

TRANSLATION INVARIANT EQUATIONS AND THE METHOD OF SANDERS

THOMAS F. BLOOM

ABSTRACT. We extend the recent improvement of Roth's theorem on three term arithmetic progressions by Sanders to obtain similar results for the problem of locating non-trivial solutions to translation invariant linear equations in many variables in both $\mathbb{Z}/N\mathbb{Z}$ and $\mathbb{F}_q[t]$.

1. INTRODUCTION

This paper concerns solutions to a linear equation in $s \geq 3$ variables,

$$(1.1) \quad c_1x_1 + \cdots + c_sx_s = 0,$$

working inside some fixed ring R , and we restrict our attention to those equations with coefficients satisfying $c_1 + \cdots + c_s = 0$. These are often referred to as translation invariant systems, since it follows that if $(x_1, \dots, x_s) \in R^s$ is a solution then so is $(x_1 + x, \dots, x_s + x)$ for any $x \in R$. Following Ruzsa [14] we define the *genus* of an equation of the shape (1.1) to be the largest number m with the property that there is a partition of $\{1, \dots, s\}$ into m disjoint nonempty sets T_j where $\sum_{i \in T_j} c_i = 0$. Note that the genus is well defined and positive, as a consequence of translation invariance. Given an equation (1.1) of genus m and a finite set A one obtains $|A|^m$ solutions $\mathbf{x} \in A^s$ by setting $x_i = x_{i'}$ whenever $i, i' \in T_j$. We call such solutions trivial, and seek an upper bound on the size of sets which contain only trivial solutions to (1.1). In this paper we generalise a recent result of Sanders which gives the best known bound in the case with $s = 3$ and $R = \mathbb{Z}/N\mathbb{Z}$ to obtain quantitative results of comparable quality for any $s \geq 3$ in circumstances where R is either $\mathbb{Z}/N\mathbb{Z}$ or the polynomial ring $\mathbb{F}_q[t]$. Despite the analogies between these rings having been well-explored throughout most of number theory, the problems of additive combinatorics have been little studied in the $\mathbb{F}_q[t]$ setting, and we hope that this paper will encourage others to obtain further results of this nature in polynomial rings. We also wish to promote the philosophy that $\mathbb{F}_q[t]$ is a useful model case for such problems, acting as a 'halfway house' between the currently popular finite field model case \mathbb{F}_p^N and $\mathbb{Z}/N\mathbb{Z}$, which is technically simpler than the latter but captures more of the interesting behaviour of the integer case than the former.

The problem of finding large sets with no non-trivial solutions to (1.1) may be posed in any ring (or indeed any module), but since the work of Roth (see [12, 13]) it has received most attention in the integers \mathbb{Z} , or rather finite truncations $\{1, \dots, N\}$ that may conveniently be viewed as the cyclic group $\mathbb{Z}/N\mathbb{Z}$. In the special case

Date: July 7, 2011.

The author is supported by an EPSRC doctoral training grant.

$s = 3$ and $(c_1, c_2, c_3) = (1, -2, 1)$ this is equivalent to finding a three term arithmetic progression $a, a + d, a + 2d$, with non-trivial solutions satisfying $d \neq 0$. If $s \geq 3$ and $\mathbf{c} \in (\mathbb{Z}/N\mathbb{Z})^s$ is such that (1.1) has genus $m \geq 1$ then we define $r(N) = r_{s,\mathbf{c}}(N)$ to be the cardinality of the largest subset of $\mathbb{Z}/N\mathbb{Z}$ which contains no non-trivial solutions to (1.1). Roth showed that $r(N) \ll_{s,\mathbf{c}} N/\log \log N$. This was improved for $s = 3$ by Heath-Brown [8] and Szemerédi [18] to $r(N) \ll N/\log^c N$ for some absolute constant $c > 0$, the value of which was subsequently improved by Bourgain (see [2, 3]), first to $c = 1/2 - o(1)$ and then to $c = 2/3 - o(1)$. Sanders further improved this to $c = 3/4 - o(1)$ in [15] before the recent breakthrough result of [16],

$$(1.2) \quad r(N) \ll N \frac{\log \log^5 N}{\log N}.$$

We expect that as the number of variables s increases non-trivial solutions to (1.1) should become easier to find, and hence we should obtain improved bounds for $r(N)$. The first result of this paper confirms this expectation, generalising the result of Sanders to handle arbitrary $s \geq 3$.

Theorem 1.1. *Let $s \geq 3$, and suppose $\mathbf{c} \in \mathbb{Z}^s$ satisfies the condition that the equation (1.1) has genus $m \geq 1$. Then*

$$r_{s,\mathbf{c}}(N) \ll_{s,\mathbf{c}} N \left(\frac{\log \log^5 N}{\log N} \right)^{s-2}.$$

We remark that the implicit constant in Theorem 1.1 in fact depends only on $\ell = \max_{1 \leq i \leq s} |c_i|$. In [14] Ruzsa showed that if the equation (1.1) has genus m then $r(N) \ll N^{1/m}$, which is far superior to Theorem 1.1 whenever $m \geq 2$. The power of the methods used to obtain the bounds listed above, which originated in Roth [12], is that they give a non-trivial result in the most difficult case $m = 1$. When $s \geq 6$, recent work of Schoen and Shkredov [17], building on other work of Sanders, shows that for some absolute positive constants C and c we have the near-optimal bound $r(N) \ll N \exp(-C \log^c N)$. Theorem 1.1, however, gives the sharpest known bounds for $s = 4$ and $s = 5$.

By partial summation we obtain the following corollary.

Corollary 1.1. *If $A \subset \mathbb{N}$ satisfies $\sum_{a \in A} a^{-1} = \infty$ then A contains infinitely many non-trivial solutions to every translation invariant equation of the form (1.1) having $s \geq 4$ variables.*

For comparison, a conjecture of Erdős asserts that if $\sum_{a \in A} a^{-1} = \infty$ then A must contain infinitely many arithmetic progressions of length k for any $k \in \mathbb{N}$. In the case $k = 3$, this would follow were Corollary 1.1 to hold with $s = 3$, which would in turn follow if one could improve the $\log \log^5 N$ factor in Theorem 1.1 to $\log \log^{-2} N$, for example.

Another popular setting for the problem of finding solutions to the equation (1.1) is \mathbb{F}_p^N , an N -dimensional vector space over the finite field \mathbb{F}_p . Here we have coefficients $c_i \in \mathbb{F}_p$ and we define $r_{s,\mathbf{c}}(N)$ to be the cardinality of the largest subset of \mathbb{F}_p^N that contains only trivial solutions to the equation (1.1). The best known bounds here are superior to those in $\mathbb{Z}/N\mathbb{Z}$. Thus, for example, the simple method of Meshulam

[11] yields $r(N) \ll p^N/N$, a conclusion comparable to $r(N) \ll N/\log N$ in the integer case. Furthermore, Bateman and Katz [1] have recently improved this upper bound to show that for some absolute constant $\epsilon > 0$, one has $r(N) \ll p^N/N^{1+\epsilon}$. The investigation of such problems in \mathbb{F}_p^N has recently become popular since it is often conceptually easier to prove results in \mathbb{F}_p^N , taking advantage of the vector space structure, and then ‘translate’ the methods to the more difficult case $\mathbb{Z}/N\mathbb{Z}$ (see [7] for a comprehensive discussion of this technique).

One setting that has received comparatively little attention is $\mathbb{F}_q[t]$, the ring of polynomials over a finite field; as in the case of \mathbb{Z} it is more convenient to deal with finite truncations of this infinite ring, so we shall work in the additive subgroup of polynomials with degree strictly less than N , which we shall denote by G_N . There are many well-known analogies between number theory in \mathbb{Z} and in $\mathbb{F}_q[t]$, and theorems in one often have a natural counterpart in the other. We should expect, in particular, some result analogous to (1.2) to hold for solutions to the equation (1.1) where $\mathbf{c} \in \mathbb{F}_q[t]^s$. Perhaps part of the reason why this problem has been largely overlooked is that G_N is additively isomorphic to \mathbb{F}_q^N , since both are N -dimensional vector spaces over \mathbb{F}_q . Hence in the classical case, with $s = 3$ and $\mathbf{c} = (1, -2, 1)$, the strong bounds of the \mathbb{F}_q^N setting are available. In general, such bounds are available whenever $\mathbf{c} \in \mathbb{F}_p^s$, where p is the characteristic of \mathbb{F}_q . In the $\mathbb{F}_q[t]$ setting, however, this restricts all coefficients to have degree zero, and hence trivial size under the usual valuation $|a| = q^{\deg a}$. Viewed in this way it is not surprising that we can obtain such impressive bounds with such a strong restriction on the coefficients of our linear equation. Note in particular that such issues of triviality do not arise in the $\mathbb{Z}/N\mathbb{Z}$ case, since $s \geq 3$ and translation invariance forces at least one of the coefficients of the equation (1.1) to have non-trivial size.

When $s \geq 3$ and $\mathbf{c} \in \mathbb{F}_q[t]^s$, denote by $r(N)$ the size of the largest subset of G_N which contains no non-trivial solutions to the equation (1.1). The sharpest bound for $r(N)$ currently available in the case where $\mathbf{c} \in \mathbb{F}_q^s$ is due to Liu and Spencer [10] who showed that $r(N) \ll q^N/N^{s-2}$. The second result of this paper adapts the method of Sanders to give a comparable bound even when the coefficients c_i have large degree. As above, this result gives a bound which improves as s increases.

Theorem 1.2. *Let $s \geq 3$, and suppose $\mathbf{c} \in \mathbb{F}_q[t]^s$ satisfies the condition that the equation (1.1) has genus $m \geq 1$. Then*

$$r_{s,\mathbf{c}}(N) \ll_{s,\ell} q^N \left(\frac{\log^4 N}{N} \right)^{s-2},$$

where $\ell = \max(\deg c_i)$.

Note here that the implicit constant depends only on the highest degree of the coefficients c_i .

It is not difficult to get some kind of quantitative bound for $r(N)$ even when the coefficients have large degree, and indeed the original method of Roth can be adapted with little effort to obtain $r(N) \ll q^N/\log N$. Theorem 1.2 goes beyond this to show that we can obtain a quantitative bound comparable to the best known in $\mathbb{Z}/N\mathbb{Z}$. Indeed, in $\mathbb{F}_q[t]$ we are able to do slightly better (by a factor of $\log N$) than the

analogue of (1.2). This improvement is a consequence of the vector space structure of $\mathbb{F}_q[t]$ which we are able to exploit even when the coefficients have large degree.

We remark that the methods of [14] are easily adapted to the $\mathbb{F}_q[t]$ setting to give $r(N) \ll q^{N/m}$ whenever (1.1) has genus m , a conclusion sharper than that of Theorem 1.2 whenever $m \geq 2$. The strength of Theorem 1.2 is that it applies when $m = 1$, which appears to be the most difficult case.

The main ideas and techniques used to establish Theorems 1.1 and 1.2 are those of the original argument in [16], but we are able to make many technical simplifications in $\mathbb{F}_q[t]$, largely thanks to the fact that the analogue of ‘Bohr sets’ are closed under addition. Indeed, we hope that this paper may also serve as an exposition of the methods in [16], for in $\mathbb{F}_q[t]$ the fundamental ideas are less obscured by technical difficulties. We further hope that this paper will indicate how to translate many results of additive combinatorics to an analogous result in $\mathbb{F}_q[t]$, which we believe to be a technically simpler model for such results.

In recent years there has been a focus on obtaining such results in \mathbb{F}_p^N and then translating the methods used so as to obtain a result in $\mathbb{Z}/N\mathbb{Z}$. We believe that the translation from \mathbb{F}_p^N to $\mathbb{F}_q[t]$ and thence to $\mathbb{Z}/N\mathbb{Z}$ is simpler to perform and more intuitive. We are confident that this new method will stimulate further progress in other problems of additive combinatorics.

Thus, for example, the approach taken in this paper is to first establish Theorem 1.2, the $\mathbb{F}_q[t]$ case, and then sketch how to adapt the proof to obtain Theorem 1.1, the $\mathbb{Z}/N\mathbb{Z}$ case. In Section 2 we explain the notation used in the rest of the paper and prove two useful lemmas. Sections 3 and 4 provide the two tools used in the approach of Sanders [16], namely a combinatorial transformation of sets inspired by the work of Katz and Koester [9] and an application of the useful probabilistic theorem due to Croot and Sisask [5]. Section 5 combines these tools in the ‘density increment’ strategy so as to conclude the proof of Theorem 1.2. Section 6 then sketches a proof of Theorem 1.1. Here we are brief, since the proofs are the same as those for Theorem 1.2 except for routine technical changes.

2. NOTATION AND PRELIMINARIES

Throughout the rest of this paper, let N be some fixed large integer and \mathbb{F}_q be the finite field with q elements. We shall write $G = G_N = \{x \in \mathbb{F}_q[t] : \deg x < N\}$, and note in particular that G is an N -dimensional vector space over \mathbb{F}_q . We will frequently use the expectation notation

$$\mathbb{E}_{x \in X} f(x) = \frac{1}{|X|} \sum_{x \in X} f(x).$$

Unless denoted otherwise, all expectations are taken over G . When considering $C \subset G$, we use μ_C to denote the measure induced by C , so that $\mu_C(A) = |A \cap C|/|C|$. For notational convenience we will also use β to denote μ_B . Given a subspace $B \subset G$ we will denote subsets of B by capital Roman letters, and their density within B by the corresponding lower case Greek letter, so that $\mu_B(A) = \alpha$, for example. By a convenient abuse of notation, we use the same letter for the set and its characteristic

function. We draw special attention to our novel notation for the balanced function of a subset $A \subset B$, defined by

$$\mathbf{A}(x) := (A - \alpha B)(x) = \begin{cases} 1 - \alpha & \text{if } x \in A, \\ -\alpha & \text{if } x \in B \setminus A, \text{ and} \\ 0 & \text{otherwise,} \end{cases}$$

so that $\mathbb{E}\mathbf{A}(x) = 0$. We define

$$\langle f, g \rangle_{\mu_X} = \mathbb{E}_{x \in X} f(x) \overline{g(x)},$$

and similarly, for any $1 \leq p < \infty$,

$$\|f\|_{p(\mu_X)} = \left(\mathbb{E}_{x \in X} |f(x)|^p \right)^{1/p} \quad \text{and} \quad \|f\|_{\infty(\mu_X)} = \sup_{x \in X} |f(x)|.$$

Where the measure is omitted these expectations should be taken over G . There is an analogue of Fourier analysis in the $\mathbb{F}_q[t]$ setting. If $x = \sum_{i < N} a_i t^i \in \mathbb{F}_q((1/t))$ then we define

$$e(x) := \exp(2\pi i \text{Tr}(a_{-1})/p)$$

where $\text{Tr} : \mathbb{F}_q \rightarrow \mathbb{F}_p$ is the familiar trace map. We may now define a character on G as the map $x \mapsto e(\xi x)$ where ξ is a member of the dual group \widehat{G} which is defined by

$$\widehat{G} = \{a_{-1}t^{-1} + \cdots + a_{-N}t^{-N} : a_i \in \mathbb{F}_q, -N \leq i \leq -1\}.$$

The usual definitions and results of Fourier analysis apply in this case, so we define the Fourier transform of $f \in L^1(G)$ by

$$\widehat{f}(\xi) = \mathbb{E} f(x) e(\xi x),$$

and we have Parseval's identity

$$\sum_{\xi} \widehat{f}(\xi) \overline{\widehat{g}(\xi)} = \langle f, g \rangle.$$

In most cases, when we use the convolution operator it will be with respect to the measure over some subspace B (which will be clear from the context), and hence we define

$$f * g(x) = \mathbb{E}_{y \in B} f(y) g(x - y).$$

In some cases, g will be a measure μ_C for some other set C , and there we define

$$f * \mu_C(x) = \mathbb{E}_{y \in C} f(x - y).$$

As usual, the Fourier transform converts convolution into multiplication, with an appropriate scaling factor:

$$\mu_G(B) \widehat{f * g}(\xi) = \widehat{f}(\xi) \widehat{g}(\xi).$$

This has the useful corollary, frequently used without mention in what follows, that

$$\mu_G(B) \mathbb{E} f * g = \mathbb{E} f \mathbb{E} g.$$

For $k \in \mathbb{N}$ and $\Gamma \subset \widehat{G}$ we define the Bohr set $B_k(\Gamma)$ to be the set

$$\{x : \deg\{x\xi\} < -k \text{ for all } \xi \in \Gamma\}$$

where if $x = \sum_{i < N} a_i t^i \in \mathbb{F}_q((1/t))$ then $\{x\} = \sum_{i < 0} a_i t^i$. We call the size of Γ the rank of B , denoted by $\text{rk}(B)$, and call k the width of B . We note the easy estimate

$$|B_k(\Gamma)| \geq q^{N-k|\Gamma|}.$$

Furthermore, note that $B_k(\Gamma)$ is itself a vector space over \mathbb{F}_q . This is analogous to the traditional Bohr set in the context of $\mathbb{Z}/N\mathbb{Z}$,

$$\text{Bohr}_\rho(\Gamma) = \{x : \|x\xi\| < \rho \text{ for all } \xi \in \Gamma\}$$

where $\|\cdot\|$ measures the distance to the nearest integer. For both concepts, the idea is to find a set on which a given set of characters is ‘approximately’ trivial. This is made easier in our setting of $\mathbb{F}_q[t]$, since a character can only take finitely many values even as $N \rightarrow \infty$, unlike the $\mathbb{Z}/N\mathbb{Z}$ case. The significant observation here is that the other important feature of Bohr sets is that they are approximately closed under dilation by the coefficients of our linear equation; in the $\mathbb{F}_q[t]$ case this follows from the inclusion $\lambda \cdot B_{k+m}(\Gamma) \subset B_k(\Gamma)$ which holds whenever $\deg \lambda \leq m$. We also desire that they be, at least approximately, closed under addition. This has obscured the dilation preservation feature in the traditional settings of $\mathbb{Z}/N\mathbb{Z}$ and \mathbb{F}_p^N , since this follows from closure under addition using crude estimates such as $2 \cdot A \subset A + A$.

In the density increment procedure used to prove Theorem 1.2 we will need to dilate a Bohr set by some $c \in \mathbb{F}_q[t]$ and ensure that this dilate is itself a Bohr set, so that the next iteration can be performed. The following lemma shows that this is true, and that the rank of this new Bohr set can only increase by some bounded amount. Recall that G_m denotes the set of polynomials in $\mathbb{F}_q[t]$ with degree strictly less than m .

Lemma 2.1. *If $c \in \mathbb{F}_q[t]$ and $B = B_k(\Gamma)$ is a Bohr set such that $B \subset G_{N-\deg c}$ then $c \cdot B$ is a Bohr set with rank at most $\text{rk}(B) + q^{\deg c}$ and width k .*

Proof. Choose $c^{-1} \in \mathbb{F}_q((1/t))$ such that $c^{-1}c = 1$ and let $\Gamma' = c^{-1}(\Gamma \cup G_{\deg c})$. Let $\tilde{\Gamma}$ be the truncation of Γ' to the first N negative coefficients, that is, the set of all $\gamma \in \widehat{G}$ such that $\gamma \equiv \{x\} \pmod{t^{-N}}$ for some $x \in \Gamma'$. We have $\tilde{\Gamma} \subset \widehat{G}$ and $|\tilde{\Gamma}| \leq |\Gamma'| \leq |\Gamma| + q^{\deg c}$, so to complete the proof it remains to show that $c \cdot B = B_k(\tilde{\Gamma})$.

We first note that, by orthogonality, we have the identity

$$\sum_{l \in G_{\deg c}} e(c^{-1}lx) = \begin{cases} 0 & \text{when } c \nmid x, \text{ and} \\ q^{\deg c} & \text{when } c \mid x. \end{cases}$$

By the definition of $e(\cdot)$, however, we also have $\sum_{l \in G_{\deg c}} e(c^{-1}lx) = q^{\deg c}$ if and only if $\text{Tr}(c^{-1}lx) = 0$ for all $l \in G_{\deg c}$, which is true if and only if $\deg\{c^{-1}lx\} < -1$ for all $l \in G_{\deg c}$ since $\mathbb{F}_q \cdot G_{\deg c} = G_{\deg c}$. It follows that $c \mid x$ if and only if $\deg\{c^{-1}lx\} < -1$ for all $l \in G_{\deg c}$. If we have some $y \in B_k(\tilde{\Gamma})$, therefore, then we must have $y = cx$ for some $x \in G_{N-\deg c}$. Furthermore, for all $\gamma \in \Gamma$ we have $\deg\{x\gamma\} = \deg\{yc^{-1}\gamma\} < -k$ and hence $B_k(\tilde{\Gamma}) \subset c \cdot B$. Finally, if $x \in B$ then for any $\gamma \in \Gamma$ we have $\deg\{cxc^{-1}\gamma\} =$

$\deg\{x\gamma\} < -k$ and for any $l \in G_{\deg c}$ we have $\deg\{cxc^{-1}l\} < -k$, so that $cx \in B_k(\tilde{\Gamma})$ and hence $B_k(\tilde{\Gamma}) = c \cdot B$ as required. \square

We will frequently use both the big- O notation and the Vinogradov \ll symbol, where the implicit constant depends on at most s , ℓ and q . We will also sometimes use a c or C to denote some positive constant, again depending on at most s , ℓ and q , which may vary from line to line. All these constants may be made explicit, though we do not do so here.

We shall now define the dual concepts of spectra and symmetry sets. As above, we suppose B is some subspace of G which shall be clear from context, and we define the η -spectrum of a function $f \in L^1(B)$ to be

$$\Delta_\eta(f) = \{\xi \in \widehat{B} : |\widehat{f}(\xi)| \geq \eta \|f\|_1\},$$

where \widehat{B} is the image of B under the canonical identification of G with \widehat{G} . Furthermore, for any $L, K \subset B$ we define the η -symmetry set, which is the physical space analogue of the spectrum, to be

$$\text{Sym}_\eta(L, K) = \{x \in B : L * K(x) \geq \eta\}.$$

We will need some structural information on spectra due to Chang [4]. The following version is proved as Lemma 4.36 in [19].

Lemma 2.2. *Let B be some subspace of G and $D \subset B$. If $\eta > 0$ then there is a subset $\tilde{\Delta} \subset \Delta_\eta(D)$ such that $|\tilde{\Delta}| \ll \eta^{-2} \log(1/\delta)$ and*

$$\Delta_\eta(D) \subset \{-1, 0, 1\}^{|\tilde{\Delta}|} \cdot \tilde{\Delta}.$$

We now use this to convert the fact that we have large Fourier coefficients over a large spectrum into a more useful density increment property. The idea goes back to the work of Heath-Brown [8] and Szemerédi [18], and was developed into the following form by Sanders [15].

Lemma 2.3. *For any Bohr set $B \subset G$ and $\eta > 0$ if $A, D \subset B$ satisfy*

$$\sum_{\gamma \in \Delta_\eta(D)} |\widehat{\mathbf{A}}(\gamma)|^2 \geq \nu \alpha^2 \mu_G(B)$$

then there is a Bohr set $B' \subset B$ with

$$\text{rk}(B') \leq \text{rk}(B) + O(\eta^{-2} \log(1/\delta)) \text{ and } \mu_B(B') \geq \exp(-C\eta^{-2} \log(1/\delta))$$

*such that $\|A * \beta'\|_\infty \geq \alpha(1 + \nu)$.*

Proof. Let $\tilde{\Delta} \subset \Delta_\eta(D)$ be the set given by Lemma 2.2, and refine the Bohr set B by setting

$$B' = \{x \in B : \deg\{x\gamma\} < -1 \text{ for all } \gamma \in \tilde{\Delta}\}.$$

It follows that $\text{rk}(B') \leq \text{rk}(B) + |\tilde{\Delta}| \leq \text{rk}(B) + O(\eta^{-2} \log(1/\delta))$, and hence $\mu_B(B') \geq \exp(-C\eta^{-2} \log(1/\delta))$. Furthermore, for all $\gamma \in \Delta_\eta(D)$ we have $\widehat{\beta}'(\gamma) = \mathbb{E}_{x \in B'} e(x\gamma) = 1$. By Parseval's identity and the initial hypothesis

$$\|A * \beta'\|_2^2 = \sum_{\gamma \in \widehat{G}} |\widehat{\mathbf{A}}(\gamma)|^2 |\widehat{\beta}'(\gamma)|^2 \geq \sum_{\gamma \in \Delta_\eta(D)} |\widehat{\mathbf{A}}(\gamma)|^2 \geq \nu \alpha^2 \mu_G(B).$$

Furthermore, note that $B * \beta'(y) = \mathbb{E}_{x \in B'} B(y-x) = B(y)$, and hence $\langle B * \beta', B * \beta' \rangle = \mu_G(B)$. Similarly,

$$\langle A * \beta', B * \beta' \rangle = \mu_G(B) \prod_{x \in B} \prod_{y \in B'} A(y-x) = \alpha \mu_G(B).$$

In particular,

$$\begin{aligned} \|\mathbf{A} * \beta'\|_2^2 &= \|A * \beta'\|_2^2 + \alpha^2 \langle B * \beta', B * \beta' \rangle - 2\alpha \langle A * \beta', B * \beta' \rangle \\ &= \|A * \beta'\|_2^2 - \alpha^2 \mu_G(B). \end{aligned}$$

It follows that $\|A * \beta'\|_2^2 \geq \alpha^2(1 + \nu) \mu_G(B)$. Combining Hölder's inequality with the equality $\|A * \beta'\|_1 = \alpha \mu_G(B)$ gives $\|A * \beta'\|_\infty \geq \alpha(1 + \nu)$ as required. \square

We make one final remark on the values of the densities α of the various $A \subset G$ we consider, where these A will be sets with no non-trivial solutions to the equation (1.1). Recall that we may assume $\alpha \leq r(N) \ll 1/\log N$ by adapting the method of Roth. In particular, for sufficiently large N , we will assume that $\alpha < 1/e^2$, so that $\log(1/\alpha) > 2$. Making this assumption avoids some notational awkwardness.

3. A COMBINATORIAL TRANSFORMATION

We begin with a generalisation of one of the two main tools in the method of Sanders, a combinatorial transformation. This has its origins in the work of Katz and Koester [9], though a similar transformation known as the Dyson e -transform has been a useful tool in additive combinatorics since [6].

The idea is to find, given subsets A_1, \dots, A_{k+1} of a Bohr set B , corresponding $L, S_1, \dots, S_k \subset B$ such that L is 'thick', the S_i are not too 'thin' and

$$L * S_1 * \dots * S_k \leq \alpha_1^{-2} A_1 * A_2 * \dots * A_{k+1}.$$

Since solutions to (1.1) are counted by $\langle (c_1 \cdot A) * \dots * (c_{s-1} \cdot A), c_s \cdot A \rangle$ to find many solutions to (1.1) in A it then suffices to find a lower bound for $\langle L * \dots * S_{s-2}, c_s \cdot A \rangle$. The large density of L may then be efficiently exploited using the recent probabilistic method of Croot and Sisask, as we shall do in the next section. When we cannot find such sets we find a strong density increment and jump to the final stage of the proof.

We first give a simplified proof of Lemma 4.2 from [16], taking advantage of the fact that our Bohr sets are subspaces in order to streamline the proof. We also observe that the proof gives a slightly stronger statement than is recorded in [16]. We then apply this technical lemma iteratively to construct sets L, S_1, \dots, S_k as above to prove the main theorem of this section. This is a generalisation of Proposition 4.1 from [16], taking advantage of the large number of convolutions to run a more efficient iterative procedure.

Lemma 3.1. *Let B be a Bohr set and $K, T, L, S \subset B$. If $\lambda < 1/4$ then either*

- (1) *there exists a Bohr set $B' \subset B$ with rank at most $rk(B) + O(\lambda \kappa^{-1} \log(1/\tau))$, density $\mu_B(B') \geq \exp(-C\lambda \kappa^{-1} \log(1/\tau))$ such that, for some $x \in B$, $\beta'(K + x) \geq \kappa/32\lambda$, or*

(2) there are $L', S' \subset B$ with $\beta(L') \geq \lambda + \kappa/2$ and $\beta(S') \geq \tau\sigma/2$ such that, for all $x \in G$,

$$L' * S'(x) \leq L * S(x) + K * T(x).$$

Proof. For an arbitrary $x \in B$, let $L_x = L \cup (K + x)$ and $S_x = S \cap (T - x)$. By construction,

$$L_x * S_x \leq L * S_x + (K + x) * S_x \leq L * S + K * T.$$

Furthermore, we have

$$\beta(L_x) = \lambda + \kappa - \beta(L \cap (K + x)) = \lambda + \kappa - L * (-K)(x)$$

and $\beta(S_x) = (-S) * T(x)$. Hence if $|\text{Sym}_{\tau\sigma/2}(-S, T)| > |\text{Sym}_{\kappa/2}(-L, K)|$ then we are in the second case by taking $L' = L_x$ and $S' = S_x$ for some $x \in \text{Sym}_{\tau\sigma/2}(-S, T) \setminus \text{Sym}_{\kappa/2}(L, -K)$. Note that

$$\beta(\text{Sym}_{\tau\sigma/2}(-S, T))\sigma + \sigma\tau/2 \geq \int_{x \in B} (-S) * T(x) = \sigma\tau,$$

so that $\beta(\text{Sym}_{\tau\sigma/2}(-S, T)) \geq \tau/2$. Hence either we are in the second case of the theorem, or we may assume that $D = \text{Sym}_{\kappa/2}(-L, K)$ has density at least $\tau/4$, say. We have

$$\langle (-L) * K, D \rangle_\beta \geq \kappa\beta(D)/2.$$

Note that

$$\langle (-L) * B, D \rangle_\beta = \lambda\beta(D).$$

Hence, by the triangle inequality,

$$|\langle (-L) * \mathbf{K}, D \rangle_\beta| \geq \kappa\beta(D)(1/2 - \lambda) \geq \kappa\beta(D)/4.$$

Taking the Fourier transform of the left hand side we see that

$$\left| \sum_{\gamma \in \widehat{G}} \widehat{(-L)}(\gamma) \widehat{\mathbf{K}}(\gamma) \widehat{D}(\gamma) \right| \geq \kappa\beta(D)\mu_G(B)^2/4.$$

The Cauchy-Schwarz inequality combined with Parseval's identity then gives

$$\sum_{\gamma \in \widehat{G}} |\widehat{\mathbf{K}}(\gamma)|^2 |\widehat{D}(\gamma)|^2 \geq \kappa^2\beta(D)^2\mu_G(B)^3/16\lambda.$$

Furthermore, if we let $\eta^2 = \kappa/32\lambda$ then by Parseval's identity again we have

$$\sum_{\gamma \notin \Delta_\eta(D)} |\widehat{\mathbf{K}}(\gamma)|^2 |\widehat{D}(\gamma)|^2 \leq \kappa^2\beta(D)^2\mu_G(B)^3/32\lambda.$$

Using the trivial upper bound $|\widehat{D}(\gamma)| \leq \beta(D)\mu_G(B)$ it follows that

$$\sum_{\gamma \in \Delta_\eta(D)} |\widehat{\mathbf{K}}(\gamma)|^2 \geq \kappa^2\mu_G(B)/32\lambda.$$

Applying Lemma 2.3, we have a Bohr set $B' \subset B$ with rank and size as required and $\beta'(K + x) \geq \kappa/32\lambda$ for some $x \in X$. \square

We now apply Lemma 3.1 iteratively to prove the main theorem of this section.

Theorem 3.1. *Let B be a Bohr set and suppose $A_1 \subset B$ with density α_1 and $A_2, \dots, A_{k+1} \subset B$ all with density at least α . Then either*

- (1) *there exists a Bohr set $B' \subset B$ with rank $\text{rk}(B') \leq \text{rk}(B) + O(\alpha_1^{-1/k} \log(1/\alpha))$, density $\mu_B(B') \geq \exp(-C\alpha_1^{-1/k} \log(1/\alpha))$ and, for some $x \in B$, we have $\beta'(A_1 + x) \geq 2\alpha_1$, or*
- (2) *there are sets $L \subset B$ and $S_1, \dots, S_k \subset B$ with $\lambda \geq 2^{-k-6}$ and $\sigma_i \geq \alpha^C \alpha_1^{-1/k}$ for $1 \leq i \leq k$ such that*

$$L * S_1 * \dots * S_k(x) \leq \alpha_1^{-2} A_1 * A_2 * \dots * A_{k+1}(x)$$

for all $x \in G$.

Proof. Let $L_0 = A_1$. We show by induction on $1 \leq j \leq k$ that we may either find $X_1, \dots, X_j \subset B$ such that $L_j = A_1 + X_1 + \dots + X_j$ satisfies

$$\lambda_j \geq \max(\alpha_1 |X_1| \dots |X_j| / 2^j, \alpha_1^{1-j/k} / 2^j)$$

and $S_1, \dots, S_j \subset B$ satisfying $\sigma_i \geq \alpha^C \alpha_1^{-1/k}$ for $1 \leq i \leq j$ such that

$$L_j * S_1 * \dots * S_j * A_{j+2} * \dots * A_{k+1}(x) \leq \alpha_1^{-2j/k} A_1 * A_2 * \dots * A_{k+1}(x)$$

for all $x \in G$, or there is a Bohr set $B' \subset B$ such that $\text{rk}(B') \leq \text{rk}(B) + O(\alpha_1^{-1/k} \log(1/\alpha))$ and for some $x \in B$ we have $\beta'(L_{j-1} + x) \geq 2\alpha_1 |X_1| \dots |X_{j-1}|$. If we are in the second case for any $1 \leq j \leq k$ then we halt the inductive procedure and by the pigeonhole principle we must have, for some $x_i \in X_i$, $\beta'(A_1 + x_1 + \dots + x_{j-1} + x) \geq 2\alpha_1$, giving us the first case of the theorem. Otherwise, setting $j = k$ and $L = L_k$ gives the second case of the theorem.

We now fix $1 \leq j \leq k$ and suppose that we have produced L_{j-1} and S_1, \dots, S_{j-1} as above. We produce a sequence of sets $L_j^{(i)}, S_j^{(i)} \subset B$ iteratively such that $\lambda_{j-1}^{(i)} \geq \lambda_j^{(i)} \geq \lambda_{j-1}^{(i-1)}/2$, $\sigma_j^{(i)} \geq (\alpha/2)^i$, and $L_j^{(i)} * S_j^{(i)} \leq i L_{j-1} * A_{j+1}$. We begin the iteration by letting $L_j^{(1)} := L_{j-1}$ and $S_j^{(1)} := A_{j+1}$.

We now repeatedly apply Lemma 3.1, with $K = L_{j-1}$ and $T = A_{j+1}$, and note that an examination of the proof shows that we take $L_j^{(i+1)} = L_j^{(i)} \cup (L_{j-1} + x_{i+1})$ for some $x_{i+1} \in B$ so that $L_j^{(i)} = L_{j-1} + X_j^{(i)}$ where $|X_j^{(i)}| \leq i$. If we can continue the process for $i' = \lceil \alpha_1^{-1/k} \rceil$ steps, then we can halt and set $L_j = L_j^{(i')}$, so that $L_j = L_{j-1} + X_j$ where

$$\lambda_j \geq \lambda_{j-1} \lceil \alpha_1^{-1/k} \rceil / 2 \geq \max(\alpha_1 |X_1| \dots |X_j| / 2^j, \alpha_1^{1-j/k} / 2^j)$$

by inductive hypothesis. Furthermore, if we set $S_j = S_j^{(i')}$ then $\sigma_j \geq (\alpha/2)^{i'} \geq \alpha^C \alpha_1^{-1/k}$ and by inductive hypothesis again,

$$\begin{aligned} L_j * \dots * S_j * A_{j+2} * \dots * A_{k+1} &\leq \lceil \alpha_1^{-1/k} \rceil L_{j-1} * \dots * S_{j-1} * A_{j+1} * \dots * A_{k+1} \\ &\leq \alpha_1^{-2j/k} A_1 * A_2 * \dots * A_{k+1} \end{aligned}$$

as required.

Otherwise, we must halt at some $i' \leq \alpha_1^{-1/k}$ with the first alternative from Lemma 3.1, so that there is a Bohr set $B' \subset B$ with rank at most

$$\text{rk}(B) + O(\lambda_j^{(i')} \lambda_{j-1}^{-1} \log(1/\alpha)) \leq \text{rk}(B) + O(\alpha_1^{-1/k} \log(1/\alpha))$$

and density $\mu_B(B') \geq \exp(-C\alpha_1^{-1/k} \log(1/\alpha))$ such that, for some $x \in B$,

$$\beta'(L_{j-1} + x) \geq \lambda_{j-1}/32\lambda_j^{(i')} \geq \alpha_1 |X_1| \cdots |X_{j-1}| / 2^{j+5} \lambda_j^{(i')} \geq 2\alpha_1 |X_1| \cdots |X_{j-1}|$$

since if $\lambda_j^{(i')} > 2^{-j-6} > 2^{-k-6}$ we may end the entire procedure with the second case of the theorem. This completes the induction. \square

4. THE CROOT-SISASK PROBABILISTIC METHOD

We now use the new Croot-Sisask method of random sampling from [5] to find a density increment that can exploit the fact that one set in a convolution is very dense, while allowing us to ignore the density loss in the other sets that results from applying Theorem 3.1. The operator τ_t is defined by $\tau_t(f)(x) = f(x-t)$. The theorem finds a large set T such that the translation operator τ_t is essentially constant on the convolution $L * S_1 * \cdots * S_k$ for all $t \in T$. This set may have no structure initially, but Croot and Sisask observed that we may recover some structure by taking repeated sumsets of T .

We first state the theorem from [5] that we require, before proving an analogue of Corollary 5.2 of [16] for multiple convolutions. The proof given here is very similar, but once again we are able to present a simplified and streamlined version using the subspace structure of our Bohr sets.

Theorem 4.1. *Let B be an additive group, $S \subset B$ and $f \in L^1(B)$. Then for any $\epsilon > 0$ and $p \geq 2$ there is some $T \subset B$ with $\beta(T) \geq \sigma^{C\epsilon^{-2p}}$ such that*

$$\|\tau_t(f * \mu_S) - f * \mu_S\|_{p(\beta)} \leq \epsilon \|f\|_{p(\beta)}$$

for all $t \in T$.

This is a global version of Theorem 5.1 from [16], which we are able to apply here because Bohr sets are additive groups. When they are only approximately closed under addition, as in $\mathbb{Z}/N\mathbb{Z}$, the local theorem of [16] must be used. A proof may be found in [5] and [16]. The following theorem applies this to obtain a useful density increment.

Theorem 4.2. *Let B be a Bohr set, and suppose $A, L, S_1, \dots, S_k \subset B$. Then either*

- (1) $\langle L * S_1 * \cdots * S_k, A \rangle_\beta \geq \lambda \sigma_1 \cdots \sigma_k \alpha / 2$, or
- (2) for any integer l there exists a Bohr set $B' \subset B$ such that

$$\text{rk}(B') \leq \text{rk}(B) + O(\lambda^{-2-1/l} \alpha^{-1/l^2} \log(1/\alpha) \log(1/\sigma_k)),$$

$$\mu_B(B') \geq \exp(-C\lambda^{-2-1/l} \alpha^{-1/l^2} \log(1/\alpha) \log(1/\sigma_k))$$

and, for some $x \in B$, we have $\beta'(A + x) \geq \alpha(1 + c\lambda)$.

Proof. Applying Theorem 4.1 with parameters p and ϵ to be chosen later, where $f = L * \mu_{S_1} * \cdots * \mu_{S_{k-1}}$, there is a set $T \subset B$ such that $|T| \geq \sigma_k^{C\epsilon^{-2}p}|B|$ and

$$\begin{aligned} \|\tau_t(L * \mu_{S_1} * \cdots * \mu_{S_k}) - L * \mu_{S_1} * \cdots * \mu_{S_k}\|_{p(\beta)} &\leq \epsilon \|L * \mu_{S_1} * \cdots * \mu_{S_{k-1}}\|_{p(\beta)} \\ &\leq \epsilon \|L\|_{p(\beta)} \end{aligned}$$

for all $t \in T$, or after rescaling,

$$\|\tau_t(L * S_1 * \cdots * S_k) - L * S_1 * \cdots * S_k\|_{p(\beta)} \leq \epsilon \sigma_1 \cdots \sigma_k \lambda^{1/p} \leq \epsilon \sigma_1 \cdots \sigma_k.$$

By the triangle inequality, for any $l \in \mathbb{N}$ and $t \in l$,

$$\|\tau_t(L * S_1 * \cdots * S_k) - L * S_1 * \cdots * S_k\|_{p(\beta)} \leq l \epsilon \sigma_1 \cdots \sigma_k.$$

If we let $g := \mu_T * \cdots * \mu_T$ with l copies of μ_T , then after recalling the definition of τ_t and applying the triangle inequality again we get

$$\|L * S_1 * \cdots * S_k * g - L * S_1 * \cdots * S_k\|_{p(\beta)} \leq l \epsilon \sigma_1 \cdots \sigma_k.$$

By Hölder's inequality,

$$\begin{aligned} |\langle L * S_1 * \cdots * S_k * g, A \rangle_\beta - \langle L * S_1 * \cdots * S_k, A \rangle_\beta| &\leq l \epsilon \sigma_1 \cdots \sigma_k \|A\|_{p/p-1(\beta)} \\ &\leq \lambda \sigma_1 \cdots \sigma_k \alpha / 4 \end{aligned}$$

if we let $p = \lceil \log(1/\alpha) \rceil$ and $\epsilon = \lambda/4el$. Hence either we are in the first case, or

$$\langle L * S_1 * \cdots * S_k * g, A \rangle_\beta \leq 3\lambda \sigma_1 \cdots \sigma_k \alpha / 4.$$

Since $\mathbb{E}_{x \in B} g = 1$ we have

$$\mathbb{E}_{x \in B} L * S_1 * \cdots * S_k * g(x) = \lambda \sigma_1 \cdots \sigma_k.$$

Hence

$$|\langle L * S_1 * \cdots * S_k * g, \mathbf{A} \rangle_\beta| \geq \lambda \sigma_1 \cdots \sigma_k \alpha / 4.$$

We complete the proof by a standard conversion into Fourier space, as in the proof of Lemma 3.1. Taking the Fourier transform of this inequality we see that

$$\left| \sum_{\gamma \in \widehat{G}} \widehat{L}(\gamma) \widehat{S}_1(\gamma) \cdots \widehat{S}_k(\gamma) \widehat{g}(\gamma) \widehat{\mathbf{A}}(\gamma) \right| \geq \lambda \sigma_1 \cdots \sigma_k \alpha \mu_G(B)^{k+1} / 4.$$

Using the Cauchy-Schwarz inequality combined with Parseval's identity and the trivial bound $|\widehat{S}_i(\gamma)| \leq \sigma_i \mu_G(B)$ gives

$$\sum_{\gamma \in \widehat{G}} \left| \widehat{g}(\gamma) \widehat{\mathbf{A}}(\gamma) \right|^2 \geq \lambda \alpha^2 \mu_G(B) / 16.$$

Recalling the definition of g this is

$$\sum_{\gamma \in \widehat{G}} |\widehat{\mu_T}(\gamma)|^{2l} \left| \widehat{\mathbf{A}}(\gamma) \right|^2 \geq \lambda \alpha^2 \mu_G(B) / 16.$$

Furthermore, if we let $\eta^{2l} = \lambda \alpha / 32$ then, by Parseval's identity,

$$\sum_{\gamma \notin \Delta_\eta(T)} |\widehat{\mu_T}(\gamma)|^{2l} \left| \widehat{\mathbf{A}}(\gamma) \right|^2 \leq \lambda \alpha^2 \mu_G(B) / 32.$$

Combining these inequalities we have

$$\sum_{\gamma \in \Delta_\eta(T)} \left| \widehat{\mathbf{A}}(\gamma) \right|^2 \geq \sum_{\gamma \in \Delta_\eta(T)} |\widehat{\mu}_T(\gamma)|^{2l} \left| \widehat{\mathbf{A}}(\gamma) \right|^2 \geq \lambda \alpha^2 \mu_G(B)/32$$

and the first case follows from Lemma 2.3. \square

5. DENSITY INCREMENT STRATEGY

We now prove Theorem 1.2 using the traditional density increment strategy: we construct a series of Bohr sets $B^{(i)}$ such that at each stage, either $A \cap B^{(i)}$ has many solutions to (1.1), or we may find a new Bohr set $B^{(i+1)} \subset B^{(i)}$ on which the density of A increases. Since the density is always trivially bounded above by 1, we can only iterate this a bounded number of times before we must end with the first alternative, from which we deduce the existence of many solutions to (1.1) in the original set A , and hence deduce that A must contain at least one non-trivial solution if α is sufficiently large.

We define $\Lambda(A)$ by letting $\Lambda(A)|G|^{s-1}$ be the number of solutions (x_1, \dots, x_s) to (1.1) with $x_i \in A$. Note in particular that we include the trivial solutions so that $\Lambda(A) \geq |A|^m |G|^{1-s}$, where m is the genus of (1.1).

We first combine the results from the previous two sections to obtain a density increment which we then iterate to provide a lower bound for $\Lambda(A)$, from which we deduce Theorem 1.2 as a corollary. This proof is a generalised form of the one found in [16].

Lemma 5.1. *Let B be a Bohr set and $A_i \subset B$ for $1 \leq i \leq s$, all with density at least α . Then either*

- (1) $\langle A_1 * A_2 \cdots * A_{s-1}, A_s \rangle_\beta \gg \alpha^{C\alpha^{-1/(s-2)}}$, or
- (2) we have a Bohr set $B' \subset B$ of rank $\text{rk}(B) + O(\alpha^{-1/(s-2)} \log^4(1/\alpha))$ and size

$$\mu_B(B') \geq \exp(-C\alpha^{-1/(s-2)} \log^4(1/\alpha))$$

such that for some $x \in B$ and $i \in \{1, s\}$ we have $\beta'(A_i + x) \geq \alpha(1 + c)$.

Proof. By Theorem 3.1, either we are in the second case of the theorem or there are $L, S_1, \dots, S_{s-2} \subset B$ with $\lambda > 2^{-s-3}$ and $\sigma_i \geq \alpha^{C\alpha^{-1/(s-2)}}$ for $1 \leq i \leq s-2$ such that

$$A_1 * A_2 \cdots * A_{s-1} \geq \alpha^2 L * S_1 * \cdots * S_{s-2}.$$

We may now apply Theorem 4.2 to get that either for any integer l we have a Bohr set $B' \subset B$ such that

$$\text{rk}(B') \leq \text{rk}(B) + O(\alpha^{-1/(s-2)-1/2l} l^2 \log^2(1/\alpha))$$

and, for some $x \in B$, we have $\beta'(A_s + x) \geq \alpha(1 + c)$, or

$$\langle L * S_1 * \cdots * S_{s-2}, A_s \rangle_\beta \geq \sigma_1 \cdots \sigma_{s-2} \lambda \alpha^2 / 2 \gg \alpha^{C\alpha^{-1/(s-2)}}.$$

Combining this with the bound above, we are either in the first or third case of the theorem, after taking a roughly optimal choice of $l = \lceil \log(1/\alpha) \rceil$. \square

We now prove the following generalisation of Theorem 1.2 using the traditional density increment strategy.

Theorem 5.1. *Let $s \geq 3$, and suppose that $\mathbf{c} \in \mathbb{F}_q[t]^s$ satisfies the condition that (1.1) has genus $m \geq 1$. Then there exists a constant $C > 0$, depending only on s, q and $\ell = \max(\deg c_i)$, such that if $A \subset G$ then*

$$\Lambda(A) \geq \exp(-C\alpha^{-1/(s-2)} \log^4(1/\alpha)).$$

Proof. We shall iteratively construct a sequence of Bohr sets $B^{(i)}$ and sets $A^{(i)} \subset G$, each a translate and dilate of A . Let $\text{rk}(B^{(i)}) = d^{(i)}$ and $\beta^{(i)}(A^{(i)}) = \alpha^{(i)}$. We shall insist that, given $B^{(i)}$ and $A^{(i)}$, we can either find a $B^{(i+1)}$ and $A^{(i+1)}$ such that for some constant $c' > 0$ (depending only on s, ℓ and q), the density satisfies $\alpha^{(i+1)} \geq \alpha^{(i)}(1 + c')$,

$$d^{(i+1)} \leq d^{(i)} + O(\alpha^{(i)(-1/(s-2))} \log^4 \alpha^{-i}),$$

and

$$\mu_{B^{(i)}}(B^{(i+1)}) \geq \exp(-C(d^{(i)} + \alpha^{(i)(-1/(s-2))} \log^4 \alpha^{-i})),$$

or we have

$$\Lambda(A) \geq \exp(-Cd^{(i)})\mu_G(B^{(i)})^{s-1}\alpha^{C\alpha^{-1/(s-2)}}.$$

We begin the iteration by letting $B^{(1)} = G_{N-s\ell}$ and $A^{(1)} = A \cap B^{(1)}$ (so that $d^{(1)} \leq q^{s\ell}$ and $\alpha^{(1)} \geq \alpha - q^{s\ell-N}$). Since we always have the trivial bound $\alpha^{(i)} \leq 1$, we can iterate this process only $O(\log(1/\alpha))$ times, hence at some $B^{(K)}$, with $K \ll \log(1/\alpha)$, we have the second alternative. Here we have

$$\begin{aligned} d^{(K)} &\ll \sum_{i=1}^K \alpha^{(i)(-1/(s-2))} \log^4 \alpha^{-i} \ll \alpha^{-1/(s-2)} \log^4(1/\alpha) \sum_{i=0}^K (1+c')^{-i} \\ &\ll \alpha^{-1/(s-2)} \log^4(1/\alpha). \end{aligned}$$

Similarly,

$$\mu_G(B^{(K)}) \geq \exp(-C\alpha^{-1/(s-2)} \log^4(1/\alpha))$$

from which the result follows.

We now discuss the iterative procedure. We are given a Bohr set $B^{(i)}$ and a set $A^{(i)}$. To reduce notation in what follows, let $\tilde{A} = A^{(i)} \cap B^{(i)}$ and $B = B^{(i)}$. We have $\tilde{A} \subset B$ with density $\tilde{\alpha} = \alpha^{(i)}$, where B is a Bohr set with rank d and width k . Note that for any $B' \subset B$ we have $\mathbb{E}_{x \in B} \tilde{A} * \beta'(x) = \mathbb{E}_{x \in B'} \tilde{A} * \beta(x) = \tilde{\alpha}$. In particular, let B' be the Bohr set with the same frequency set as B but width $k + s\ell$, and for $1 \leq j \leq s$ define $B_j = c_1 \cdots c_{j-1} c_{j+1} \cdots c_s \cdot B'$. Since $B_j \subset B$ for all $1 \leq j \leq s$, we have

$$\mathbb{E}_{x \in B} (\tilde{A} * \beta_1 + \cdots + \tilde{A} * \beta_s)(x) = s\tilde{\alpha}$$

and in particular there is some $x \in B$ such that

$$\tilde{A} * \beta_1(x) + \cdots + \tilde{A} * \beta_s(x) \geq s\tilde{\alpha}.$$

Note that if

$$\tilde{A} * \beta_j(x) \geq \tilde{\alpha} \left(1 + \frac{c}{2(s-1)}\right)$$

for any $1 \leq j \leq s$ then we may proceed to the next stage of the iteration, letting $B^{(i+1)} = B_j$ and $A^{(i+1)} = A^{(i)}$. Here we use Lemma 2.1 which ensures B_j is a

genuine Bohr set, after noting that since we chose $B^{(1)} = G_{N-sl}$ the condition $B' \subset G_{N-\deg\{c_1 \cdots c_{j-1} c_{j+1} \cdots c_s\}}$ will always be satisfied. Otherwise we must have $\tilde{A} * \beta_j(x) \geq \tilde{\alpha} \left(1 - \frac{\epsilon}{2}\right)$ for all $1 \leq j \leq s$.

Hence if we let $B'' = c_1 \cdot B_1 = \cdots = c_s \cdot B_s$, and similarly $A_j = c_j(A^{(i)} - x)$, then we have $\beta''(A_j) \geq \tilde{\alpha} \left(1 - \frac{\epsilon}{2}\right)$ for $1 \leq j \leq s$. By Lemma 2.1 we see that B'' is a genuine Bohr set and furthermore $B'' \subset B$. We now satisfy the conditions of Lemma 5.1; hence either

$$\langle A_1 * A_2 \cdots * A_{s-1}, A_s \rangle_{\beta''} \geq \alpha^{C\alpha^{-1/(s-2)}},$$

or there is a Bohr set $B^* \subset B''$ of rank $d + O(\alpha^{-1/(s-2)} \log^4(1/\alpha))$ and density

$$\mu_B(B^*) = \mu_B(B'')\mu_{B''}(B^*) \geq \exp(-C(d + \alpha^{-1/(s-2)} \log^4(1/\alpha)))$$

such that either A_1 or A_s has density increment at least $\tilde{\alpha}(1+c)(1-\frac{\epsilon}{2}) \geq \tilde{\alpha}(1+\frac{\epsilon}{4})$. In this case we are done, letting $B^{(i+1)} = B^*$ and $A^{(i+1)}$ be A_1 or A_s as appropriate.

Otherwise, note that

$$\Lambda(A) \geq \mu_G(B'')^{s-1} \langle A_1 * A_2 \cdots * A_{s-1}, A_s \rangle_{\beta''}.$$

Here is where we make essential use of the translation invariant condition $\sum c_i = 0$, for the inner product on the right hand side only counts solutions in a set which may be translated and dilated many times from our original A , but each solution can be retranslated and dilated to give a solution in the original set. Putting this all together and noting the estimate $\mu_B(B'') \geq \exp(-Cd)$ we have

$$\Lambda(A) \geq \mu_G(B'')^{s-1} \alpha^{C\alpha^{-1/(s-2)}} \geq \exp(-Cd) \mu_G(B)^{s-1} \alpha^{C\alpha^{-1/(s-2)}}$$

as required. \square

We finally prove Theorem 1.2 as a corollary.

Proof of Theorem 1.2. Suppose $A \subset G$ with density α has no non-trivial solutions to (1.1), so that $\Lambda(A) = \alpha^m q^{(1+m-s)N}$. Theorem 5.1 then implies $N \ll \alpha^{-1/(s-2)} \log^4(1/\alpha)$. If $\alpha \ll N^{-s+2}$ then we are done; otherwise, we have $\log(1/\alpha) \ll \log N$, and hence

$$\alpha \ll \left(\frac{\log^4 N}{N} \right)^{s-2}$$

as required. \square

6. THE INTEGER CASE

We now sketch how to prove Theorem 1.1 using the ideas from previous sections. The results here will be much closer in spirit to [16], and we focus on explaining the process of translating a proof in $\mathbb{F}_q[t]$ to one in $\mathbb{Z}/N\mathbb{Z}$.

Given $\rho > 0$ and $\Gamma \subset \mathbb{Z}/N\mathbb{Z}$ we define the Bohr set $B_\rho(\Gamma)$ to be

$$B_\rho(\Gamma) = \{x \in \mathbb{Z}/N\mathbb{Z} : |e(\gamma x) - 1| < \rho \text{ for all } \gamma \in \Gamma\}$$

where we are now using the character given by $e(x) = \exp(2\pi i x/N)$. As in the $\mathbb{F}_q[t]$ case, we define $\text{rk}(B) = |\Gamma|$ and call ρ the width of $B_\rho(\Gamma)$. We have the analogous size estimate $|B_\rho(\Gamma)| \geq \rho^{|\Gamma|} N$.

The reason for the extra technical complications in the $\mathbb{Z}/N\mathbb{Z}$ case is that it is no longer true that $B + B = B$ for a Bohr set B . In general, the best we can do is $B_\rho(\Gamma) + B_\rho(\Gamma) \subset B_{2\rho}(\Gamma)$, which may be exponentially larger. The key insight here, originating with Bourgain [2], is that if we restrict ourselves to instead consider $B + B'$ for some smaller Bohr set $B' \subset B$, and if B is sufficiently well-behaved, then we know that $B + B'$ is *mostly* contained within B , and this is enough to run the previous sort of arguments while incurring only a small error term.

To make this precise we call a Bohr set $B = B_\rho(\Gamma)$, where $\text{rk}(B) = k$, regular if

$$\frac{|B|}{1 + k|\eta|/100} \leq |B_{(1+\eta)\rho}(\Gamma)| \leq (1 + k|\eta|/100)|B|$$

whenever $|\eta| \leq 1/100k$. Given any Bohr set B we can always find a regular Bohr set that closely approximates it. The following is given as Lemma 4.25 in [19].

Lemma 6.1. *Given a Bohr set $B_\rho(\Gamma)$ there is $\epsilon \in [1/2, 1)$ such that $B_{\epsilon\rho}(\Gamma)$ is regular.*

Any statement which relies on the subspace structure of Bohr sets B in the previous arguments can then be replaced with an approximate version, incurring an error term dependent on the width parameter ρ , which can be controlled by choosing ρ sufficiently small. In many cases this means ρ must depend on α , which results in the extra logarithmic factors in Theorem 1.1. For example, the appeal to the identity $\langle A, B * \beta' \rangle = \alpha \mu_G(B)$ in the proof of Lemma 2.3 can be replaced by an appeal to the following lemma.

Lemma 6.2. *If $B = B_\rho(\Gamma)$ is a regular Bohr set with rank k and $B' \subset B_{\epsilon\rho}(\Gamma)$ for some $\epsilon \ll 1/k$ then for any $f \in L^1(G)$ with $\|f\|_\infty \leq 1$ we have*

$$\langle f, B * \beta' \rangle = \left(\mathbb{E}_{x \in B} f(x) + O(\epsilon k) \right) \mu_G(B).$$

Proof. Since $\text{supp}(B * \beta') = B + B' \subset B_{(1+\epsilon)\rho}(\Gamma)$, we have

$$|G| \langle f, B * \beta' \rangle = \sum_{x \in B_{(1+\epsilon)\rho}(\Gamma)} f(x) B * \beta'(x).$$

Furthermore, whenever $x \in B_{(1-\epsilon)\rho}(\Gamma)$, we have $B' + x \subset B$ and hence $B * \beta'(x) = 1$. It follows that the sum above is

$$\sum_{x \in B} f(x) + \sum_{x \in B \setminus B_{(1-\epsilon)\rho}(\Gamma)} f(x)(B * \beta'(x) - 1) + \sum_{x \in B_{(1+\epsilon)\rho}(\Gamma) \setminus B} f(x) B * \beta'(x).$$

By regularity, however, in each of the second sums there are $O(\epsilon k)|B|$ summands, and the result follows by the triangle inequality. \square

We restrict ourselves to stating the main tools needed for the proof of Theorem 1.1, which closely follows the proof of Theorem 1.2. Making the required technical changes is a simple but lengthy matter, and is easily done by comparing the proofs given in the rest of this paper to the proofs for the $s = 3$ and $\mathbb{Z}/N\mathbb{Z}$ case given in [16], which works in this more complicated setting from the beginning.

The following are the analogues of Theorems 3.1 and 4.2 (and also Proposition 4.1 and Corollary 5.2 from [16]) respectively. Note that the only change is that we restrict

certain sets to be inside a sub-Bohr set of the main Bohr set, with width suitably restricted.

Theorem 6.1. *Let $B_\rho(\Gamma)$ and $B' = B_{\rho'}(\Gamma')$ be regular rank d Bohr sets such that $B' \subset B_{\epsilon\rho}(\Gamma)$. Suppose further that $A_1 \subset B$ and $A_2, \dots, A_{k+1} \subset B'$ all with density at least α . Finally, suppose $\epsilon \leq c\alpha_1/d$ and $\epsilon' \leq c\alpha/d$ for some sufficiently small absolute constant $c > 0$. Then either*

- (1) *there exists a regular Bohr set $B'' \subset B'$ such that*

$$rk(B'') \leq rk(B) + O(\alpha_1^{-1/k} \log(1/\alpha)),$$

$$\mu_{B'}(B'') \geq \left(\frac{\alpha_1}{d \log(1/\alpha)} \right)^{C(d + \alpha_1^{-1/k} \log(1/\alpha))}$$

and, for some $x \in B$, we have $\beta''(A_1 + x) \geq 2\alpha_1$, or

- (2) *there are sets $L \subset B$ and $S_1, \dots, S_k \subset B''$ with $\lambda \gg 1$ and $\sigma_i \geq \alpha^{C\alpha_1^{-1/k}}$ for $1 \leq i \leq k$ such that*

$$L * S_1 * \dots * S_k(x) \leq \alpha_1^{-2} A_1 * A_2 * \dots * A_{k+1}(x)$$

for all $x \in G$, where all convolutions are taken over the measure β' .

Theorem 6.2. *Let $B_\rho(\Gamma)$ and B' be regular rank d Bohr sets such that $B' \subset B_{\epsilon\rho}(\Gamma)$ and suppose $A, L \subset B$ and $S_1, \dots, S_k \subset B'$. Suppose further that $\epsilon \leq c\lambda\alpha/d$ for some sufficiently small absolute constant $c > 0$. Then either*

- (1) *$\langle L * S_1 * \dots * S_k, A \rangle_\beta \gg \lambda \sigma_1 \dots \sigma_k \alpha$, or*
(2) *for any integer l there exists a regular Bohr set $B'' \subset B'$ and an m satisfying*

$$m \ll \lambda^{-2-1/2l} \alpha^{-1/2l} l^2 \log(1/\alpha) \log(1/\sigma_k)$$

with rank $rk(B'') \leq rk(B) + m$, density $\mu_{B'}(B'') \geq (1/dm)^{C(d+m)}$ and, for some $x \in B$, we have $\beta''(A + x) \geq \alpha(1 + c\lambda)$.

The following is the main density increment lemma needed to prove Theorem 1.1, and is proved by combining Theorems 6.1 and 6.2 as in the proof of the analogous Lemma 5.1.

Lemma 6.3. *Let $B_\rho(\Gamma)$ and $B' = B_{\rho'}(\Gamma')$ such that $B' \subset B_{\epsilon\rho}(\Gamma)$. Suppose further that $A_1, A_s \subset B$ and $A_2, \dots, A_{s-1} \subset B'$ all with density at least α . Finally, suppose that $\epsilon \leq c/d \min(\alpha_1, \alpha_s)$ and $\epsilon' \leq c\alpha/d$ for some sufficiently small absolute constant $c > 0$. Then either*

- (1) *$\langle A_1 * A_2 * \dots * A_{s-1}, A_s \rangle_\beta \gg \alpha^{C\alpha_1^{-1/(s-2)}} \alpha_s$, or*
(2) *we have a Bohr set $B'' \subset B'$ such that*

$$rk(B'') \leq rk(B) + O(\alpha_1^{-1/(s-2)} \log^3(1/\alpha_s) \log(1/\alpha)),$$

$$\mu_{B'}(B'') \geq \left(\frac{\alpha_1}{d \log(1/\alpha)} \right)^{C(d + \alpha_1^{-1/(s-2)} \log^3(1/\alpha_s) \log(1/\alpha))}$$

such that for some $x \in B$ and $i \in \{1, s\}$, $\beta''(A_i + x) \geq \alpha(1 + c)$ for some absolute constant $c > 0$.

Theorem 1.1 is then proved by repeated applications of Lemma 6.3, following the strategy of the proof of Theorem 5.1.

ACKNOWLEDGEMENTS

I would to thank my PhD supervisor Trevor Wooley for pointing me towards this area of research and many useful discussions.

REFERENCES

1. M. Bateman and N. H. Katz, *New bounds on cap sets*, arXiv:1101.5851, 2011.
2. J. Bourgain, *On triples in arithmetic progression*, *Geom. Funct. Anal.* **9** (1999), 968–984.
3. ———, *Roth’s theorem on progressions revisited*, *J. Anal. Math.* **104** (2008), 155–206.
4. M.-C. Chang, *A polynomial bound in Freiman’s theorem*, *Duke Math. J.* **113** (2002), 399–419.
5. E. S. Croot and O. Sisask, *A probabilistic technique for finding almost-periods of convolutions*, *Geom. Funct. Anal.* **20** (2010), 1367–1396.
6. F. J. Dyson, *A theorem on the densities of sets of integers*, *J. London. Math. Soc.* **20** (1945), 8–14.
7. B. Green, *Finite field models in additive combinatorics*, *Surveys in Combinatorics 2005*, Cambridge University Press, 2005, London Mathematical Society Lecture Note Series (No. 327).
8. D. R. Heath-Brown, *Integer sets containing no arithmetic progressions*, *J. London Math. Soc.* (2) **35** (1987), 385–394.
9. N. H. Katz and P. Koester, *On additive doubling and energy*, *SIAM J. Discrete Math.* **24** (2010), 1684–1693.
10. Y.-R. Liu and C. V. Spencer, *A generalization of Roth’s theorem in function fields*, *Int. J. Number Theory* **5** (2009), 1149–1154.
11. R. Meshulam, *On subsets of finite abelian groups with no 3-term arithmetic progressions*, *J. Combin. Theory Ser. A* **71** (1995), 168–172.
12. K. F. Roth, *On certain sets of integers*, *J. London Math. Soc.* **28** (1953), 104–109.
13. ———, *On certain sets of integers (II)*, *J. London Math. Soc.* **29** (1954), 20–26.
14. I. Ruzsa, *Solving a linear equation in a set of integers I*, *Acta Arith.* **65** (1993), 259–282.
15. T. Sanders, *On certain other sets of integers*, *J. Anal. Math.* (to appear).
16. ———, *On Roth’s theorem on progressions*, *Annals of Math.* (to appear).
17. T. Schoen and I. Shkredov, *Roth’s theorem in many variables*, arXiv:1106.1601, 2011.
18. E. Szemerédi, *Integer sets containing no arithmetic progressions*, *Acta Math. Hungar.* **56** (1990), 155–158.
19. T. Tao and V. Vu, *Additive combinatorics*, Cambridge Studies in Advanced Mathematics 105, Cambridge University Press, 2006.

THOMAS BLOOM, DEPARTMENT OF MATHEMATICS, UNIVERSITY OF BRISTOL, UNIVERSITY WALK, CLIFTON, BRISTOL BS8 1TW, UNITED KINGDOM
E-mail address: matfb@bristol.ac.uk