}} = \prod_\omega \hat{A}_i$ the proof is complete. $\square$

Corollary 7.2 gives a useful restriction on the structure of the group on which the limit function of a convergent seqence is defined. For example if A_i are growing groups of prime order, then the limit function is defined on a compact group whose dual group is torsion-free. On the other hand, if p is a fix prime and f_i is defined on $\mathbb{Z}_p^i$, then the limit function is defined on the compact group $\mathbb{Z}_p^\infty$.

8. Proofs of Theorems 3, 4, 5

For the proofs of Theorem 3 and Theorem 4 assume that $\{f_i\}_{i=1}^{\infty}$ is a convergent sequence in $\mathcal{H}(K)$ for some convex compact set $K \subseteq \mathbb{C}$. Corollary 7.1 implies that the limit is $\mathcal{Q}(f)$ where $f = \lim_\omega f_i$. Note that f takes its values in K. We have that $\mathcal{Q}(f) = \mathcal{Q}(g)$ where $g = \mathbb{E}(f|\mathcal{F}(\mathbf{A}))$. It follows by Theorem 6 that $g = h \circ \phi$ for some measure preserving homomorphism $\phi : \mathbf{A} \to A$ and the isomorphism class of h is $\mathcal{Q}(g)$. Since g is a projection of f to a σ-algebra we have that g (and thus h) takes its values in K. This completes the proof of Theorem 3.

For the proof of Theorem 4 assume that f_i is tightly convergent and $K = \{x : x \in \mathbb{C}, \|x\| \leq a\}$. Then, using the above notation we have that $\|g\|_2 = \|h\|_2 = \lim_{i\to\infty} \|f_i\|_2 = \lim_\omega \|f_i\|_2 = \|f\|_2$ where we use tightness in the second equality. This is only possible if $f = g$ and thus $\mu_h = \mu_f = \lim_\omega \mu_{f_i}$ holds. Since this is true for every ultrafilter ω we obtain that $\lim_{i\to\infty} \mu_{f_i} = \mu_h$ holds with respect to weak convergence of measures.

To prove Theorem 5 assume that $\mathcal{L}$ has complexity 1 and f_i is a d convergent sequence as above. Using the above notation and Proposition 6.2 we have that

$\lim_\omega t(\mathcal{L}, f_i) = t(\mathcal{L}, f) = t(\mathcal{L}, \mathcal{Q}(f))$ where (using Corollary 7.1) $\mathcal{Q}(f)$ is equal to the d-limit of the sequence $\{f_i\}_{i=1}^\infty$. Since this is true for every ultrafilter ω the proof is complete.

9. Proof of Theorem 2

For the proof of Theorem 2 we will need the next proposition which is interesting in its own right.

Proposition 9.1. *Let B be a compact abelian group with torsion-free dual group and let $f : B \to [0,1]$ be an arbitrary measurable function. Then there are subsets $S_p \subseteq \mathbb{Z}_p$ for every prime number p such that the functions 1_{S_p} converge to f.*

Lemma 9.1. *For every $\epsilon > 0$ there is $N(\epsilon)$ such that if A is a finite abelian group with $|A| \geq N(\epsilon)$ and $f : A \to [0,1]$ is a function, then there is a function $h : A \to \{0,1\}$ such that $\|f - h\|_{U_2} \leq \epsilon$.*

Proof. Let us fix $\epsilon > 0$. Let $f : A \to [0,1]$ be a function on a finite abelian group. Let h be the random function on A whose distribution is uniquely determined by the following properties: 1.) h is $\{0,1\}$-valued, 2.) $\{h(a) \mid a \in A\}$ is an independent system of random variables, and 3.) $\mathbb{E}(h(a)) = f(a)$ holds for every $a \in A$. We claim that with a large probability the function $h - f$ has U_2-norm at most ϵ if $|A|$ is large enough. Obsereve that $X_a := h(a) - f(a)$ is a random variable for each $a \in A$ with 0 expectation and $\|X_a\|_\infty \leq 1$. The random variables X_a are all independent. Let $\chi : A \to \mathbb{C}$ be a linear character. Then we have that $(h - f, \chi) = |A|^{-1} \sum_{a \in A} X_a \chi(a)$. By Chernoff's bound we have that $\mathbb{P}(|(h - f, \chi)| \geq \epsilon^2)$ is exponentially small in $|A|$. This implies that if $|A|$ is large enough, then with probability close to 1 we have that $\|\hat{h} - \hat{g}\|_\infty \leq \epsilon^2$ and thus by Lemma 6.2 we get $\|h - g\|_{U_2} \leq \epsilon$ holds in these cases. ☐

Proof of Proposition 9.1. For a number n let $a(n)$ denote the minimum of $d(1_S, f)$ where S is a subset in $\mathbb{Z}_n$. The statement of the proposition is equivalent with $\lim_{p\to\infty} a(n) = 0$ where p runs through the prime numbers. Assume by contradiction that there is $\epsilon > 0$ and a growing infinite sequence $\{p_i\}_{i=1}^\infty$ of prime numbers with $a(p_i) > \epsilon$. Let $A_i = \mathbb{Z}_{p_i}$ and $\mathbf{A} = \prod_\omega A_i$. We have that $\hat{\mathbf{A}} = \prod_\omega \hat{A}_i \simeq \prod_\omega A_i = \mathbf{A}$. Since $\mathbf{A}$ is not only an abelian group but a field of 0 characteristic with uncountably many elements we have that $\mathbf{A}$ (and thus $\hat{\mathbf{A}}$) as an abelian group is isomorphic to an infinite direct sum of $\mathbb{Q}^+$. It follows that the torsion-free group $\hat{B}$ has an embedding $\hat{\phi} : \hat{B} \to \hat{\mathbf{A}}$ into $\hat{\mathbf{A}}$. This embedding induces a continuous homomorphsim $\phi : \mathbf{A} \to B$ in the way that $\phi(x)$ denotes the unique element in B such that $\chi(\phi(x)) = \hat{\phi}(\chi)(x)$ holds for every $\chi \in \hat{B}$.

Let $g = f \circ \phi$. We have that $g : \mathbf{A} \to [0,1]$ is a measurable function and thus $g = \lim_\omega g_i$ for a system of functions $\{g_i : A_i \to [0,1]\}_{i=1}^\infty$. By Lemma 9.1 for every i we can find a $0-1$ valued function g_i' such that $\lim_{i\to\infty} \|g_i' - g_i\|_{U_2} = 0$. By choosing a subsequence we can assume that both $\{g_i'\}_{i=1}^\infty$ and $\{g_i\}_{i=1}^\infty$ are d-convergent. Let $g' = \lim_\omega g_i'$. We have that $\|g - g'\|_{U_2} = 0$ and thus since g is measurable in $\mathcal{F}(\mathbf{A})$ we have that $g = \mathbb{E}(g'|\mathcal{F}(\mathbf{A}))$. By Corollary 7.1 we obtain that the d limit of $\{g_i'\}_{i=1}^\infty$ is f. This implies that $0 = \lim d(g_i', f) \geq \liminf a(p_i) \geq \epsilon$ which is a contradiction. ☐

Now we are ready to prove Theorem 2. First observe that in Proposition 9.1 we can assume with no additional cost that the sets S_p have density at least $\mathbb{E}(f)$.

This follows from the fact that their densities converge to $\mathbb{E}(f)$ and so it is enough to set a few values to 1 (with density tending to 0). This observation together with Proposition 9.1 and Theorem 1 imply that if $f : A \to [0,1]$ is a measurable function with $\mathbb{E}(f) = \delta$ on an abelian group with torsion-free dual, then $\rho(\delta, \mathcal{L}) \leq t(\mathcal{L}, f)$. It remains to find a function where equality holds. For every p prime let $S_p \subseteq \mathbb{Z}_p$ be such that $|S_p|/p \geq \delta$ and that $t(\mathcal{L}, 1_{S_p})$ is minimal possible. We can choose a d-convergent subsequence $\{f_i\}_{i=1}^{\infty}$ from 1_{S_p} such that $\lim_{i\to\infty} t(\mathcal{L}, f_i) = \rho(\delta, \mathcal{L})$. Let f be the limit of $\{f_i\}_{i=1}^{\infty}$. By Theorem 1 we have that $t(\mathcal{L}, f) = \lim_{i\to\infty} t(\mathcal{L}, f_i) = \rho(\delta, \mathcal{L})$. Corollary 7.2 guarantees that f is defined on a group whose dual is torsion-free.

10. Connection to dense graph limit theory and concluding remarks

Let H and G be finite graphs. The density of H in G is the probability that a random map from $V(H)$ to $V(G)$ takes edges to edges. We denote this quantity by $t(H, G)$. One can generalize this notion of density for the case when G is replaced by a symmetric bounded measurable function $W : \Omega^2 \to \mathbb{C}$ where (Ω, μ) is a probability space. Then $t(H, W)$ is defined by

$$t(H, W) := \int_{x_1, x_2, \ldots, x_n \in \Omega} \prod_{(i,j)\in E(H)} W(x_i, x_j) \, d\mu^n,$$

where the vertices of H are indexed by $\{1, 2, \ldots, n\}$. It is easy to check that if $\Omega = V(G)$, μ is the uniform distribution on $V(G)$ and $W : V(G)^2 \to \{0,1\}$ is the adjacency matrix of G, then $t(H, G) = t(H, W)$.

In the framework of dense graph limit theory, a sequence of graphs $\{G_i\}_{i=1}^{\infty}$ is called convergent if for every fixed graph H the sequence $\{t(H, G_i)\}_{i=1}^{\infty}$ is convergent. It was proved in [10] that for a convergent graph sequence $\{G_i\}_{i=1}^{\infty}$ there is a limit object of the form of a symmetric measurable function $W : \Omega^2 \to [0,1]$ (called a *graphon*) such that for every graph H we have $\lim_{i\to\infty} t(H, G_i) = t(H, W)$.

In the above theorem Ω can be chosen to be $[0,1]$ with the uniform measure, however, in many cases it is more natural to use other probability spaces. We investigate the case when (Ω, μ) is a compact abelian group A with the normalized Haar measure. Let $f : A \to \mathbb{C}$ be a bounded measurable function and let $W_f : A^2 \to \mathbb{C}$ be defined by $W_f(x, y) := f(x + y)$. As it was pointed out in the introduction, for a finite graph H the density $t(H, W_f)$ is equal to $t(\mathcal{L}, f)$ where $\mathcal{L}_H := \{x_i + x_j \ : \ (i,j) \in E(H)\}$. Using this correspondence and our results in this paper we get the following theorem on graph limits.

Theorem 8. *Let $\{f_i : A_i \to K\}_{i=1}^{\infty}$ be a sequence of measurable functions on compact abelian groups with values in a compact convex set $K \subseteq \mathbb{C}$. Assume that* $\lim_{i\to\infty} t(H, W_{f_i})$ *exists for every graph H. Then there is a measurable function $f : A \to K$ on a compact abelian group A such that* $\lim_{i\to\infty} t(H, W_{f_i}) = t(H, W_f)$ *holds for every graph H.*

Proof. By choosing a subsequence we can assume by Theorem 3 that $\{f_i\}_{i=1}^{\infty}$ is convergent in d with limit $f : A \to K$. Then by Theorem 1 we obtain that $\lim_{i\to\infty} t(\mathcal{L}_H, f_i) = t(\mathcal{L}_H, f)$ holds for every graph H. This completes the proof. $\square$

Theorem 8 is closely related to the results in [12]. Let $f : G \to [0,1]$ be a measurable function on a compact but not necessarily commutative group. Assume that the technical condition $f(g) = f(g^{-1})$ holds for every $g \in G$. Then the function $W : G^2 \to [0,1]$ defined by $W(x,y) = f(xy^{-1})$ is symmetric. We call graphons of this type Cayley graphons. It was proved in [12] that limits of Cayley graphons are also Cayley graphons. This theorem implies that one can talk about limits of functions on compact topological groups and the limit objects are also functions on compact topological groups. Another direction of generalization in the commutative setting is when we consider densities of linear configurations of higher complexity. As it was shown in [13], this refinement of the limit concept leads to more complicated limit objects that are measurable functions on nilmanifolds and nilspaces.

As we promised in the introduction we finish the paper by showing that Theorem 2 implies the qualitative version of Roth's theorem. Assume by contradiction that $\rho(\delta,\mathcal{L}) = 0$ holds for some $\delta > 0$. Then there is a function $f : A \to [0,1]$ such that $t(\mathcal{L}, f) = 0$ with $\mathbb{E}(f) = \delta$. It is easy to see that if S is the support of f, then $t(\mathcal{L}, 1_S) = 0$ also holds and $\mathbb{E}(1_S) \geq \delta$. Since A is the inverse limit of finite dimensional torus groups we have that there is a factor map $\tau : A \to \mathbb{T}_n$ to a finite dimensional torus such that $\mathbb{E}(1_S|\tau) > 3/4$ holds on a positive measure set $\tau^{-1}(Q)$ where $Q \subseteq \mathbb{T}_n$ is Borel measurable. We have that

$$0 = t(\mathcal{L}, 1_S) \geq t(\mathcal{L}, 1_S 1_{\tau^{-1}(Q)}) \geq t(\mathcal{L}, 1_{\tau^{-1}(Q)})/4 = t(\mathcal{L}, 1_Q)/4,$$

where the only nontrivial inequality is the second one. To see this observe that for almost every 3-term arithmetic progression inside $\tau^{-1}(Q)$ a random translate with some element from $\ker(\tau)$ is with probability at least $1/4$ inside $\tau^{-1}(Q) \cap S$. This is true because $\mathbb{E}(1_S 1_{\tau^{-1}(Q)}) > 3/4$ holds inside $\tau^{-1}(Q)$. It remains to show that on $\mathbb{T}_n$ there is no positive density set Q with 0 density copies of $\mathcal{L}$. By the Lebesgue density theorem we can find intervals $I_1, I_2, \dots, I_n \subseteq \mathbb{T}_1$ for every $\epsilon > 0$ such that Q intersects $C := \times_{i=1}^n I_i$ in a way that it has density at least $1-\epsilon$ in C. It is easy to see that if ϵ is small enough, then $C \cap Q$ must contain positive density copies of $\mathcal{L}$.

References

[1] Pablo Candela and Olof Sisask, *Convergence results for systems of linear forms on cyclic groups and periodic nilsequences*, SIAM J. Discrete Math. **28** (2014), no. 2, 786–810. MR3211031

[2] C. Borgs, J. T. Chayes, L. Lovász, V. T. Sós, and K. Vesztergombi, *Convergent sequences of dense graphs. I. Subgraph frequencies, metric properties and testing*, Adv. Math. **219** (2008), no. 6, 1801–1851. MR2455626

[3] F. R. K. Chung, R. L. Graham, and R. M. Wilson, *Quasi-random graphs*, Combinatorica **9** (1989), no. 4, 345–362. MR1054011

[4] Gábor Elek and Balázs Szegedy, *A measure-theoretic approach to the theory of dense hypergraphs*, Adv. Math. **231** (2012), no. 3-4, 1731–1772. MR2964622

[5] W. T. Gowers, *A new proof of Szemerédi's theorem*, Geom. Funct. Anal. **11** (2001), no. 3, 465–588. MR1844079

[6] W. T. Gowers, *Fourier analysis and Szemerédi's theorem*, Proceedings of the International Congress of Mathematicians, Vol. I (Berlin, 1998), Doc. Math. **Extra Vol. I** (1998), 617–629. MR1660658

[7] W. T. Gowers and J. Wolf, *The true complexity of a system of linear equations*, Proc. Lond. Math. Soc. (3) **100** (2010), no. 1, 155–176. MR2578471

[8] Benjamin Green and Terence Tao, *Linear equations in primes*, Ann. of Math. (2) **171** (2010), no. 3, 1753–1850. MR2680398

[9] László Lovász, *Large networks and graph limits*, American Mathematical Society Colloquium Publications, vol. 60, American Mathematical Society, Providence, RI, 2012. MR3012035

[10] László Lovász and Balázs Szegedy, *Limits of dense graph sequences*, J. Combin. Theory Ser. B **96** (2006), no. 6, 933–957. MR2274085

[11] László Lovász and Balázs Szegedy, *Szemerédi's lemma for the analyst*, Geom. Funct. Anal. **17** (2007), no. 1, 252–270. MR2306658

[12] László Lovász and Balázs Szegedy, *The automorphism group of a graphon*, J. Algebra **421** (2015), 136–166. MR3272377

[13] B. Szegedy, *On higher order Fourier analysis*, arxiv:1203.2260

[14] B. Szegedy, *Higher order Fourier analysis as an algebraic theory I.*, arxiv:0903.0897

[15] B. Szegedy, *Higher order Fourier analysis as an algebraic theory III.*, arxiv:1001.4282

[16] B. Szegedy, *Gowers norms, regularization and limits of functions on abelian groups*, arxiv:1010.6211

[17] Walter Rudin, *Fourier analysis on groups*, Wiley Classics Library, John Wiley & Sons, Inc., New York, 1990. Reprint of the 1962 original; A Wiley-Interscience Publication. MR1038803

[18] Terence Tao, *Higher order Fourier analysis*, Graduate Studies in Mathematics, vol. 142, American Mathematical Society, Providence, RI, 2012. MR2931680

[19] E. Warner, *Ultraproducts and the foundations of higher order Fourier analysis*, Ph.D. Thesis, 2012

Alfred Renyi Institute of Mathematics, HU 1053 Budapest, Reáltanoda utca 13-15, 1053 Hungary