Math. Proc. Camb. Phil. Soc. (2025), **178**, 177–192
doi:10.1017/S0305004125000076
First published online 4 April 2025

Improved bounds for skew corner-free sets

BY ADRIAN BEKER
Department of Mathematics, Faculty of Science, University of Zagreb, Bijenička cesta 30, 10000 *Zagreb, Croatia.*
e-mail: adrian.beker@math.hr

(*Received* 20 *March* 2024*; revised* 20 *March* 2024*; accepted* 08 *January* 2025)

Abstract

We construct skew corner-free subsets of $[n]^2$ of size $n^2 \exp(-O(\sqrt{\log n}))$, thereby improving on recent bounds of the form $\Omega(n^{5/4})$ obtained by Pohoata and Zakharov. We also prove that any such set has size at most $O(n^2(\log n)^{-c})$ for some absolute constant $c > 0$. This improves on the previously best known upper bound $O(n^2(\log\log n)^{-c})$, coming from Shkredov's work on the corners theorem.

2020 Mathematics Subject Classification: 11B30 (Primary); 05D99 (Secondary)

1. *Introduction*

A *skew corner* is a configuration of points of the form $(x, y), (x, y + d), (x + d, y')$ with $x, y, y', d \in \mathbb{Z}$. We say that a skew corner is *trivial* if $d = 0$ and *non-trivial* otherwise. A subset of $\mathbb{Z}^2$ is said to be *skew corner-free* if it contains no non-trivial skew corners. These notions were recently introduced by Pratt [18] (see also [19]), who asked the following interesting question: what is the largest size of a skew corner-free subset of $[n]^2$?

Letting $s(n)$ denote the answer to this question, one immediately sees that $s(n)$ cannot exceed $r(n)$, the largest size of a corner-free subset of $[n]^2$. We recall that a *corner* is a configuration of points of the form $(x, y), (x, y + d), (x + d, y)$ with $x, y, d \in \mathbb{Z}$, which is what one gets by requiring that $y' = y$ in the definition of a skew corner. The quantity $r(n)$ is an important object of study in additive combinatorics. Ajtai and Szemerédi [1] were the first to establish that $r(n) = o(n^2)$, a result now known as the *corners theorem*. A lot of work has since been done on this topic, see e.g. [23] for an elegant proof using the triangle removal lemma. The best known quantitative bounds on $r(n)$ are due to Shkredov [22] and Green [9], who showed that $r(n) = O(n^2(\log\log n)^{-1/73})$ and $r(n) \geq 2^{-(c+o(1))\sqrt{\log_2 n}} n^2$, where $c = 2\sqrt{2\log_2 \frac{4}{3}}$, respectively. For a version of Shkredov's argument in the finite field setting, which preserves the main ideas but greatly simplifies the technical details, we refer the reader to Green's excellent survey [8, Section 8] and the companion note [7].

Returning to skew corners, the additional degree of freedom in the definition makes it seem rather hard to construct large sets lacking these structures. This led Pratt to conjecture that, for any $\varepsilon > 0$, one has $s(n) = O(n^{1+\varepsilon})$ ([18, Conjecture 1·2]). The significance of this conjecture is that, were it to be true, it would rule out the the possibility of obtaining algorithms for matrix multiplication of complexity $O(n^{2+o(1)})$ using certain group-theoretic approaches. However, this prediction turned out to be false – in very recent work [17],

© The Author(s), 2025. Published by Cambridge University Press on behalf of Cambridge Philosophical Society

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

Pohoata and Zakharov showed that $s(n)$ grows at least as fast as $n^{5/4}$. Nonetheless, they hypothesised that one should still have $s(n) = O(n^{2-c})$ for some absolute constant $c > 0$. In this paper, we prove to the contrary that there exist skew corner-free subsets of $[n]^2$ of size $n^{2-o(1)}$. More precisely, we prove the following lower bound, which has the same shape as the corresponding bound for corner-free sets.

THEOREM 1·1. *There exists a constant $c_1 > 0$ such that*

$$s(n) \geq \frac{n^2}{2^{(c_1+o(1))\sqrt{\log_2 n}}}.$$

The proof of Theorem 1·1 is based on a Behrend-type construction (see [2]). In Section 2, we first give an example which suffices to disprove Pratt's conjecture, but is, in our opinion, somewhat simpler than that of [17]. We then proceed to give a full proof of Theorem 1·1 and briefly comment on how it relates to the arguments of [17].

In the other direction, it is natural to wonder whether one can improve the bound on $s(n)$ obtained from Shkredov's bounds on $r(n)$. Furthermore, Shkredov's argument is rather involved, so it would be desirable to have a simpler way of obtaining bounds for skew corner-free sets. We do so by proving the following result.

THEOREM 1·2. *There exists a constant $c_2 > 0$ such that*

$$s(n) = O\left(\frac{n^2}{(\log n)^{c_2}}\right).$$

The proof of Theorem 1·2 is based on the traditional Fourier-analytic approach, as opposed to [22], which uses a different, more subtle notion of uniformity. More precisely, we adapt Heath-Brown's [10] and Szemerédi's [24] proof of Roth's theorem. For a very readable exposition of these arguments, the reader may consult [6] (see also [16]). Even though the general scheme of the argument transfers fairly smoothly to our setting, we have to make some innovations on a technical level in order to make it work. This is done in Section 3, whereas in Section 4, we discuss how this problem relates to other Szemerédi-type problems and highlight some questions that remain open.

It will be useful to consider skew corner-free sets in the setting of abelian groups other than $\mathbb{Z}$, namely $\mathbb{Z}^d$ and finite cyclic groups. An observation that turns out to be convenient in proofs of both lower and upper bounds is that a subset of $[n]^2$ retains the property of being skew corner-free if we regard it as a subset of $(\mathbb{Z}/2n\mathbb{Z})^2$. Another feature of this property, of which we make crucial use in the proof of Theorem 1·2, is that it is translation-invariant in a rather strong sense. More specifically, if G is a finite cyclic group then the skew corner-freeness of a subset of G^2 is preserved under horizontal and vertical translations, where we are allowed to perform a different translation inside each column of the toroidal grid G^2. This fact will enable us to carry out the density increment strategy similarly as in the proof of Roth's theorem.

Notation and normalisations. Given a positive integer M, we abbreviate the set $\{1, \dots, M\}$ to $[M]$. We use standard asymptotic notation. Given quantities A and B that depend (in a possibly implicit way) on a positive integer parameter n, we write $A = O(B)$ or $A \ll B$ if there is a constant $C \in (0, \infty)$ such that $|A| \leq C|B|$ for all sufficiently large n. We also use

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

the equivalent notation $B = \Omega(A)$ and $B \gg A$. We write $A = o(B)$ to mean that $A/B \to 0$ as $n \to \infty$.

We now introduce some notation and conventions pertaining to Fourier analysis. Let G be a finite abelian group. Throughout the paper, G will be equipped with the uniform probability measure, whereas its dual $\widehat{G}$ will carry the counting measure. Thus, given functions $f\colon G \to \mathbb{C}$, $\phi\colon \widehat{G} \to \mathbb{C}$ and a parameter $p \in [1, \infty)$, the $\mathbb{L}^p$-norms of f and ϕ are

$$\|f\|_p = \left(\mathbb{E}_{x\in G}|f(x)|^p\right)^{\frac{1}{p}}, \quad \|\phi\|_p = \left(\sum_{\gamma\in\widehat{G}}|\phi(\gamma)|^p\right)^{\frac{1}{p}},$$

whereas the $\mathbb{L}^\infty$-norms of f and ϕ are

$$\|f\|_\infty = \sup_{x\in G}|f(x)|, \quad \|\phi\|_\infty = \sup_{\gamma\in\widehat{G}}|\phi(\gamma)|.$$

For functions $f, g\colon G \to \mathbb{C}$, their inner product is

$$\langle f, g\rangle = \mathbb{E}_{x\in G} f(x)\overline{g(x)}.$$

and the difference convolution of f and g is defined as

$$f \circ g\colon G \to \mathbb{C}, \quad x \mapsto \mathbb{E}_{y\in G} f(y)\overline{g(y-x)}.$$

The Fourier transform of f is defined by

$$\widehat{f}\colon \widehat{G} \to \mathbb{C}, \quad \gamma \mapsto \langle f, \gamma\rangle = \mathbb{E}_{x\in G} f(x)\overline{\gamma(x)}.$$

We will make use of the Fourier inversion formula

$$f(x) = \sum_{\gamma\in\widehat{G}} \widehat{f}(\gamma)\gamma(x),$$

Parseval's identity

$$\mathbb{E}_{x\in G} f(x)\overline{g(x)} = \sum_{\gamma\in\widehat{G}} \widehat{f}(\gamma)\overline{\widehat{g}(\gamma)},$$

and the fact that $\widehat{f \circ g} = \widehat{f}\widehat{g}$. When working with characters in explicit form, we will find it convenient to use the standard notation $e(\theta) := e^{2\pi i\theta}$ and $\|\theta\|_{\mathbb{R}/\mathbb{Z}} := \min\{|x| \mid x \in \theta\}$ for $\theta \in \mathbb{R}/\mathbb{Z}$, as well as the elementary estimates

$$4\|\theta\|_{\mathbb{R}/\mathbb{Z}} \leq |1 - e(\theta)| = 2|\sin(\pi\theta)| \leq 2\pi\|\theta\|_{\mathbb{R}/\mathbb{Z}}.$$

We next fix some notation regarding sets and associated functions. Given a subset $E \subseteq X$, where X is some finite ambient set, its indicator function is denoted by 1_E. If $E \subseteq G$ is non-empty, we denote by μ_E its normalised indicator function. Given a function $f\colon G^2 \to \mathbb{C}$, we define

$$S(f)\colon G \to \mathbb{C}, \quad x \mapsto \mathbb{E}_{y\in G} f(x, y).$$

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

Thus, $S(f)(x)$ coincides with $\|f(x,\cdot)\|_1$ if f is non-negative. We also define the following normalised variant of f:

$$\widetilde{f}\colon G^2 \to \mathbb{C}, \quad (x,y) \mapsto \begin{cases} \frac{f(x,y)}{\|f(x,\cdot)\|_1} & \text{if } \|f(x,\cdot)\|_1 \neq 0 \\ 0 & \text{otherwise} \end{cases}.$$

Finally, in Section 2, we will be using the notation $\langle\cdot,\cdot\rangle$ and $\|\cdot\|$ to denote the standard inner product and Euclidean norm on $\mathbb{R}^d$ respectively, but this should not cause any confusion (in particular, these are defined with respect to the counting measure as usual).

2. *Proof of Theorem 1·1*

We first present a simple proof of the weaker bound $s(n) \gg n^{1+c_1}$ for some constant $c_1 > 0$. We will be fairly brief since this is superseded by the results of [**17**]. The idea is to choose a fixed skew corner-free set $S \subseteq (\mathbb{Z}/b\mathbb{Z})^2$ of size $|S| > b$, for some positive integer b. That one can do so follows from the discussion at the end of the Section 1 since it is known that $s(n)$ is not $O(n)$ (see [**18**, Proposition 4·16]). Given a positive integer n, choose $k = \lfloor \log_b n \rfloor$ and let $A \subseteq [n]^2$ consist of all points $(x,y) \in [b^k]^2$ such that $(x_j, y_j) \in S$ for all $j \in [k]$, where z_j denotes the j-th digit in the base-b representation of $z-1$ for $z \in [b^k]$. It is straightforward to check that A is skew corner-free and that it has size $\Omega(n^{1+c_1})$ with $c_1 = \log_b |S| - 1 > 0$. We remark that this kind of product construction is analogous to the binary-ternary construction for sets free of three-term arithmetic progressions due to Erdős and Turán (see [**3**]).

We now turn to the main business of this section, which is to establish Theorem 1·1. Given a positive integer n, consider the box $B = [m]^d \subseteq \mathbb{Z}^d$, where m,d are positive integers such that $(2m)^d \leq n$. For positive integers r,t, consider the set

$$A_{r,t} = \{(x,y) \in B \times B \mid \|x\|^2 = r,\ \langle x,y\rangle = t\}.$$

We claim that $A_{r,t}$ is skew corner-free as a subset of $(\mathbb{Z}^d)^2$. Indeed, suppose that it contains three points of the form (x,y), (x,y''), (x',y') with $y''-y = x'-x$. Since $\langle x,y\rangle = t = \langle x,y''\rangle$, it follows that $\langle x, y''-y\rangle = 0$, whence $\langle x, x'-x\rangle = 0$. Hence, Pythagoras' theorem implies that

$$r = \|x'\|^2 = \|x\|^2 + \|x'-x\|^2 = r + \|x'-x\|^2.$$

It follows that $\|x'-x\| = 0$, that is, $x' = x$, as desired. In order to transfer this example to the integer setting, we use the map

$$\varphi\colon \mathbb{Z}^d \to \mathbb{Z}, \quad (x_1,\ldots,x_d) \mapsto 1 + \sum_{j=1}^{d} (2m)^{j-1}(x_j - 1),$$

which, when restricted to B, becomes a Freiman isomorphism into $[n]$. Thus, $(\varphi \times \varphi)(A_{r,t})$ is a skew corner-free subset of $[n]^2$ of the same size as $A_{r,t}$. Therefore, all that remains is to choose r,t so that $A_{r,t}$ is of the required size. But as r,t range over $[dm^2]$, the sets $A_{r,t}$ form a partition of $B \times B$. Hence, by the pigeonhole principle, there exists a choice of r,t such that

$$|A_{r,t}| \geq \frac{|B \times B|}{(dm^2)^2} = \frac{m^{2d}}{d^2 m^4} = \frac{m^{2d-4}}{d^2}.$$

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

Finally, choosing $m = \lfloor \frac{1}{2} n^{1/d} \rfloor$ and $d = \lfloor (2 \log_2 n)^{1/2} \rfloor$, it is not hard to check that we get

$$|A_{r,t}| \geq \frac{n^2}{2^{(4\sqrt{2}+o(1))\sqrt{\log_2 n}}}.$$

Remark. Our construction draws inspiration from the classical Behrend construction of sets lacking three-term arithmetic progressions [2] (in particular, the projection of our set onto the first coordinate is precisely such a Behrend set). Additionally, it turns out to be quite similar in spirit to the construction of Pohoata and Zakharov [17, Theorem 1·1]. Their construction is based on (the affine version of) the Hermitian unital over $\mathbb{F}_{p^2}^2$, which can be thought of as an analogue of a sphere in the finite field setting. By working directly in the integer setting, we take the advantage of being able to consider higher dimensional spheres, which results in superior bounds.

Remark. Both constructions described in this section can be easily modified so as to produce bi-skew corner-free sets of comparable size (for a precise definition, see [18, Definition 4·17]). Indeed, for the first construction, one merely has to replace the set S with a suitable bi-skew corner-free set, e.g.

$$\{(0,0),(0,1),(2,0),(2,3),(3,1),(3,3),(3,5),(4,0)\} \subseteq (\mathbb{Z}/6\mathbb{Z})^2.$$

For the second construction, one can modify the definition of the set $A_{r,t}$ by requiring that the second point also lie on the sphere of radius $\sqrt{r}$ centred at the origin:

$$A'_{r,t} = \{(x,y) \in B \times B \mid \|x\|^2 = \|y\|^2 = r,\ \langle x, y\rangle = t\}.$$

To ensure that $A'_{r,t}$ is large, one first uses the pigeonhole principle to find a value of r such that $\sqrt{r}S^{d-1}$ contains at least $m^d/dm^2 = m^{d-2}/d$ points of B. A further application of the pigeonhole principle then yields a set $A_{r,t}$ of size at least

$$\frac{(m^{d-2}/d)^2}{dm^2} = \frac{m^{2d-6}}{d^3},$$

which can then be optimised to give a bi-skew corner-free subset of $[n]^2$ of size at least

$$n^2 2^{-(4\sqrt{3}+o(1))\sqrt{\log_2 n}}.$$

3. *Proof of Theorem 1·2*

We begin by giving a high-level overview of our approach. We use the density increment/iteration method that falls into the broad conceptual framework outlined in [8, Section 4]. More precisely, our strategy is based on the approach of Heath-Brown [10] and Szemerédi [24] to Roth's theorem. The starting point is to write down a trilinear form that counts skew corners in a given set A. By performing standard manipulations, one can express this form in terms of the Fourier coefficients of A. An application of Hölder's inequality then shows that this quantity is controlled by various $\mathbb{L}^p$-norms associated to the Fourier transform of A. Thus, if the count of skew corners in A deviates from what we would expect for a random set of the same density, one of these norms must be large. From this, one can deduce either a large Fourier $\mathbb{L}^2$-mass along characters oscillating in the horizontal or vertical direction. These alternatives are then shown to lead to density increments on either an

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

arithmetic progression of rows or an arithmetic progression of columns. Since we are working in a two-dimensional setting, we need one additional averaging argument to obtain a density increment on a square substructure. The final density increment can then be iterated in the usual way to obtain the stated bound.

We now begin the argument proper. Let G be a finite abelian group. We define the trilinear form

$$\Lambda(f,g,h) := \mathbb{E}_{x,y,y',d\in G} f(x,y)g(x,y+d)h(x+d,y'),$$

where f,g,h are real-valued functions on G^2. In particular, if $A \subseteq G^2$, then $N^4\Lambda(1_A,1_A,1_A)$ is the number of skew corners in A. By employing the Fourier inversion formula and the fact that the Fourier transform diagonalises convolution, one can derive the Fourier representation of Λ:

$$\begin{aligned}\Lambda(f,g,h) &= \mathbb{E}_{x,d\in G}(g(x,\cdot)\circ f(x,\cdot))(d)S(h)(x+d)\\ &= \mathbb{E}_{x\in G}[S(h)\circ(g(x,\cdot)\circ f(x,\cdot))](x)\\ &= \mathbb{E}_{x\in G}\sum_{\gamma\in\widehat{G}} S(h)\circ(\widehat{g(x,\cdot)}\circ f(x,\cdot))(\gamma)\gamma(x)\\ &= \sum_{\gamma\in\widehat{G}} \widehat{S(h)}(\gamma)\mathbb{E}_{x\in G}\widehat{f(x,\cdot)}(\gamma)\overline{\widehat{g(x,\cdot)}(\gamma)}\gamma(x).\end{aligned}$$

In particular, the triangle inequality implies that

$$|\Lambda(f,g,h)| \le \sum_{\gamma\in\widehat{G}} |\widehat{S(h)}(\gamma)|\sigma(f,g)(\gamma),$$

where we define

$$\sigma(f,g)\colon \widehat{G}\to\mathbb{R},\quad \gamma\mapsto \mathbb{E}_{x\in G}|\widehat{f(x,\cdot)}(\gamma)||\widehat{g(x,\cdot)}(\gamma)|.$$

Thus, on applying applying Hölder's inequality, we obtain that

$$|\Lambda(f,g,h)| \le \|\widehat{S(h)}\|_p\|\sigma(f,g)\|_q \tag{1}$$

for all $1\le p,q\le\infty$ such that $1/p+1/q=1$. It is also useful to note that, by the Cauchy–Schwarz inequality and Parseval's identity, we have the following bound:

$$\|\sigma(f,g)\|_1 = \mathbb{E}_{x\in G}\sum_{\gamma\in\widehat{G}}|\widehat{f(x,\cdot)}(\gamma)||\widehat{g(x,\cdot)}(\gamma)| \le \mathbb{E}_{x\in G}\|f(x,\cdot)\|_2\|g(x,\cdot)\|_2. \tag{2}$$

To motivate the main part of the argument, we first briefly discuss some of the core ideas in the finite field setting. The key observation is that we can use Fourier analysis to bound the number of skew corners in a given set from below. This is made precise by the following result, which we call a 'generalised von Neumann theorem', borrowing the terminology from [**8**]. Given a parameter $\eta\in[0,1]$, we say that a subset of a finite abelian group is *η-uniform* if all non-trivial Fourier coefficients of its indicator function are at most η in absolute value.

PROPOSITION 3·1 (Generalised von Neumann theorem). *Let G be a finite abelian group and let $A\subseteq G^2$ be a subset of density α. Let $f_A = 1_A - \alpha 1_{G^2}$ be the balanced function of A.*

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

Then

$$\Lambda(1_A, 1_A, 1_A) \geq \alpha^3 - \alpha \|\widehat{S(f_A)}\|_\infty. \tag{3}$$

In particular, if A is η-uniform, then A contains at least $(\alpha^3 - \alpha\eta)N^4$ skew corners.

Proof. Consider the decomposition

$$\Lambda(1_A, 1_A, 1_A) = \Lambda(1_A, 1_A, f_A) + \alpha\Lambda(1_A, 1_A, 1_{G^2}).$$

By the Cauchy–Schwarz inequality, we have the following lower bound for the second term on the right-hand side:

$$\Lambda(1_A, 1_A, 1_{G^2}) = \mathbb{E}_{x,y,d\in G} 1_A(x, y)1_A(x, y + d) = \mathbb{E}_{x\in G} S(1_A)(x)^2 \geq \Big(\mathbb{E}_{x\in G} S(1_A)(x)\Big)^2 = \alpha^2.$$

On the other hand, by (1) with $p = \infty$, $q = 1$ and (2), the first term can be upper bounded in magnitude as follows:

$$|\Lambda(1_A, 1_A, f_A)| \leq \|\widehat{S(f_A)}\|_\infty \mathbb{E}_{x\in G} \|1_A(x, \cdot)\|_2^2 = \|\widehat{S(f_A)}\|_\infty \mathbb{E}_{x\in G} S(1_A)(x) = \alpha \|\widehat{S(f_A)}\|_\infty,$$

so (3) follows. The second statement follows on observing that $\widehat{S(f_A)}(\gamma) = \widehat{f_A}(\gamma, 1)$ for $\gamma \in \widehat{G}$, where we naturally identify $\widehat{G^2}$ with $\widehat{G}^2$.

Using Proposition 3·1 as a starting point, one can give an argument in the finite field setting in the same vein as for Meshulam's theorem [**14**]. Indeed, if $G = \mathbb{F}_2^n$ and $A \subseteq G^2$ is skew corner-free of density α, then one can easily verify that $\Lambda(1_A, 1_A, 1_A) \leq \alpha/N$, where $N = 2^n$. Thus, either $\alpha \leq \sqrt{2/N}$ or there exists $\gamma \in \widehat{G} \setminus \{1\}$ such that $|\widehat{S(f_A)}(\gamma)| \geq \alpha^2/2$. The latter case leads to a density increment of $S(1_A)$ on a coset of a codimension-1 subspace $H \leqslant G$, and an extra averaging argument produces a density increment of A on a coset of H^2. This density increment can then be iterated to obtain a bound of the form $|A| \ll 4^n/n$. This is already better than the best known bound [**13**] for corner-free sets in this setting, and the argument is substantially simpler.

However, we do not have the luxury of working in the finite field setting, so we have to work a bit harder. To begin, we set up some notation that will be used throughout the remainder of this section. Let $A \subseteq [n]^2$ be a subset of density α. Set $N = 2n$ and view $[n]^2$ as a subset of the group G^2, where $G = \mathbb{Z}/N\mathbb{Z}$. Let $f_A = 1_A - \alpha 1_{[n]^2}$ be the balanced function of A. Given $x \in G$, we let $A_x := \{y \in G \mid (x, y) \in A\}$ be the slice of A at x.

In the above setting, the following result plays a role analogous to that of Proposition 3·1. For comparison, the corresponding result in the setting of three-term arithmetic progressions is roughly [**6**, Proposition 2].

PROPOSITION 3·2. *Let $A \subseteq [n]^2$ be skew corner-free of density α. Then:*

(i) either $\alpha \leq \frac{8}{n}$;
(ii) or $\|\sigma(\widetilde{1_A}, 1_A)\|_{3/2} \geq 2^6\alpha$;
(iii) or $\|\widehat{S(f_A)}\|_3 \geq 2^{-12}\alpha$.

Proof. As in the proof of Proposition 3·1, we decompose

$$\Lambda(\widetilde{1_A}, 1_A, 1_A) = \Lambda(\widetilde{1_A}, 1_A, f_A) + \alpha\Lambda(\widetilde{1_A}, 1_A, 1_{[n]^2}).$$

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

Since A is skew corner-free when viewed as a subset of G^2, the left-hand side equals

$$\frac{1}{N}\mathbb{E}_{x,y,y'\in G}\widetilde{1_A}(x,y)1_A(x,y)1_A(x,y') = \frac{1}{N}\mathbb{E}_{x\in G}\Big(\mathbb{E}_{y\in G}\widetilde{1_A}(x,y)\Big)\Big(\mathbb{E}_{y'\in G}1_A(x,y')\Big) = \frac{\alpha}{4N}.$$

To obtain a lower bound on the second term on the right-hand side, we write it as follows:

$$\Lambda(\widetilde{1_A}, 1_A, 1_{[n]^2}) = \frac{1}{2N^2}\sum_{x\in[n],\ A_x\neq\varnothing}\frac{|\{(y_1,y_2)\in A_x^2 \mid y_2-y_1\in[n]-x\}|}{|A_x|}.$$

Since all pairs $(y_1,y_2)\in A_x^2$ such that either $1\le y_2\le y_1\le x$ or $x<y_1\le y_2\le n$ are counted in the numerator of the summand, we get that this quantity is at least

$$\begin{aligned}\frac{1}{2N^2}\sum_{x\in[n],\ A_x\neq\varnothing}\frac{1}{|A_x|}\left(\frac{|A_x\cap[1,x]|^2}{2}+\frac{|A_x\cap[x+1,n]|^2}{2}\right) &\ge \frac{1}{2N^2}\sum_{x\in[n],\ A_x\neq\varnothing}\frac{1}{|A_x|}\cdot\frac{|A_x|^2}{4}\\ &= \frac{1}{8N^2}\sum_{x\in[n],\ A_x\neq\varnothing}|A_x|\\ &= \frac{\alpha}{32},\end{aligned}$$

where we made use of convexity in the first inequality. Hence, if (i) fails to hold, it follows by (1) with $p=3$, $q=3/2$ that

$$\|\widehat{S(f_A)}\|_3\|\sigma(\widetilde{1_A},1_A)\|_{3/2}\ge\frac{\alpha^2}{64},$$

so either (ii) or (iii) must hold.

In the following lemma we record two bounds based on Parseval's identity that will be important for the success of the $\mathbb{L}^2$ density increment strategy.

LEMMA 3·3. *For a subset $A\subseteq[n]^2$ of density α, we have the following bounds*:

(i) $\|\sigma(\widetilde{1_A},1_A)\|_1\le 1$;
(ii) $\sum_{\gamma\in\widehat{G}}|\widehat{S(f_A)}(\gamma)|^2\le\alpha$.

Proof. To prove (i), observe that if $A_x\neq\varnothing$, then

$$\|\widetilde{1_A}(x,\cdot)\|_2\|1_A(x,\cdot)\|_2 = \frac{\|1_A(x,\cdot)\|_2^2}{S(1_A)(x)} = 1,$$

whereas if $A_x=\varnothing$, then the left-hand side trivially vanishes. The conclusion now follows from the estimate (2). Likewise, Parseval's identity implies that

$$\sum_{\gamma\in\widehat{G}}|\widehat{S(f_A)}(\gamma)|^2 = \|S(f_A)\|_2^2\le\|S(1_A)\|_2^2\le\|S(1_A)\|_\infty\|S(1_A)\|_1\le\alpha,$$

which proves (ii).

We note in passing that, assuming the first alternative in Proposition 3·2 doesn't hold, we could have instead used the estimate (1) with $p=\infty$, $q=1$ coupled with part (i) of Lemma 3·3 to obtain the existence of a character $\gamma\in\widehat{G}\setminus\{1\}$ such that $|\widehat{S(f_A)}(\gamma)|\ge\alpha^2/64$. We could then carry on essentially in the same fashion as in Roth's original argument [**20**] to obtain a

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

bound of the form $s(n) \ll n^2/\log\log n$. However, as mentioned before, we pursue a different route, which leads to a better bound.

The following lemma is technical in nature and is a variant of [6, Lemma 4]. It is a device that will be used to convert information about large Fourier $\mathbb{L}^p$-mass into that about large Fourier $\mathbb{L}^2$-mass, which in turn serves as input to the $\mathbb{L}^2$ density increment strategy.

LEMMA 3·4. *Let $\phi\colon \widehat{G} \to \mathbb{C}$ be such that $\|\phi\|_{3/2} \geq \beta$ for some $\beta > 0$. Then there exists a non-empty set $\Gamma \subseteq \widehat{G}$ such that $|\Gamma| \leq (2\beta^{-1}\|\phi\|_1)^3$ and $\sum_{\gamma\in\Gamma} |\phi(\gamma)| \geq 2^{-3}|\Gamma|^{1/4}\beta$.*

Proof. Enumerate $\widehat{G}$ as $\{\gamma_1, \ldots, \gamma_N\}$ so that the sequence $(b_j)_{j\in[N]}$ defined by $b_j = |\phi(\gamma_j)|$ for $j \in [N]$ is non-increasing. Then note that for each $k \in [N]$ we have

$$kb_k \leq \sum_{j=1}^{k} b_j \leq \|\phi\|_1,$$

whence $b_k \leq k^{-1}\|\phi\|_1$. Hence, for all $k \in [N]$ we get

$$\sum_{j=k+1}^{N} b_j^{3/2} \leq b_k^{1/2} \sum_{j=k+1}^{N} b_j \leq b_k^{1/2}\|\phi\|_1 \leq k^{-1/2}\|\phi\|_1^{3/2}.$$

In particular, taking $k = \min\left\{\left\lceil 4(\beta^{-1}\|\phi\|_1)^3\right\rceil, N\right\}$, it follows that $\sum_{j=1}^{k} b_j^{3/2} \geq \frac{1}{2}\beta^{3/2}$. Now if we had $b_j \leq 2^{-3}j^{-3/4}\beta$ for all $j \in [k]$, then it would follow that

$$\sum_{j=1}^{k} b_j^{3/2} \leq 2^{-9/2}\beta^{3/2} \sum_{j=1}^{k} j^{-9/8} < 2^{-9/2}\beta^{3/2}\left(1 + \int_1^{\infty} x^{-9/8}\, dx\right) < \frac{1}{2}\beta^{3/2},$$

which is absurd. Therefore, there must exist $m \in [k]$ such that $b_m \geq 2^{-3}m^{-3/4}\beta$, whence

$$\sum_{j=1}^{m} b_j \geq mb_m \geq 2^{-3}m^{1/4}\beta.$$

Thus, we may take $\Gamma = \{\gamma_1, \ldots, \gamma_m\}$, as desired.

Following [11], for a set of characters $\Gamma \subseteq \widehat{G}$ and a parameter $\nu \in (0, 2]$, we will say that Γ is *ν-annihilated* by a non-empty set $X \subseteq G$ if $|\gamma(x) - 1| \leq \nu$ for all $\gamma \in \Gamma$ and $x \in X$. The following simple lemma is a variant of [11, Lemma 12].

LEMMA 3·5. *If $\Gamma \subseteq \widehat{G}$ is ν-annihilated by X, then $|\widehat{\mu_X}(\gamma)| \geq 1 - \nu^2/2$ for all $\gamma \in \Gamma$.*

Proof. If $\gamma(x) = e(x\theta)$ for $\theta \in \mathbb{R}/\mathbb{Z}$, then for all $x \in X$ we have $2|\sin(\pi x\theta)| = |1 - e(x\theta)| \leq \nu$, whence

$$\mathrm{Re}\gamma(x) = \cos(2\pi x\theta) = 1 - 2\sin^2(\pi x\theta) \geq 1 - \frac{\nu^2}{2}.$$

Therefore, we have

$$\mathrm{Re}\widehat{\mu_X}(\gamma) = \mathbb{E}_{x\in X}\mathrm{Re}\gamma(x) \geq 1 - \frac{\nu^2}{2},$$

so the conclusion follows.

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

The following two propositions, whose names are partly inspired by [11, Section 8], constitute the heart of the proof. In the first proposition, we obtain a density increment on a suitable set of rows as a consequence of the second alternative in Proposition 3·2.

PROPOSITION 3·6 (Horizontal $\mathbb{L}^2$ density increment). *Let $A \subseteq [n]^2$ be a subset of density $\alpha > 0$ such that $\|\sigma(\widetilde{1_A}, 1_A)\|_{3/2} \geq 2^6\alpha$. Then there exists a non-empty set $\Gamma \subseteq \widehat{G}$ of size $|\Gamma| \ll \alpha^{-3}$ such that, for any set $P \subseteq G$ that 1-annihilates Γ,*

$$\mathbb{E}_{x\in G} \max_{y\in G} \frac{|(A_x - y) \cap P|}{|P|} \geq 2|\Gamma|^{1/4}\alpha.$$

Proof. By Lemma 3·4 applied with $\phi = \sigma(\widetilde{1_A}, 1_A)$, $\beta = 2^6\alpha$ and part (i) of Lemma 3·3, we obtain a set $\Gamma \subseteq \widehat{G}$ of the required size such that

$$\sum_{\gamma\in\Gamma} \mathbb{E}_{x\in G} |\widehat{\widetilde{1_A}(x,\cdot)}(\gamma)||\widehat{1_A(x,\cdot)}(\gamma)| \geq 2^3|\Gamma|^{1/4}\alpha.$$

Hence, if Γ is 1-annihilated by P, we obtain using Parseval's identity and Lemma 3·5 that

$$\begin{aligned}\mathbb{E}_{x\in G}\langle \widetilde{1_A}(x,\cdot)\circ\mu_P, 1_A(x,\cdot)\circ\mu_P\rangle &= \mathbb{E}_{x\in G}\sum_{\gamma\in\widehat{G}} |\widehat{\widetilde{1_A}(x,\cdot)}(\gamma)||\widehat{1_A(x,\cdot)}(\gamma)||\widehat{\mu_P}(\gamma)|^2\\ &\geq 2^{-2}\sum_{\gamma\in\Gamma} \mathbb{E}_{x\in G}|\widehat{\widetilde{1_A}(x,\cdot)}(\gamma)||\widehat{1_A(x,\cdot)}(\gamma)|\\ &\geq 2|\Gamma|^{1/4}\alpha.\end{aligned}$$

On the other hand, this quantity can be upper bounded as follows:

$$\begin{aligned}\mathbb{E}_{x\in G}\langle \widetilde{1_A}(x,\cdot)\circ\mu_P, 1_A(x,\cdot)\circ\mu_P\rangle &\leq \mathbb{E}_{x\in G}\|\widetilde{1_A}(x,\cdot)\circ\mu_P\|_1\|1_A(x,\cdot)\circ\mu_P\|_\infty\\ &= \mathbb{E}_{x\in G}\|\widetilde{1_A}(x,\cdot)\|_1\|\mu_P\|_1\|1_A(x,\cdot)\circ\mu_P\|_\infty\\ &\leq \mathbb{E}_{x\in G}\|1_A(x,\cdot)\circ\mu_P\|_\infty,\end{aligned}$$

whence the desired conclusion follows.

The reason why we use the 'locally normalised' indicator function $\widetilde{1_A}$ instead of 1_A lies in the last part of the proof of Proposition 3·6. At that point, we have to upper bound the average additive energy of a vertical slice of A with P. Had we used 1_A instead of $\widetilde{1_A}$, we would end up with a weighted average having the column densities of A as weights. Since these densities may vary a lot from column to column, this would not necessarily imply an adequate density increment. This is in contrast to the setting of Roth's theorem, where all normalisations are of a 'global' nature and hence are purely a matter of taste.

The following proposition has the third alternative in Proposition 3·2 as a starting assumption and produces a density increment on a suitable set of columns.

PROPOSITION 3·7 (Vertical $\mathbb{L}^2$ density increment). *Let $A \subseteq [n]^2$ be a subset of density $\alpha > 0$ such that $\|\widehat{S(f_A)}\|_3 \geq 2^{-12}\alpha$. Then there exists a non-empty set $\Gamma \subseteq \widehat{G}$ of size $|\Gamma| \ll \alpha^{-3}$ such that, for any set $P \subseteq [n]$ of diameter at most $2^{-30}|\Gamma|^{1/4}\alpha n$ that 1-annihilates Γ,*

$$\max_{t\in G} \mathbb{E}_{x\in P}|A_{x+t}|/n \geq (1 + 2^{-27}|\Gamma|^{1/4})\alpha.$$

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

Proof. We start by applying Lemma 3·4 with $\phi = |\widehat{S(f_A)}|^2$, $\beta = 2^{-24}\alpha^2$ and part (ii) of Lemma 3·3 to obtain a set $\Gamma \subseteq \widehat{G}$ of the required size such that

$$\sum_{\gamma\in\Gamma} |\widehat{S(f_A)}(\gamma)|^2 \geq 2^{-27}|\Gamma|^{1/4}\alpha^2.$$

Thus, for any $P \subseteq [n]$ that 1-annihilates Γ, we get using Parseval's identity and Lemma 3·5 that

$$\|S(f_A)\circ\mu_P\|_2^2 = \sum_{\gamma\in\widehat{G}} |\widehat{S(f_A)}(\gamma)|^2|\widehat{\mu_P}(\gamma)|^2 \geq 2^{-2}\sum_{\gamma\in\Gamma} |\widehat{S(f_A)}|^2 \geq 2^{-29}|\Gamma|^{1/4}\alpha^2.$$

Next, we perform the usual passage from balanced to indicator functions. By expanding out inner products and using the adjoint property of convolution, one obtains that

$$\|S(f_A)\circ\mu_P\|_2^2 = \|S(1_A)\circ\mu_P\|_2^2 - \alpha\langle S(1_A)+S(f_A), S(1_{[n]^2})\circ(\mu_P\circ\mu_P)\rangle.$$

But note that we have the estimates

$$\|S(1_A)+S(f_A)\|_\infty \leq 1, \quad \|S(1_{[n]^2})\circ(\mu_P\circ\mu_P) - S(1_{[n]^2})\|_1 \leq \frac{\mathrm{diam}(P)}{n},$$

where the second one holds because $S(1_{[n]^2})\circ(\mu_P\circ\mu_P)$ is $1/2$ on $(\mathrm{diam}(P), n-\mathrm{diam}(P)]$, 0 outside $(-\mathrm{diam}(P), n+\mathrm{diam}(P)]$ and has values in $[0,1/2]$ everywhere. Hence, Hölder's inequality implies that

$$|\langle S(1_A)+S(f_A), S(1_{[n]^2})\circ(\mu_P\circ\mu_P) - S(1_{[n]^2})\rangle| \leq \frac{\mathrm{diam}(P)}{n}.$$

Since we have

$$\langle S(1_A)+S(f_A), S(1_{[n]^2})\rangle = \frac{\alpha}{8},$$

it follows that

$$\|S(1_A)\circ\mu_P\|_2^2 \geq \|S(f_A)\circ\mu_P\|_2^2 + \Big(\frac{\alpha}{8} - \frac{\mathrm{diam}(P)}{n}\Big)\alpha.$$

Hence, if $\mathrm{diam}(P) \leq 2^{-30}|\Gamma|^{1/4}\alpha n$, we obtain that

$$\|S(1_A)\circ\mu_P\|_2^2 \geq \Big(\frac{1}{8} + 2^{-30}|\Gamma|^{1/4}\Big)\alpha^2.$$

We now conclude similarly as in the proof of Proposition 3·6, by observing the upper bound

$$\|S(1_A)\circ\mu_P\|_2^2 \leq \|S(1_A)\circ\mu_P\|_1\|S(1_A)\circ\mu_P\|_\infty = \frac{1}{4}\alpha\|S(1_A)\circ\mu_P\|_\infty$$

and unfolding the definition of the difference convolution of $S(1_A)$ and μ_P.

To complete the density increment argument, we require the following well-known result (for a proof, see e.g. [**21**]):

LEMMA 3·8 (Dirichlet's approximation theorem). *Let $\theta_1,\dots,\theta_m \in \mathbb{R}/\mathbb{Z}$ and let Q be a positive integer. Then there exists $q \in [Q]$ such that $\|q\theta_j\|_{\mathbb{R}/\mathbb{Z}} \leq 1/Q^{1/m}$ for all $j \in [m]$.*

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

Lemma 3·8 has the following straightforward consequence, which will be used to furnish arithmetic progressions P as input to Propositions 3·6 and 3·7. In order for the final averaging argument to succeed, we also need to keep track of the *span* of these progressions, which is defined to be the product of the length and the common difference (this piece of terminology is not standard).

COROLLARY 3·9. *Let $\alpha \in (0,1]$ and let $\Gamma \subseteq \widehat{G}$ be a set of characters of size $m \ll \alpha^{-3}$. Then*:

(i) either $\alpha \ll (\log n)^{-1/4}$;

(ii) or Γ is 1-annihilated by an arithmetic progression $P \subseteq [n]$ of length $|P| \gg \alpha n^{\frac{1}{m+1}}$ and span at most $2^{-30}\alpha n$.

Proof. Let $\ell = \left\lfloor \frac{1}{6}(2^{-30}\alpha n)^{\frac{1}{m+1}} \right\rfloor$ and $Q = (6\ell)^m$. Here, we are assuming that $\alpha n \geq 2^{30}6^{m+1}$, which is certainly the case unless (i) holds. Thus, $\ell \geq 1$ and $\ell \gg \alpha n^{\frac{1}{m+1}}$. Enumerating Γ as $\{\gamma_1, \dots, \gamma_m\}$, where $\gamma_j(x) = e(x\theta_j)$ for $j \in [m]$ and $\theta_1, \dots, \theta_m \in \mathbb{R}/\mathbb{Z}$, Lemma 3·4 implies the existence of $q \in [Q]$ such that $\|q\theta_j\|_{\mathbb{R}/\mathbb{Z}} \leq 1/6\ell$ for all $j \in [m]$. Consequently, if we let P be the arithmetic progression $q \cdot [\ell]$, it follows that $\|x\theta_j\|_{\mathbb{R}/\mathbb{Z}} \leq 1/6$ for all $j \in [m]$ and $x \in P$. Thus, Γ is 1-annihilated by P and P has span $\ell q \leq 2^{-30}\alpha n$, as desired.

Finally, we are ready to assemble all the pieces we have developed thus far into a single density increment result:

PROPOSITION 3·10. *Let $A \subseteq [n]^2$ be skew corner-free of density α. Then at least one of the following holds*:

(i) $\alpha \ll (\log n)^{-1/4}$;

(ii) there is a positive integer $m \ll \alpha^{-3}$, an arithmetic progression $P \subseteq [n]$ of length $|P| \gg \alpha n^{\frac{1}{m+1}}$ and span at most $\alpha m^{1/4} n$ and a skew corner-free set $A' \subseteq [n] \times P$ of density at least $4m^{1/4}\alpha$;

(iii) there is a positive integer $m \ll \alpha^{-3}$, an arithmetic progression $P \subseteq [n]$ of length $|P| \gg \alpha n^{\frac{1}{m+1}}$ and span at most $2^{-30}\alpha m^{1/4} n$ and a skew corner-free set $A'' \subseteq P \times [n]$ of density at least $(1 + 2^{-27}m^{1/4})\alpha$.

Proof. Assume that (i) doesn't hold. In particular, we must be in one of the cases (ii) or (iii) of Proposition 3·2. If we are in case (ii), then we apply Proposition 3·6 together with Corollary 3·9. We obtain an arithmetic progression $P \subseteq [n]$ of length and span as in case (ii) of the present result such that

$$\mathbb{E}_{x \in G} \max_{y \in G} \frac{|(A_x - y) \cap P|}{|P|} \geq 2m^{1/4}\alpha,$$

where $m \ll \alpha^{-3}$ is a positive integer. Hence, we may take

$$A' = \bigcup_{x \in [n]} [\{x\} \times ((A_x - y_x) \cap P)],$$

where for each $x \in [n]$ we choose $y_x \in G$ so as to maximise the quotient $|(A_x - y_x) \cap P|/|P|$.

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

Now suppose that we are in case (iii) of Proposition 3·2. Then we apply Proposition 3·7 combined with Corollary 3·9. This yields an arithmetic progression $P \subseteq [n]$ of length and span as in case (iii) of the present result such that

$$\max_{t \in G} \mathbb{E}_{x \in P} |A_{x+t}|/n \geq (1 + 2^{-27} m^{1/4})\alpha,$$

where $m \ll \alpha^{-3}$ is a positive integer. But we may now take

$$A'' = \bigcup_{x \in P} (\{x\} \times A_{x+y}),$$

where $t \in G$ is chosen so as to maximise $\mathbb{E}_{x \in P} |A_{x+t}|/n$. This concludes the proof.

In order to obtain a density increment that can be iterated, we have to make one last manoeuvre:

LEMMA 3·11. *Let $P \subseteq [n]$ be an arithmetic progression of length ℓ and common difference d. If $B \subseteq P \times [n]$ is a subset of density β, then there is a translate $P' \subseteq [n]$ of P such that B has density at least $\beta - \ell d/n$ on $P \times P'$. Similarly, if $B \subseteq [n] \times P$ has density β, then there is a translate $P' \subseteq [n]$ of P such that B has density at least $\beta - \ell d/n$ on $P' \times P$.*

Proof. We only prove the first statement; the proof of the second statement is completely analogous. Partition $[\ell d]$ into d translates of P and then partition $[n]$ into translates of $[\ell d]$ and an error set E of size less than ℓd. Then $P \times E$ occupies at most $\ell^2 d$ elements of B, so $B \setminus (P \times E)$ has density at least $\beta - \ell d/n$ on $P \times [n]$. The conclusion now follows by averaging.

By putting together Proposition 3·10 and Lemma 3·11 and using translation- and dilation-invariance of skew corner-free sets, we obtain the final form of our density increment result:

PROPOSITION 3·12. *Let $A \subseteq [n]^2$ be skew corner-free of density α. Then*:

(i) either $\alpha \ll (\log n)^{-1/4}$;

(ii) or there exist positive integers $m \ll \alpha^{-3}$, $n' < n$ such that $n' \gg \alpha n^{\frac{1}{m+1}}$ and there exists a skew corner-free set $A' \subseteq [n']^2$ of density at least $(1 + c' m^{1/4})\alpha$, where $c' > 0$ is an absolute constant.

To establish Theorem 1·2, all that remains is to iterate Proposition 3·6. In analogy with the deduction of [**6**, Theorem 1] from [**6**, Proposition 7], we show by induction that for $n > 1$, any skew corner-free subset of $[n]^2$ has density at most $C'(\log n)^{-c_2}$, where $C' > 0$ is an absolute constant. To this end, we apply Proposition 3·12; if (i) holds, we are immediately done by taking $c_2 \leq \frac{1}{4}$. Otherwise, there exist positive integers m, n' as in (ii) and a skew corner-free subset of $[n']^2$ of density at least $(1 + c' m^{1/4})\alpha$. Here, we may assume that n is sufficiently large and

$$m \leq \frac{1}{4}(\log n)^{3/4}, \quad n' \geq n^{\frac{1}{m+1}} (\log n)^{-1/4} > 1.$$

Hence, by the induction hypothesis, we have

$$(1 + c' m^{1/4})\alpha \leq C'(\log n')^{-c_2}.$$

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

Assuming for the sake of contradiction that $\alpha > C'(\log n)^{-c_2}$, it follows that

$$(1 + c'm^{1/4})\Big(\frac{\log n'}{\log n}\Big)^{c_2} < 1. \tag{4}$$

Now observe that for all $x > 0$ we have

$$\exp(3x^{1/3}) > \frac{(3x^{1/3})^3}{3!} > 4x,$$

so in particular $\log(4m) \le 3m^{1/3}$. Since the function $x \mapsto \log x/x$ is decreasing on the interval $[e, \infty)$ and $\log n \ge (4m)^{4/3} > e$, it follows that

$$\frac{\log\log n}{4\log n} \le \frac{\log(4m)}{3(4m)^{4/3}} \le \frac{1}{4m}.$$

Therefore, we have

$$\frac{\log n'}{\log n} \ge \frac{1}{m+1} - \frac{\log\log n}{4\log n} \ge \frac{1}{4m},$$

so (4) implies that

$$1 + c'm^{1/4} < (4m)^{c_2}.$$

If we take $c_2 \le 1/5$ say, the above fails for all $m > m_0$, where $m_0 \ge 1$ is some constant depending on c'. The remaining range $m \le m_0$ can now be dealt with by taking c_2 small enough so that $(4m_0)^{c_2} < 1 + c'$. We thus arrive at a contradiction, which concludes the proof.

4. *Concluding remarks and open problems*

The proofs of both Theorem 1·1 and Theorem 1·2 exhibit many parallels between the problem of estimating $s(n)$ and that of determining the largest size of a subset of $[n]$ lacking three-term arithmetic progressions. The latter problem is one of the cornerstones of additive combinatorics and has been a subject of intense study in the last seventy years. Recently, a huge breakthrough on this problem was made by Kelley and Meka [**12**], who proved upper bounds matching the shape of the well-known lower bound of Behrend [**2**]. This came as somewhat of a surprise since Kelley and Meka largely abandon the Fourier-analytic approach, relying instead on methods based in physical space. Thus, it is an interesting challenge to investigate the scope of these methods and explore potential applications to other related problems. In the very recent survey [**15**], some optimism has been expressed as to the possibility of applying Kelley–Meka techniques to improve bounds for corner-free sets. Given the apparently even larger similarity of skew corners to three-term arithmetic progressions, we think that a good intermediate step would be to investigate whether these techniques could be used to the improve bounds for the problem considered in this paper. Specifically, we ask the following question:

Question 4·1. Is it true that

$$s(n) \ll \frac{n^2}{\exp(c(\log n)^{\beta})}$$

for some constants $c, \beta > 0$?

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

There is another rather intriguing parallel with the study of corner-free sets which we think deserves more attention. As mentioned before, the usual notion of Fourier uniformity is not sufficient to get a handle on the count of corners in a given set. A concrete example explaining this phenomenon is given by [8, example 5·3], which also appears in [22, Section 1]. In this example, a random set of $B \subseteq \mathbb{F}_2^n$ of density β is taken and the product set $A = B^2 \subseteq (\mathbb{F}_2^n)^2$ is considered. Being the product of a random set with itself, A is a highly uniform set of density $\alpha = \beta^2$. However, the count of corners in A is equal to the number of triples $(x, y, d) \in (\mathbb{F}_2^n)^3$ such that $x, y, x + d, y + d \in B$, which is $\beta^4 N^3 = \alpha^2 N^3$ in expectation. This is significantly more than $\alpha^3 N^3$ corners, what one would expect for a truly random subset of $(\mathbb{F}_2^n)^2$ of density α.

Taking a closer look at this example, it becomes clear that a similar reasoning applies to our problem: if $B \subseteq \mathbb{F}_2^n$ is random of density β, then $A = B^2$ contains roughly $\beta^5 N^4 = \alpha^{5/2} N^4$ skew corners, whereas a truly random subset of density α would contain more like $\alpha^3 N^4$ skew corners. However, this apparent deficiency did not prevent us from using Fourier analysis to obtain bounds for skew corner-free sets. Indeed, we have shown in Proposition 3·1 that any uniform subset has roughly *at least* as many skew corners as a random set of the same density. In light of this, we do not find the example discussed in the previous paragraph as entirely satisfactory evidence as to the impossibility of applying Fourier analysis to the study corner-free sets. This motivates us to ask the following question:

Question 4·2. Do there exist Fourier uniform sets with significantly fewer than expected corners? More precisely, do there exist $\alpha, \delta > 0$ such that, for infinitely many N, there exists a set $A \subseteq (\mathbb{Z}/N\mathbb{Z})^2$ of density $\alpha + o(1)$ which is $o(1)$-uniform, but has at most $(\alpha^3 - \delta)N^3$ corners?

We believe that Question 4·2 should have an affirmative answer. However, showing the existence of such sets might not be a simple task. Indeed, comparing this to the situation for four-term arithmetic progressions, it has been known since the seminal work of Gowers [4] that there exist uniform sets with *more* than expected arithmetic progressions of length four. In spite of this, constructing such sets with *fewer* than expected 4-progressions turned out to be a non-trivial problem on its own (see [5]).

Acknowledgements. This work was supported by the Croatian Science Foundation under the project number HRZZ-IP-2022-10-5116 (FANAP). The author is grateful to Vjekoslav Kovač and Rudi Mrazović for bringing [18] and the questions therein to his attention, as well as for fruitful discussions and encouragement. He would also like to thank Cosmin Pohoata and Dmitrii Zakharov for helpful comments on a draft of the paper, Zach Hunter for pointing out the reference [3] and the anonymous referee for several useful suggestions.

REFERENCES

[1] M. Ajtai and E. Szemerédi. Sets of lattice points that form no squares. *Studia Sci. Math. Hungar.* **9** (1974), 9–11.
[2] F. A. Behrend. On sets of integers which contain no three terms in arithmetical progression. *Proc. Nat. Acad. Sci. U.S.A.* 32 (1946), 331–332.
[3] P. Erdös and P. Turán. On Some Sequences of Integers. *J. London Math. Soc.* **11** (1936), 261–264.
[4] W. T. Gowers. A new proof of Szemerédi's theorem. *Geom. Funct. Anal.* **11** (2001), 465–588.
[5] W. T. Gowers. A uniform set with fewer than expected arithmetic progressions of length 4. *Acta Math. Hungar.* **161** (2020), 756–767.

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press

[6] B. GREEN. Progressions of length 3 following Szemerédi, (1999). Manuscript, available at: https://people.maths.ox.ac.uk/greenbj/papers/szemeredi-roth.pdf
[7] B. GREEN. An argument of Shkredov in the finite field setting, (2004). Manuscript, available at: https://people.maths.ox.ac.uk/greenbj/papers/corners.pdf.
[8] B. GREEN. Finite field models in additive combinatorics. In *Surveys in combinatorics*. London Math. Soc. Lecture Note Ser. vol. 327 (Cambridge University Press, Cambridge, 2005), pp. 1-27.
[9] B. GREEN. Lower bounds for corner-free sets. *New Zealand J. Math.* **51** (2021), 1–2.
[10] D. R. HEATH-BROWN. Integer sets containing no arithmetic progressions. *J. London Math. Soc.* (2) **35** (1987), 385–394.
[11] K. HENRIOT. Bourgain's bounds for Roth's theorem, (2013). Manuscript, available at: https://khenriot.wordpress.com/wp-content/uploads/2016/03/bourgainroth.pdf.
[12] Z. KELLEY and R. MEKA. Strong bounds for 3-progressions. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science*, pp. 933–973. IEEE Computer Soc., Los Alamitos, CA.
[13] M. T. LACEY and W. MCCLAIN. On an argument of Shkredov on two-dimensional corners. *Online J. Anal. Comb.* (2007), art. 2, 21.
[14] R. MESHULAM. On subsets of finite abelian groups with no 3-term arithmetic progressions. *J. Combin. Theory Ser. A* **71** (1995), 168–172.
[15] S. PELUSE. Finite field models in arithmetic combinatorics – twenty years on. ArXiv:2312.08100.
[16] S. PELUSE. Recent progress on bounds for sets with no three terms in arithmetic progression. *Astérisque* (2022), no. 1196, 581.
[17] C. POHOATA and D. ZAKHAROV. On skew corner-free sets. ArXiv:2401.17507.
[18] K. PRATT. On generalised corners and matrix multiplication. ArXiv:2309.03878.
[19] K. PRATT. A variant of the corners problem. MathOverflow (2023), https://mathoverflow.net/questions/451580/a-variant-of-the-corners-problem.
[20] K. F. ROTH. On certain sets of integers. *J. London Math. Soc.* **28** (1953), 104–109.
[21] W. M. SCHMIDT. *Diophantine Approximation*, volume 785 of Lecture Notes in Math. vol. 785 (Springer, Berlin, 1980).
[22] I. D. SHKREDOV. On a generalisation of Szemerédi's theorem. *Proc. London Math. Soc.* (3) 93 (2006), 723–760.
[23] J. SOLYMOSI. Note on a generalisation of Roth's theorem. In *Discrete Comput. Geom.* Algorithms Combin. vol 25 (Springer, Berlin, 2003), pp. 825–827.
[24] E. SZEMERÉDI. Integer sets containing no arithmetic progressions. *Acta Math. Hungar.* **56** (1990), 155–158.

https://doi.org/10.1017/S0305004125000076 Published online by Cambridge University Press