ON THE DISTRIBUTION OF THREE-TERM ARITHMETIC PROGRESSIONS IN SPARSE SUBSETS OF $\mathbf{F}_p^n$

HOI H. NGUYEN

Abstract. We give a short proof for the following result on the distribution of three-term arithmetic progressions in sparse subsets of $\mathbf{F}_p^n$: for every $\alpha > 0$ there exists a constant $C = C(\alpha)$ such that the following holds for all $r \geq Cp^{n/2}$ and for almost all sets R of size r of $\mathbf{F}_p^n$. Let A be any subset of R of size at least αr, then A contains a non-trivial three-term arithmetic progression. This is an analog of a hard theorem by Kohayakawa, Łuczak, and Rödl. The proof uses a version of Green's regularity lemma for subsets of a typical random set, which is of interest of its own.

1. Introduction

A well-known result of Roth asserts that for any $\alpha > 0$ there exists $n_0 = n_0(\alpha)$ such that every subset of size αn of $[n]$, where $n > n_0$, contains a non-trivial three-term arithmetic progression. This result was generalized by Brown and Buhler [3] (see also [2]) to any large abelian group of odd order. As a special case, their result implies

Theorem 1.1 (Roth's theorem for $\mathbf{F}_p^n$)**.** *For every $\alpha > 0$ and odd prime p there exists $n_0 = n_0(\alpha, p)$ such that every subset of size αp^n of $\mathbf{F}_p^n$, where $n > n_0$, contains a non-trivial three-term arithmetic progression.*

We say that a subset R of $\mathbf{F}_p^n$ is $(\alpha, 3AP)$-*dense* if any subset of size $\alpha|R|$ of R contains a non-trivial three-term arithmetic progression in $\mathbf{F}_p^n$.

One observes that, by Theorem 1.1, any subset of size ϵp^n, where $\epsilon > 0$ is given and n is sufficiently large, is $(\alpha, 3AP)$-dense. In other words, any subset of positive density in $\mathbf{F}_p^n$ is $(\alpha, 3AP)$-dense. One of the main goal of this note is to show that this phenomenon continues to hold for almost all subsets of size of order $p^{n/2}$.

Theorem 1.2 (Main theorem)**.** *For every $\alpha > 0$ there exists a constant $C = C(\alpha)$ such that the following holds for all $r \geq Cp^{n/2}$. Let R be a random subset of size r of $\mathbf{F}_p^n$, then the probability that R fails to be $(\alpha, 3AP)$-dense is $o(1)$.*

Note that the assumption $r \geq Cp^{n/2}$ is tight up to a constant factor. Indeed, a typical random set of size r in $\mathbf{F}_p^n$ contains about $\Theta(r^3/p^n)$ three-term arithmetic progressions. Thus, if $(1-\alpha)r \gg r^3/p^n$, then there is a subset of size αr of R which does not contain any non-trivial three-term arithmetic progression.

Notation. Here and later, the asymptotic notations, such as o, O, Ω, Θ, and so forth, are used under the assumption that $n \to \infty$ and p is fixed. A notation such as $O_C(.)$ emphasizes

that the hidden constant in O depends on C. If $a = \Omega(b)$, we write $b \ll a$ or $a \gg b$. All logarithms have a natural base, if not specified otherwise.

Theorem 1.2 is an analog of a hard theorem by Kohayakawa, Łuczak, and Rödl [4] on the distribution of three-term arithmetic progressions in $\mathbf{Z}_n$, where n is odd. In fact, we will follow their approach very closely to prove Theorem 1.2. However, with our additional structural result (Theorem 2.3 of Section 2), we are able to get around many technical difficulties to provide a much simpler proof.

The note is organized as follows. In the next two sections 2, 3 we will mainly discuss Green's result for sparse Cayley graphs. Some additional lemmas for proving Theorem 1.2 will be given in Section 4. The proof of Theorem 1.2 is established in Section 5.

2. Green's regularity lemma for (b, σ)-sparse Cayley graphs

Let G be a graph and let A, B be two subsets of V_G. We define the density $d(A, B)$ of $G(A, B)$ to be

$$d(A, B) := e(A, B)/|A||B|.$$

Let ϵ be a positive constant. We say that the pair (A, B) is *ϵ-regular* if

$$|d(A', B') - d(A, B)| \leq \epsilon$$

for any $A' \subset A$ and $B' \subset B$ satisfying $|A'| \geq \epsilon|A|$ and $|B'| \geq \epsilon|B|$.

Szemerédi's regularity lemma, a fundamental result in combinatorics, states that the vertex set of any dense graph can be partitioned into not-too-small pieces so that almost all pairs of pieces are regular.

Theorem 2.1 (Szemerédi's regularity lemma)**.** *Let $\epsilon > 0$. There exists $M = M(\epsilon)$ such that the vertex set can be partitioned into $1/\epsilon \leq m \leq M$ sets V_i with sizes differing by at most 1, such that at least $(1 - \epsilon)m^2$ of the pairs (V_i, V_j) are ϵ-regular.*

Consider a vector space $V = \mathbf{F}_p^n$, where p is a fixed odd prime and n is a large integer. Let A be a subset of V, we define the (bipartite, directed) Cayley graph generated by A to be $G_A = G(V_1, V_2)$, where V_1, V_2 are two copies of V, and $(v_1, v_2) \in E(G_A)$ if $v_2 - v_1 \in A$.

It is clear that G_A is a regular graph of degree $|A|$. Hence if A is dense enough, then Szemerédi's regularity lemma is applicable to G_A. Furthermore, since G_A has additional algebraic structure, it is natural to expect a stronger result than Theorem 2.1. Indeed, a result of Green [1, Section 9] confirms this intuition:

Assume that $|A| = \Omega(|V|)$. Then one can partition $V(G_A)$ into affine subspaces of large dimension and so that almost all pairs of subspaces are ϵ-regular.

Szemerédi's regularity lemma is not meaningful for sparse graphs in general. However, it can be extended to certain graph families. Let ϵ be a positive constant. We say that the pair (A, B) is *relatively ϵ-regular* if

$$|d(A', B') - d(A, B)| \leq \epsilon d(G)$$

for any $A' \subset A$ and $B' \subset B$ satisfying $|A'| \geq \epsilon|A|$ and $|B'| \geq \epsilon|B|$.

Let be given $b > 2$ and $\sigma > 0$. We say that a graph G is (b, σ)-*sparse* if

$$d(X, Y) \leq bd(G)$$

for any $|X| \geq \sigma|V_G|$ and $|Y| \geq \sigma|V_G|$. The following result extends Szemererédi' s regularity lemma for (b, σ)-sparse graphs.

Theorem 2.2 (Szemerédi's regularity lemma for sparse graphs, [4, Lemma 4])**.** *Let $b > 0$. For $\epsilon > 0$ there exists $\sigma = \sigma(b, \epsilon)$ such that the following holds for all (b, σ)-sparse graphs. There exists $M = M(\epsilon, b)$ such that the vertex set can be partitioned into $1/\epsilon \leq m \leq M$ sets V_i with sizes differing by at most 1, such that at least $(1-\epsilon)m^2$ of the pairs (V_i, V_j) are relatively ϵ-regular.*

As to how Theorem 2.2 extends Theorem 2.1, our result below shows that the result of Green can be extended easily to "(b, σ)-sparse" Cayley graphs in $\mathbf{F}_p^n$. We first need the notion of regular sets.

σ-regular set. Let σ be a positive constant. We say that a subset R of V is *σ-regular* if the number of edges between X and Y in the Cayley graph G_R is bounded by,

$$e_{G_R}(X, Y) \leq 2|R||X||Y|/N,$$

for all X, Y satisfying $|X|, |Y| \geq \sigma N$, where $N := |V|$.

We will also need a tower-type function $W(t)$ defined recursively below.

$$W(1) = 2p \text{ and } W(t) := (2p)^{W(t-1)}.$$

We now state our structural result.

Theorem 2.3 (Green's regularity lemma for sparse Cayley graphs)**.** *Let a prime p be fixed. For all $0 < \alpha < 1$ and $0 < \epsilon < 1$ there exist σ and n_0 such that for all $n > n_0$ the following holds. Let R be a σ-regular subset of $\mathbf{F}_p^n$ and let A be a subset of R satisfying $|A| \geq \alpha|R|$. Then, there is a subspace H of $\mathbf{F}_p^n$ of index $\leq W(4\epsilon^{-9}\alpha^{-2})$ such that the partition of V into the affine translates of H, $V = \cup_1^K H_i$, has at least $(1-\epsilon)K^2$ relatively ϵ-regular pairs (H_i, H_j) in G_A.*

This theorem will play a key ingredient to the establishment of Theorem 1.2. To prove it, we first pass to a simpler version to be discussed below.

Fourier transform.(cf. [5, Chapter 4.]) Let H be a subspace of V, let f be a real-valued function defined on V. Then the Fourier transform of f with respect to H is

$$\widehat{f}(\xi) := \mathbf{E}_{x\in H} f(x) e(-\langle x, \xi\rangle),$$

where $\langle x, \xi\rangle = \sum_{i=1}^{n} x_i\xi_i/p$, and $e(z) = e^{2\pi i z}$.

Convolution. Let f and g be two real-valued functions defined on V. The convolution of f and g with respect to H is

$$f * g(h) := \mathbf{E}_{x\in H} f(x) g(h-x).$$

The following basic properties for real-valued functions will be used several times.

- (Parseval's identity) $\mathbf{E}_{x\in H} f^2(x) = \sum_{\xi\in H} |\widehat{f}(\xi)|^2$.
- (Plancherel's formula) $\mathbf{E}_{x\in H} f(x)g(x) = \sum_{\xi\in H} \widehat{f}(\xi)\overline{\widehat{g}(\xi)}$.
- (Fourier inversion formula) $f(x) = \sum_{\xi\in H} \widehat{f}(\xi) e(\langle x, \xi\rangle)$ for any $x \in H$.
- $\widehat{f * g}(\xi) = \widehat{f}(\xi)\widehat{g}(\xi)$.

Let A be a subset of V, and let v be an element of V. We define A_H^v to be the set $A + v \cap H$. Sometimes we also write A_H^v as its characteristic function. Following are some simple properties:

- $\widehat{A_H^v}(\xi) = |A_H^v|/|H|$ if $\xi \in H^\perp$;
- $\widehat{A_H^{v'}}(\xi) = e(\langle v - v', \xi\rangle)\widehat{A_H^v}(\xi)$ if $v - v' \in H$; in particular, $|\widehat{A_H^{v'}}(\xi)| = |\widehat{A_H^v}(\xi)|$.

ϵ-regular vector. Let ϵ be a positive constant. Let A be a given set. We say that a vector v is *ϵ-regular for A* with respect to H if

$$\sup_{\xi\notin H^\perp} |\widehat{A_H^v}(\xi)| \le \epsilon|A|/N.$$

(It is more natural to use the upper bound $\epsilon|A_H^v|/|H|$ in the definition above, but we find our definition more convenient to use, and $\epsilon|A|/N|$ is the average value for $\epsilon|A_H^v|/|H|$.)

Notice that if v is an ϵ-regular vector, then so is any element of $v + H$.

We say that a subspace H is **ϵ-regular for** A if the number of v's which fail to be ϵ-regular is at most ϵN.

We are now ready to state an analog of Theorem 2.3.

Theorem 2.4 (Green's regularity lemma for sparse sets)**.** *Let a prime p be fixed. For all $0 < \alpha < 1$ and $0 < \epsilon < 1$ there exist σ and n_0 such that for all $n > n_0$ the following holds. Let R be a σ-regular subset of $\mathbf{F}_p^n$ and let A be a subset of R satisfying $|A| \geq \alpha|R|$. Then there is a subspace H of $\mathbf{F}_p^n$ of index $\leq W(4\epsilon^{-3}\alpha^{-2})$ which is ϵ-regular for A.*

For the rest of this section, we deduce Theorem 2.3 from Theorem 2.4.

Let $V = \cup_{i=1}^K H_i$ be the partition of V into affine translates of H. Let $v_1, \ldots, v_K$ be representatives of the coset subgroups V/H. Then by definition, all but at most ϵK vectors $v_1, \ldots, v_K$ are ϵ-regular vectors with respect to H.

Next, assume that $H_i = v_i + H$ and $H_j = v_j + H$ are two affine translates of H such that $v_j - v_i$ is an ϵ-regular vector. We will show that the subgraph $G_A(H_i, H_j)$ is relatively $\epsilon^{1/3}$-regular.

It is clear that $e_{G_A}(H_i, H_j) = |H||A_H^{v_j - v_i}|$; thus

$$d_{G_A}(H_i, H_j) = |A_H^{v_j - v_i}|/|H|.$$

Let $X \subset H_i$ and $Y \subset H_j$ be any two subsets of H_i and H_j respectively, which satisfy $|X|, |Y| \geq \epsilon^{1/3}|H|$. We shall estimate the number of edges generated by X and Y. We have

$$\begin{aligned} e_{G_A}(X, Y) &= \sum_{x \in H_i, y \in H_j} A(y - x)X(x)Y(y) \\ &= \sum_{x', y' \in H} A_H^{v_j - v_i}(y' - x')X(x' + v_i)Y(y' + v_j) \\ &= \sum_{x', y' \in H} A_H^{v_j - v_i}(y' - x')\,(X - v_i)\,(x')(Y - v_j)(y') \end{aligned}$$

Now we apply the Fourier inversion formula to the last sum to obtain (after canceling some zero sums)

$$\begin{aligned} e_{G_A}(X, Y) &= |H|^2 \sum_{\xi \in H} \widehat{A_H^{v_j - v_i}}(\xi)(\widehat{X - v_i})(\xi)(\widehat{Y - v_j})(-\xi) \\ &= |A_H^{v_j - v_i}||X||Y|/|H| + \sum_{\xi \in H \setminus \{0\}} \widehat{A_H^{v_j - v_i}}(\xi)(\widehat{X - v_i})(\xi)(\widehat{Y - v_j})(-\xi). \end{aligned}$$

Since $v_j - v_i$ is an ϵ-regular vector with respect to H, we infer that

$$\left|e_{G_A}(X,Y) - |A_H^{v_j-v_i}||X||Y|/|H|\right| \le (\epsilon|A^{v_j-v_i}|/N)\sum_{\xi}\left|(\widehat{X-v_i})(\xi)(\widehat{Y-v_j})(-\xi)\right|.$$

By Parseval's identity and by the Cauchy-Schwarz inequality we thus have

$$\begin{aligned}\left|e_{G_A}(X,Y) - |A_H^{v_j-v_i}||X||Y|/|H|\right| &\le |H|(\epsilon|A^{v_j-v_i}|/N)(|X||Y|)^{1/2}\\ &\le \epsilon|A^{v_j-v_i}||H|^2/N.\end{aligned}$$

It follows that

$$\begin{aligned}|d_{G_A}(X,Y) - d_{G_A}(H_i,H_j)| &\le \epsilon|A^{v_j-v_i}||H|^2/(|X||Y|N)\\ &\le \epsilon^{1/3}|A^{v_j-v_i}|/N\\ &= \epsilon^{1/3}d(G_A).\end{aligned}$$

Hence $G_A(H_i,H_j)$ is indeed relatively $\epsilon^{1/3}$-regular.

One observes that $v_j - v_i$ is an ϵ-regular vector for all but at most ϵK^2 pairs (i,j). Hence there are $(1-\epsilon)K^2$ relatively $\epsilon^{1/3}$-regular pairs (H_i,H_j) in G_A. This concludes our deduction of Theorem 2.3 after modifying ϵ to $\epsilon^{1/3}$.

Remark 2.5. It is crucial to note that the definition of ϵ-regular vector above works for any type of (linear) Cayley graph. For instance, assume that $(v_1+v_2)/2$ is an ϵ-regular vector with respect to H and define a bipartite Cayley graph G'_A on $(H-v_1, H-v_2)$ by connecting (h_1-v_1) with (h_2-v_2) if $((h_1-v_1)+(h_2-v_2))/2 = (h_1+h_2)/2-(v_1+v_2)/2 \in A$. Then G'_A is also $\epsilon^{1/3}$-regular. Indeed, it is clear that

$$d_{G'_A}(H_1,H_2) = |A_H^{(v_1+v_2)/2}|/|H|.$$

Next, by the Fourier inversion formula,

$$\begin{aligned}e_{G'_A}(X,Y) &= \sum_{x',y'\in H} A_H^{(v_1+v_2)/2}((y'-x')/2)X(x'+v_1)Y(y'+v_2)\\ &= \sum_{x',y'\in H} A_H^{(v_1+v_2)/2}((y'-x')/2)\,(X-v_1)\,(x')(Y-v_2)(y')\\ &= |H|^2\sum_{\xi\in H}\widehat{A_H^{(v_1+v_2)/2}}(2\xi)(\widehat{X-v_i})(\xi)(\widehat{Y-v_j})(-\xi)\\ &= |A_H^{(v_1+v_2)/2}||X||Y|/|H| + \sum_{\xi\in H\setminus\{0\}}\widehat{A_H^{(v_1+v_2)/2}}(2\xi)(\widehat{X-v_i})(\xi)(\widehat{Y-v_j})(-\xi).\end{aligned}$$

Thus, because $(v_1 + v_2)/2$ is an ϵ-regular vector with respect to H, we also have

$$\left| e_{G_A}(X,Y) - |A_H^{v_j - v_i}||X||Y|/|H| \right| \leq \epsilon |A^{(v_1+v_2)/2}||H|^2/N.$$

It then follows that

$$\begin{aligned} |d_{G'_A}(X,Y) - d_{G'_A}(H_1,H_2)| &\leq \epsilon |A^{(v_1+v_2)/2}||H|^2/(|X||Y|N) \\ &\leq \epsilon^{1/3}|A^{(v_1+v_2)/2}|/N \\ &= \epsilon^{1/3} d(G'_A), \end{aligned}$$

completing the remark.

3. Proof of Theorem 2.4

Define $d(A,H)$ by

$$d(A,H) := \frac{1}{N}\sum_{v\in V}\left(\frac{|A_H^v|}{H}\right)^2 / \left(\frac{|A|}{N}\right)^2.$$

Observe that $d(A,H)$ is the mean of the squares of the normalized densities of the $G_A(H_i,H_j)$'s. We show that this quantity is always bounded.

Claim 3.1. We have $d(A,H) \leq 4/\alpha^2$ for any $|H| \geq \sigma N$.

Proof. (of Claim 3.1). Since $H \geq \sigma N$, by the σ-regularity of R, for any v we have,

$$|H||R_H^v| = e_{G_R}(H, H-v) \leq 2|H||H||R|/N.$$

Hence $|A_H^v|/|H| \leq |R_H^v|/|H| \leq 2|R|/N \leq (2/\alpha)|A|/N$. As a result,

$$d(A,H) \leq \frac{1}{N}\sum_{v\in V}(2/\alpha)^2 \leq 4/\alpha^2.$$

□

As in the proof of Szemerédi's regularity lemma, when a partition with too many irregular pairs comes into play, then we pass to a finer partition, and by so the mean square of the densities will increase. What we are going to do is similar, the only difference is we restrict ourselves to a special family of partitions.

Lemma 3.2. *Let $\epsilon \in (0,1)$ and suppose that H is a subspace of V, which is not ϵ-regular for A. Then there is a subspace $H' \leq H$ such that $|V/H'| \leq (2p)^{|V/H|}$ and $d(A,H') \geq d(A,H) + \epsilon^3$.*

Proof. (of Lemma 3.2). Since H is not ϵ-regular for A, there are ϵN vectors v such that $\sup_{\xi \notin H^\perp} |\widehat{A_H^v}(\xi)| \geq \epsilon |A|/N$. In other words, there exists a positive integer m satisfying $\epsilon N/|H| \leq m \leq N/|H|$ together with m coset representatives $v_1, \dots, v_m \in V/H$ and vectors $\xi_1, \dots, \xi_m \in H^\perp$, such that

$$|\widehat{A_H^{v_i}}(\xi_i)| \geq \epsilon |A|/N.$$

Now let $H' \subset H$ be the annihilator of all ξ_i's. It is clear that

$$|H'| \geq |H|/p^m \geq |H|/p^{|V/H|}$$

Hence,

$$|V/H'| \leq |V/H| p^{|V/H|} < (2p)^{|V/H|}.$$

Set $S := N|H'|^2 (|A|/N)^2 |H| d(A,H')$. It is obvious that

$$S = |H| \sum_{v \in V} |A_{H'}^v|^2 = \sum_{v \in V, h \in H} |A_{H'}^{v+h}|^2.$$

Notice that $|A_{H'}^{v+h}| = \sum_{x \in H} (A+v)(x-h)H'(x) = \sum_{x \in H} (A+v)(x) H'(x+h) = |H| (A_H^v * H')(-h)$. We rewrite S and then use Plancherel's formula,

$$\begin{aligned} S &= |H|^2 \sum_{v \in V, h \in H} |A_H^v * H'(h)|^2 \\ &= |H|^3 \sum_{v \in V, \xi \in H} \left| \widehat{A_H^v * H'}(\xi) \right|^2 \\ &= |H|^3 \sum_{v \in V, \xi \in H} |\widehat{A_H^v}(\xi)|^2 |\widehat{H'}(\xi)|^2. \end{aligned}$$

In the last sum, the contribution of the $\xi = 0$ term gives

$$\begin{aligned} S_0 &= |H|^3 \sum_{v\in V} \left(|A_H^v|/|H|\right)^2 \left(|H'|/|H|\right)^2 \\ &= |H||H'|^2 \sum_{v\in V} \left(|A_H^v|/|H|\right)^2 \\ &= N|H||H'|^2(|A|/N)^2 d(A,H); \end{aligned}$$

while the sums contributed from $\xi \in H\backslash\{0\}$ is bounded from below by

$$S_{\neq 0} \geq |H|^3 \sum_{i=1}^{m} \sum_{v\in H+v_i} |\widehat{A_H^{v_i}}(\xi_i)|^2 |\widehat{H'}(\xi_i)|^2.$$

But since $\xi_i \in H'^{\perp}$, we have $\widehat{H'}(\xi_i) = |H'|/|H|$. Use the bound $|\widehat{A}_H^{v_i}(\xi_i)| \geq \epsilon|A|/N$ for all $1 \leq i \leq m$, we obtain

$$\begin{aligned} S_{\neq 0} &\geq |H|^3 m|H|(\epsilon|A|/N)^2 \left(|H'|/|H|\right)^2 \\ &\geq |H|^3(\epsilon N/|H|)|H|\epsilon^2 \left(|A|/N\right)^2 \left(|H'|/|H|\right)^2 \\ &= \epsilon^3|H||H'|^2 N \left(|A|/N\right)^2. \end{aligned}$$

From the estimate for S_0 and $S_{\neq 0}$ we conclude that $d(A,H') \geq d(A,H) + \epsilon^3$.

□

To complete the proof of Theorem 2.4 we keep applying Lemma 3.2. Since $d(A,H) \leq 4/\alpha^2$, the iteration stops after at most $4\epsilon^{-3}\alpha^{-2}$ steps. During the iteration, $|H'|$ is always bounded below by $N/W(4\epsilon^{-3}\alpha^{-2})$, thus we may choose $\sigma = (2W(4\epsilon^{-3}\alpha^{-2}))^{-1}$.

Let us conclude this section by an important corollary to the proof of Theorem 2.4.

Theorem 3.3. *Let $\alpha, \epsilon \in (0,1)$ and let m be a positive integer. There is a constant $\sigma = \sigma(\epsilon,\alpha,m)$ such that if R is a σ-regular set of V and A is a subset of R of cardinality $\alpha|R|$, then the following holds. Assume that $A = \cup_{i=1}^m A_i$ is a partition of A into m distinct sets of size $|A|/m$. Then there is a subspace $H \leq V$ of index bounded by $W(4m^3\epsilon^{-3}\alpha^{-2})$ which is ϵ-regular for all A_i's.*

To prove Theorem 3.3 we first let $d(A_1,\ldots,A_m,H) := \sum_{i=1}^m d(A_i,H)$. Next, keep iterating Lemma 3.2 if H is not ϵ-regular for some A_i. Since $d(A_1,\ldots,A_m,H) \leq \sum_{i=1}^m 4/(\alpha/m)^2 = 4m^3/\alpha^2$, the iteration will stop after at most $4m^3\epsilon^{-3}\alpha^{-2}$ steps.

4. Main lemmas for applications

In this section we will provide some important supporting lemmas to be used in the proof of Theorem 1.2.

4.1. **Regularity of a random set.**

Lemma 4.2. *For $0 < \sigma < 1/2$ there is a constant $C(\sigma)$ such that if $r \geq C(\sigma)N^{1/2}$ and R is a random subset of size r of V, then R is a σ-regular set asymptotically almost surely.*

To start with, we consider a slightly different model as follows.

Lemma 4.3. *For $\sigma > 0$ there is a constant $C(\sigma)$ such that if $r \geq C(\sigma)N^{1/2}$ and $q = r/N$, and R is a random subset of V whose elements are equally selected with probability q, then the following holds asymptotically almost surely for R: $e_R(X,Y) \leq 1.5|R||X||Y|/N$ for all $|X|, |Y| \geq \sigma N$.*

Proof. (of Lemma 4.3) Let $X, Y \subset V$, of cardinality at least σN. The number of edges of G_R generated by X and Y is

$$e_R(X,Y) = \sum_{x,y \in V} 1_R(y-x)1_X(x)1_Y(y) = N^2 \sum_{\xi \in V} \widehat{1_R}(\xi)\widehat{1_X}(\xi)\widehat{1_Y}(-\xi)$$

where the Fourier transform is defined with respect to V, and the latter identity comes from Fourier inversion formula. Thus we have

$$e_R(X,Y) = |R||X||Y|/N + N^2 \sum_{\xi \in V, \xi \neq 0} \widehat{1_R}(\xi)\widehat{1_X}(\xi)\widehat{1_Y}(-\xi).$$

Let us pause to estimate $\widehat{1_R}(\xi)$.

Lemma 4.4. $\sup_{\xi \neq 0} |\widehat{1_R}(\xi)| < |R|/(N \log N)$ *asymptotically almost surely for R.*

The proof of this lemma is routine by applying the exponential moment method. For the sake of completeness, we prove it in Appendix A.

Assuming Lemma 4.4, then by the Cauchy-Schwarz inequality and Parseval's identity we have

$$|e_R(X,Y) - |R||X||Y|/N| \le N^2 \sup_{\xi\neq 0} |\widehat{1_R}(\xi)| \left(\sum_{\xi\in V} |\widehat{1_X}(\xi)|^2 \sum_{\xi\in V} |\widehat{1_Y}(\xi)|^2 \right)^{1/2}$$
$$\le N^2 \sup_{\xi\neq 0} |\widehat{1_R}(\xi)| \left(|X||Y|/N^2\right)^{1/2}$$
$$= \sup_{\xi\neq 0} |\widehat{1_R}(\xi)| (|X||Y|)^{1/2} N.$$

On the other hand, as $|X|, |Y| \ge \sigma N$ and $\sup_{\xi\neq 0} |\widehat{1_R}(\xi)| \le |R|/(N \log N)$, we have

$$\sup_{\xi\neq 0} |\widehat{1_R}(\xi)| (|X||Y|)^{1/2} N = o\left(|R||X||Y|/N\right),$$

completing the proof of Lemma 4.3.

□

Next we show that the two models, of Lemma 4.2 and of Lemma 4.3, are similar.

Proof. (of Lemma 4.2). Let $q = (1-\sigma^4)|R|/N$. We first consider a random set R_1 by selecting each element of V with probability q. It is obvious that the size of this random set belongs to $[(1-2\sigma^4)|R|, |R|]$ asymptotically almost surely. We restrict ourself to this event by renormalizing the probability space. Hence the random set R_1 is chosen uniformly from the collection of subsets of size $[(1-2\sigma^4)|R|, |R|]$. Next we pick uniformly a set R_2 of size $|R| - |R_1|$ from $V \backslash R_1$ and set $R = R_1 \cup R_2$.

Suppose that $X, Y \subset V$ and $|X|, |Y| \ge \sigma N$. By Lemma 4.3, we have $e_{R_1}(X,Y) \le 1.5|R_1||X||Y|/N \le 1.5|R||X||Y|/N$. On the other hand, it is obvious that

$$e_R(X,Y) \le e_{R_1}(X,Y) + |R_2|N$$
$$\le 1.5|R||X||Y|/N + 2\sigma^4|R|N$$
$$\le 2|R||X||Y|/N.$$

Hence $e_R(X,Y) \le 2|R||X||Y|/N$ asymptotically almost surely, completing the proof of Lemma 4.2. □

4.5. **Edge distribution of quasi-random graphs.** Roughly speaking, if we choose randomly a large number of vertices of a dense quasi-random graph, then the chance of obtaining an edge is very high. This simple observation, as a strong tool to exploit structure for counting, was used in [4], and will play a key role in our proof of Theorem 1.2.

Let $G = G(u, \rho, \epsilon)$ be an ϵ-regular bipartite graph, $V(G) = U_1 \cup U_2$, where $|U_1| = |U_2| = u$ and $d(G) = e(G)/u^2 = \rho$. Let $t_1, t_2 < u/2$ be two given positive integers. We select a random subgraph of G as follows. First, an adversary chooses a set $S_1 \subset U_1$ with $|S_1| \leq u/2$. Then we pick a set $T_1 \subset U_1 \backslash S_1$ with $|T_1| = t_1$ from the collections of all t_1-subsets of $U_1 \backslash S_1$ with equal probability. Next, our adversary picks a set $S_2 \subset U_2$ with $|S_2| \leq u/2$, and we pick a set $T_2 \subset U_2 \backslash S_2$ with $T_2 = t_2$ from the collections of all t_2-subsets of $U_2 \backslash S_2$ with equal probability. Let us call the outcome of the above procedure a random (t_1, t_2)-subgraph of G.

Lemma 4.6. [4, Lemma 11] *For every constant $0 < \eta < 1$, there exist a constant $0 < \epsilon < 1$ and a natural number u_0 such that, for any real $t \geq 2(u/\epsilon)^{1/2}$ and any given graph $G = G(u, \rho, \epsilon)$ as above with $u \geq u_0$ and $\rho \geq t/u$, the following holds. If $t_1, t_2 \geq t$, regardless of the choices for S_1 and S_2 of our adversary, the probability that a random (t_1, t_2)-subgraph of G fails to contain an edge is at most η^t.*

The proof of Lemma 4.6 is simple, the interested reader may read [4].

4.7. **Roth's theorem for $\mathbf{F}_p^n$.** Another important ingredient is the following (equivalent) form of Theorem 1.1.

Theorem 4.8. *For any $\delta > 0$ there is a number $c(\delta) > 0$ such that if B is a subset of V of size $\delta|V|$, then B contains at least $c(\delta)|V|^2$ three-term arithmetic progressions.*

In the next section, we shall put every thing together to establish Theorem 1.2.

5. Proof of Theorem 1.2

We shall work with several constants throughout this section, so let us mention briefly here to avoid confusion.

$$\alpha, c(\alpha) \to \eta \to \epsilon \to \sigma \to C$$

Firstly, α is the constant that we fix all the time. The constants $c(\alpha)$'s depend only on α. Secondly, η will be chosen to be small enough depending on α and the $c(\alpha)$'s. The constant ϵ will depend on α and η. Last but not least, σ depends on α and ϵ. We shall choose η, ϵ, σ to be small enough, while the constants $C = C(\alpha, \eta, \epsilon, \sigma)$ are often very large.

Now we discuss the proof of Theorem 1.2 in details.

By Theorem 4.8, it is enough to work with the case

$$r = o_\alpha(N).$$

We say that a set A is (α, σ)-*bad* if it contains no non-trivial three-term arithmetic progression and there exists a σ-regular set R such that $A \subset R$ and $|A| = \alpha|R| = \alpha r$.

Our main goal is to give an upper bound for the number of bad sets of a given size.

Theorem 5.1. *For all $\alpha > 0$ there exists $c > 0$ such that for all $\eta > 0$ there exist $C > 0$ and $\sigma > 0$ such that for all $s \geq C(\alpha, \eta)N^{1/2}$, the number of (α, σ)-bad sets of size s is at most $\eta^{c(\alpha)s}\binom{N}{s}$.*

Proof. (of Theorem 1.2 assuming Theorem 5.1). We choose $\eta = \eta(\alpha)$ to be small enough. Let $s \geq C(\alpha, \eta)N^{1/2}$ and put $r = s/\alpha$. Pick a random set R among all r-subsets of V. Then by Theorem 4.2, R is σ-regular almost surely. Among these σ-regular r-sets, by Theorem 5.1, the number of sets that contain at least one (α, σ)-bad subset is at most

$$\eta^{c(\alpha)s}\binom{N}{s}\binom{N-s}{r-s}.$$

Observe that, as η is small enough, this amount is $o\left(\binom{N}{r}\right)$. Hence almost all r-sets of V contain no bad subsets at all. To finish the proof, we note that if R contains no (α, σ)-bad subset, then it is $(\alpha, 3AP)$-dense.

□

We shall concentrate on proving Theorem 5.1 by localizing some properties of A. Our approach follows that of [4] closely, but the key difference here is that we shall exploit the rich structure obtained from Theorem 2.3 and Theorem 3.3.

Let R be a σ-regular set of fixed size $C(\sigma)N^{1/2} \leq r = o(N)$ such that $A \subset R$. Let $m = m(\alpha)$ be a large number to be defined later.

From now on we shall view A as an ordered m-set-tuple, $A = (A_1, \ldots, A_m)$, where $|A_i| = |A|/m$ for all i and $A = \cup A_i$. We shall choose $\epsilon = \epsilon(\alpha)$ to be small enough. By Theorem 3.3, there exists a subspace H of V which has index bounded by $W(4m^3\alpha^{-2}\epsilon^{-9})$ and which is ϵ^3-regular for all A_i's .

Let $v_1, \ldots, v_K$ be representatives of the quotient space $V' := V/H$. For each A_i, let us consider a set B_i of vectors $v = v_j$ that satisfy the following conditions:

- v is ϵ^3-regular with respect to A_i and H.
- $|(A_i)_H^v| \geq (1/4)|A_i||H|/N$.

It is clear that $|(A_i)_H^v| \leq |A_H^v| \leq |R_H^v|$. But by definition of R, $|R_H^v| \leq 2|R||H|/N$; thus we have

$$\begin{aligned}\sum_{v \in B_i} |(A_i)_H^v| &\geq |A_i| - (\epsilon K)(2|R||H|/N) - K((1/4)|A_i||H|/N) \\ &\geq (1 - (\epsilon m)/\alpha - 1/4)|A_i| \geq |A_i|/2,\end{aligned}$$

provided that $\epsilon \leq \alpha/4m$. We infer that the size of B_i is large,

$$|B_i| \geq (|A_i|/2)/(2|R||H|/N) \geq \frac{\alpha}{4m}K.$$

By a truncation if needed, we assume that B_i has cardinality $(\alpha/4m)K$ for all i. Notice that these sets are not necessarily disjoint. We shall show that there are many three-term arithmetic progressions (in V') with the property that all 3 terms belong to different B_i's.

Now we set $B := \{v \in V' : v \in B_i \cap B_j \cap B_k \text{ for some } i < j < k\}$ and consider the following two cases.

Case 1. $|B| \geq (\alpha/8m)K = (\alpha/8m)|V'|$. Applying Theorem 4.8 we obtain $c(\alpha/8m)K^2$ three-term arithmetic progressions in B. By the definition of B, it follows that there are $c(\alpha/8m)K^2$ three-term arithmetic progressions with the property that all three terms belong to three different sets B_i.

Case 2. $|B| \leq (\alpha/8m)K = |B_i|/2$. We let $B' = \cup_{i=1}^m B_i \backslash B$. By an elementary counting argument, it follows that $|B'| \geq m|B_i|/4 = (\alpha/16)K$. Let us write $B' = \cup_{i=1}^m B'_i$, where $B'_i \subset B_i$ and all B'_i are disjoint.

By Theorem 4.8, the set B' contains $c(\alpha/16)K^2$ three-term arithmetic progressions. Among them, since each three-term arithmetic progression is defined by two parameters, the number of three-term arithmetic progressions that consist of at least two terms from the same B'_i is bounded by $3\sum_{i=1}^m |B'_i|^2$. The latter quantity is bounded by $3|B_i|(\sum_1^m |B_i|) \leq 3(\alpha/4m)(\alpha/4)K^2$; which is negligible compared to $c(\alpha/16)K^2$ by letting $m = m(\alpha)$ large.

In both cases, the number of three-term arithmetic progressions with the property that all three terms belong to three different sets B_i is at least $c'(\alpha)K^2$, where

$$c'(\alpha) = \min(c(\alpha/8m), c(\alpha/16)/2).$$

By an averaging argument, there exist three indices $i_0 < j_0 < k_0$ such that the number of three-term arithmetic progressions in $B_{i_0} \times B_{j_0} \times B_{k_0}$ is at least $c'(\alpha)K^2/m^3 = c''(\alpha)K^2$. In particular, there exist a vector $v_{i_0} \in B_{i_0}$ and $c''(\alpha)K$ pairs $(v^l_{j_0}, v^l_{k_0}) \in B_{j_0} \times B_{k_0}$ such that each triple $(v_{i_0}, v^l_{j_0}, v^l_{k_0})$ is a three-term arithmetic progression.

Let us summarize what have been achieved.

(1) There exists a subspace H of index bounded by a function of α and ϵ, and there exist $A_{i_0}, A_{j_0}, A_{k_0}$ and triples $(v_{i_0}, v^l_{j_0}, v^l_{k_0})$, where $1 \leq l \leq c''(\alpha)K$, such that the following holds:

(2) v_{i_0} is an ϵ^3-regular vector for A_{i_0}, and $|(A_{i_0})_H^{v_{i_0}}| \geq (1/4)|A_{i_0}||H|/N$;

(3) $|(A_{j_0})_H^{v^l_{j_0}}| \geq (1/4)|A_{j_0}||H|/N = (1/4m)s|H|/N$, where we recall that $s = |A| = \alpha r$;

(4) $|(A_{k_0})_H^{v_{k_0}^l}| \geq (1/4)|A_{k_0}||H|/N = (1/4m)s|H|/N$;

(5) $|(v_{i_0}, v_{j_0}^l, v_{k_0}^l)|$ is a three-term arithmetic progression in V/H.

One also observes that $v_{j_0}^l, v_{k_0}^l$ are ϵ^3-regular vectors with respect to A_{j_0} and A_{k_0}; but we do not need this fact. Let us call this configuration an $\left(\alpha, \epsilon^3, H, i_0, j_0, k_0, v_{i_0}, v_{j_0}^l, v_{k_0}^l\right)$-*flower*. Roughly speaking, the reader may visualize a flower with a center $A_{i_0} + v_{i_0} \cap H$, where $A_{i_0} + v_{i_0} \cap H$ sits nicely in H, and with $c''(\alpha)K$ petals $(A_{j_0} + v_{j_0}^l \cap H, A_{k_0} + v_{k_0}^l \cap H)$.

We now show that the probability that a random set A of size s (which is chosen by first choosing a random set A_1 of size s/m from V, and then a random set A_2 of size s/m from $V \backslash A_1$, and so on), with the properties above, fails to contain a three-term arithmetic progression is very small.

Proposition 5.2. *Let α, η be given. Then there exist constants $c = c(\alpha) > 0, C = C(\alpha, \eta)$ and $\epsilon = \epsilon(\alpha, \eta) > 0$ such that the probability that a random set A of size s, viewed as an ordered m-set-tuples $(A_1, \ldots, A_m)$, where $s \geq C(\alpha, \eta)N^{1/2}$, that contains a flower but not any non-trivial three-term arithmetic progression is at most $\eta^{c(\alpha)s}$.*

It is clear that Theorem 5.1 follows from Proposition 5.2. Therefore it suffices to prove Proposition 5.2.

First, we shall estimate the probability that A has no non-trivial three-term arithmetic progressions but contains a given $\left(\alpha, \epsilon^3, H, i_0, j_0, k_0, v_{i_0}, v_{j_0}^l, v_{k_0}^l\right)$-flower. We will condition on A_{i_0} and the size t_{1l} of $A_{j_0}^{v_{j_0}^l} \cap H$ and t_{2l} of $A_{k_0}^{v_{k_0}^l} \cap H$ for all l. Thus, v_{i_0} is an ϵ^3-regular vector with respect to a fixed set A_{i_0}; and the sets A_{j_0} and A_{k_0} vary in such a way that all $v_{j_0}^l$ and $v_{k_0}^l$ satisfy (3) and (4), and A_{j_0}, A_{k_0} intersect $H - v_{j_0}^l, H - v_{k_0}^l$ in sets of size t_{1l} and t_{2l} respectively.

Set

$$S_1^l := \bigcup_{1 \leq m < j_0} (A_m \bigcap (H - v_{j_0}^l)); \quad T_1^l := A_{j_0} \bigcap (H - v_{j_0}^l),$$

and

$$S_2^l := \bigcup_{1 \leq m < k_0} (A_m \bigcap (H - v_{k_0}^l)); \quad T_2^l := A_{k_0} \bigcap (H - v_{k_0}^l).$$

Without loss of generality, we assume that $2v_{i_0} = v_{j_0}^l + v_{k_0}^l$. Define a Cayley graph between $H - v_{j_0}^l$ and $H - v_{k_0}^l$ by connecting $v_1 \in H - v_{j_0}^l$ to $v_2 \in H - v_{k_0}^l$ if $(v_1 + v_2)/2 \in A_{i_0}$. Since v_{i_0} is ϵ^3-regular with respect to A_{i_0}, by the observation made in Remark 2.5, this Cayley graph is ϵ-regular. (Note that for other cases such as $v_{i_0} = 2v_{j_0}^l - v_{k_0}^l$, we can define a

similar linear Cayley graph: we connect $v_1 \in H - v^l_{j_0}$ to $v_2 \in H - v^l_{k_0}$ if $2v_1 - v_2 \in A_{i_0}$. Again, due to the argument given in Remark 2.5, this new graph is also ϵ-regular.)

Recall that $i_0 < j_0 < k_0$, and we expose the sets A_i in order. In the ϵ-regular graph defined above, the sets $S^l_1, T^l_1, S^l_2, T^l_2$ play the role of S_1, T_1, S_2, T_2 of Lemma 4.6.

By choosing $\epsilon = \epsilon(\alpha, \eta) = \epsilon(\alpha)$ to be small enough and $C = C(\alpha, \eta)$ to be large enough, one may check easily by using (3) and (4) that for each bipartite graph $(H - v^l_{j_0}, H - v^l_{k_0})$, the assumptions of Lemma 4.6 are satisfied. For instance, by (3), the random set T^l_1 has size $t_{1l} \geq (s/4m)|H|/N$, which is greater that $2(|H|/\epsilon)^{1/2}$ if C is sufficiently large. Also, as $|S^l_1|, |S^l_2| \leq s = o(N) = o(|H|)$, the conditions $|S^l_1|, |S^l_2| \leq |H|/2$ are satisfied automatically.

It then follows from Lemma 4.6 that the probability each petal fails to contain a non-trivial three-term arithmetic progression is less than $\eta^{(1/4)s/(mK)}$. Hence, because there are $c''(\alpha)K$ petals, the probability that A contains no non-trivial three-term arithmetic progression is bounded by $\eta^{(1/4)c''(\alpha)s/m} = \eta^{c'''(\alpha)s}$.

Now we bound the number of flowers. As the number of choices for H is bounded by $N^{W(4m^3\alpha^{-2}\epsilon^{-9})}$, and the number of choices for $(i_0, j_0, k_0, v_{i_0}, v^l_{j_0}, v^l_{k_0})$ is bounded by $K^{4+2c''(\alpha)K}$ (which is independent of N), there are at most $N^{C(\alpha)}$ flowers.

Putting everything together, we infer that the probability that A contains some flower but not any non-trivial three-term arithmetic progression is bounded by

$$N^{C(\alpha)}\eta^{c'''(\alpha)s} \leq \eta^{c''''(\alpha)s}.$$

This completes our proof of Proposition 5.2.

Appendix A. Proof of Lemma 4.4

Without loss of generality, we just work with the real part of $\widehat{1}_R$. We shall prove $\mathbf{P}_R(\sup_{\xi \neq 0} |\Re\widehat{1}_R(\xi)| \geq \lambda/N) = o(1)$ for some appropriately chosen λ. Since the treatment for other cases is similar, we just show that $\mathbf{P}(\Re\widehat{1}_R(\xi) \geq \lambda/N)$ is very small for each fixed $\xi \neq 0$. For convenience, put

$$X = N\Re\widehat{1}_R(\xi) = \sum_{v \in V} 1_R(v)\Re e(-\langle v, \xi\rangle) := \sum_{v \in V} X_v.$$

One observes that X is a sum of N independent real variables X_v's. Choosing t to be a positive number smaller than 1, we have

$$\begin{aligned}\mathbf{P}_R(X \geq \lambda) &= \mathbf{P}_R(\exp(tX) \geq \exp(t\lambda)) \\ &\leq \mathbf{E}(\exp(tX))/\exp(t\lambda) \\ &= \prod \mathbf{E}(\exp(tX_v))/\exp(t\lambda) \\ &= \exp(t\mathbf{E}X)\prod \mathbf{E}(\exp(tX_v - t\mathbf{E}(X_v)))/\exp(t\lambda) \\ &= \exp(t\mathbf{E}X)\prod \mathbf{E}(\exp(tY_v))/\exp(t\lambda),\end{aligned}$$

where $Y_v := X_v - \mathbf{E}(X_v) = (1_R(v) - q)\Re e(-\langle v, \xi\rangle)$.

Notice that $|Y_v| \leq 1$ and $0 < t \leq 1$. We thus have $\exp(tY_v) \leq 1 + tY_v + t^2Y_v^2$. Hence

$$\mathbf{E}(\exp(tY_v)) \leq 1 + \mathbf{E}(t^2Y_v^2) \leq \exp(\mathbf{E}(t^2Y_v^2)).$$

Also, because $\mathbf{E}X = q\Re\sum_{v\in V} e(-\langle v, \xi\rangle) = 0$, it follows that

$$\mathbf{P}(X \geq \lambda) \leq \prod \mathbf{E}(\exp(tY_v))/\exp(t\lambda) \leq \exp(t^2\sum_{v\in V}\mathbf{E}(Y_v^2))/\exp(t\lambda).$$

On the other hand, it is clear from the definition of Y_v that $\sum_{v\in V}\mathbf{E}(Y_v^2) \leq qN$. Thus

$$\mathbf{P}(X > \lambda) \leq \exp(t^2qN - t\lambda).$$

By choosing $\lambda = |R|/\log N$ and $t = \lambda/(2qN) = 1/(2\log N)$ (thus $t < 1$), we deduce that

$$\mathbf{P}(X \geq |R|/\log N) \leq \exp(-|R|/(4\log^2 N)).$$

Hence

$$\mathbf{P}(\sup_{\xi\neq 0}\Re\widehat{1}_S(\xi) > |R|/(N\log N)) \leq N\exp(-|R|/(4\log^2 N)) = o(1).$$

(Note that the choice for λ above is not optimal, but it is enough for our goal.)

Acknowledgement. The author would like to thank Van Vu for useful discussions and encouragement. He is also grateful to Philip Wood and the referees for carefully reading this manuscript and providing very helpful remarks.

References

[1] **B. Green**, *A Szemerédi-type regularity lemma in Abelian groups, with applications*, GAFA 15 (2005), 340-376.
[2] **P. Frank, R. L. Graham** and **V. Rödl**, *On subsets of abelian groups with no three-term arithmetic progression*, Journal of Combinatorial Theory, Series A 45 (1987), 157-161.
[3] **T. C. Brown** and **J. P. Buhler**, *Lines implies spaces in density Ramsey theory*, Journal of Combinatorial Theory, Series A 36 (1984), 214-220.
[4] **Y. Kohayakawa, T. Łuczak** and **V. Rödl**, *Arithmetic progressions of length three in subsets on a random sets*, Acta Arith. 75 (1996), 133-163.
[5] **T. Tao** and **V. Vu**, *Additive Combinatorics*, Cambridge University Press, 2006.

Department of Mathematics, Rutgers University, Piscataway, NJ 08854, USA