

Sampling-based proofs of almost-periodicity results and algorithmic applications

Eli Ben-Sasson* Noga Ron-Zewi† Madhur Tulsiani‡ Julia Wolf§

October 25, 2012

Abstract

We give new combinatorial proofs of known almost-periodicity results for sumsets of sets with small doubling in the spirit of Croot and Sisask [CS10], whose almost-periodicity lemma has had far-reaching implications in additive combinatorics. We provide an alternative (and L^p -norm free) point of view, which allows for proofs to easily be converted to probabilistic algorithms that decide membership in almost-periodic sumsets of dense subsets of \mathbb{F}_2^n .

As an application, we give a new algorithmic version of the *quasipolynomial Bogolyubov-Ruzsa lemma* recently proved by Sanders [San10]. Together with the results by the last two authors [TW11], this implies an algorithmic version of the *quadratic Goldreich-Levin theorem* in which the number of terms in the quadratic Fourier decomposition of a given function is quasipolynomial in the error parameter ε , compared with an exponential dependence previously proved by the authors. It also improves the running time of the algorithm to have quasipolynomial dependence on ε instead of an exponential one.

We also give an application to the problem of finding large subspaces in sumsets of dense sets. Green showed in [Gre05] that the sumset of a dense subset of \mathbb{F}_2^n contains a large subspace. Using Fourier analytic methods, Sanders [San11a] proved that such a subspace must have dimension $\Omega(\alpha n)$. We provide an alternative (and L^p norm-free) proof of a comparable bound, which is analogous to a recent result of Croot, Laba and Sisask [CLS11] in the integers.

*Department of Computer Science, Technion, Haifa, Israel and MIT, Cambridge, MA. eli@cs.technion.ac.il. The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement number 240258.

†Department of Computer Science, Technion, Haifa. nogaz@cs.technion.ac.il, Research supported by a scholarship from the Israel Ministry of Science and Technology and by the US-Israel Binational Science Foundation. Part of the research was conducted while the author was a visiting researcher at Microsoft Research New England, Cambridge, MA. The research leading to these results has also received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement number 257575.

‡Toyota Technological Institute at Chicago.

§Centre de Mathématiques Laurent Schwartz, École Polytechnique, 91128 Palaiseau, France.

Contents

1	Introduction	1
2	Preliminaries	3
3	Sampling-based proofs of almost-periodicity results	4
3.1	Croot-Sisask almost-periodicity	5
3.2	Almost-periodicity over a subspace	8
4	An improved quadratic Goldreich-Levin theorem	10
4.1	The quasipolynomial Bogolyubov-Ruzsa lemma	13
4.2	The quasipolynomial Bogolyubov-Ruzsa lemma – algorithmic version	14
5	Sumsets of dense sets contain large subspaces	19
A	Appendix: Proof of Corollary 5.3	24

1 Introduction

When Croot and Sisask introduced “A probabilistic technique for finding almost-periods of convolutions” in 2009 [CS10], it created quite a splash in the additive combinatorics community. Roughly speaking, their main result says that if $A \subseteq \mathbb{F}_2^n$ is a set whose sumset $A + A = \{a + a' : a, a' \in A\}$ is small, then there exists a dense set T such that the convolution $\mathbb{1}_A * \mathbb{1}_A(\cdot)$ of the indicator function of A with itself and its translate $\mathbb{1}_A * \mathbb{1}_A(\cdot + t)$ are almost indistinguishable in the L^2 norm (or higher L^p norms) for all $t \in T$. This set T may then be referred to as the set of “almost-periods”.

Croot and Sisask’s original proof used a simple sampling technique combined with tailbounds for a multinomial distribution, which Sanders replaced by the Marcinkiewicz-Zygmund inequality. Both made crucial use of L^p norms, where in applications p is taken to be very large (a function of the density α of a set $A \subseteq \mathbb{F}_2^n$ under investigation, such as $\log \alpha^{-1}$).

Here we give a different proof of the Croot-Sisask lemma that proceeds entirely without recourse to L^p norms, instead only relying on Chernoff-type tail estimates for sampling. It is our hope that this proof will appeal to a larger part of the theoretical computer science community than the currently existing ones, thereby increasing the likelihood of further novel applications of this lemma.

In the present paper we illustrate the use of this new technique by new and simplified proofs of several known results as well as an algorithmic application. Let us describe these in more detail.

Applications. In its original form, the Bogolyubov-Ruzsa lemma states that if $A \subseteq \mathbb{F}_2^n$ is a set of density α , then $4A := A + A + A + A$ contains a subspace of codimension at most $2\alpha^{-2}$. One of the first applications Croot and Sisask gave of their new technique was a *weak* Bogolyubov-Ruzsa lemma, which asserted the existence of iterated sumsets of a dense set inside $4A$. It was quickly recognized by Sanders [San10] that the latter result could be boot-strapped, using a little Fourier analysis, to a quasipolynomial version of the Bogolyubov-Ruzsa lemma in which the codimension of the subspace that is found within $4A$ is polylogarithmic in the density of the set A (so the size of this subspace is quasipolynomial in the density). This result has important implications

for the bounds in Freiman’s theorem which describes the structure of sets of integers with small sumsets [Ruz99], and the inverse theorem for the Gowers U^3 norm [GT08]. It was also a crucial ingredient in Sanders’s groundbreaking upper bound of $C(\log \log N)^5 N / \log N$ for the size of a subset of $\{1, \dots, N\}$ not containing any 3-term arithmetic progressions [San11b].

In Section 4.1 we give a straightforward proof of Sanders’s quasipolynomial Bogolyubov-Ruzsa lemma, specifically adapted to the setting of \mathbb{F}_2^n , which avoids the use of higher-order L^p norms and instead relies exclusively on Chernoff-type tail bounds.

Next we present an algorithmic application. The original motivation for this paper lies with work by the last two authors on quadratic decomposition theorems. The aim of such theorems is to decompose any bounded function $f : \mathbb{F}_2^n \rightarrow \mathbb{C}$ as a sum $g + h$, where g is quadratically uniform, in the sense that the Gowers U^3 norm $\|g\|_{U^3}$ is small, and h is quadratically structured, in the sense that it is a bounded sum of quadratically structured objects. These types of decompositions constitute a higher-order analogue of classical Fourier decompositions, and had previously been obtained in an abstract and non-constructive way (either using a form of the Hahn-Banach theorem [GW12], or a so-called energy increment approach [Gre07]).

In [TW11], the authors gave a probabilistic algorithm that, given any function $f : \mathbb{F}_2^n \rightarrow \mathbb{C}$, would with high probability compute, in time polynomial in n , a quadratic decomposition for that function with a specified U^3 error ε . This essentially amounts to computing a “quadratic Fourier decomposition” for f , and was therefore termed a *quadratic Goldreich-Levin theorem* in analogy with the well-known linear case [GL89]. The quadratic Goldreich-Levin algorithm consisted of two parts: a deterministic part which is able to construct the quadratically structured part of f under the assumption that we have an algorithm which provides some quadratic phase functions that f correlates with (if there is no such phase function, we just set $g = f$). The algorithm for finding a quadratic phase function, which constitutes the second part of the overall algorithm, is basically an algorithmic version of the proof of the inverse theorem for the U^3 norm, which states that if a bounded function f has large U^3 norm, then it correlates with a quadratic phase.

As stated above, the Bogolyubov-Ruzsa lemma is crucial in the proof of the inverse theorem, and it should not come as surprise that a new proof with a quantitative improvement has implications for the efficiency of the quadratic Goldreich-Levin algorithm outlined above. In Section 4.2 we tie the techniques developed in the earlier sections of the paper into the algorithm given in [TW11] to obtain an improvement in the running time (from exponential to quasipolynomial in the quadratic uniformity parameter ε) as well as in the number of terms that are obtained in the final quadratic decomposition (with a similar improvement in the dependence on ε). One of the main difficulties, encountered already in [TW11], is that the individual subroutines in the quadratic Goldreich-Levin algorithm, which correspond to algorithmic versions of theorems in additive combinatorics, are probabilistic in nature. Since they are applied in sequence, this means that the input for the next subroutine comes with a certain amount of noise, and it is therefore necessary to prove robust algorithmic versions of the theorems from additive combinatorics. This applies in particular to the quasipolynomial Bogolyubov-Ruzsa lemma, for which we give robust version in this paper. The detailed introduction of the key concepts in additive combinatorics and quadratic Fourier analysis is postponed to the start of Section 4.

Our final application concerns the problem of finding large subspaces within sumsets of a dense set. Green [Gre05] had shown that if $A \subseteq \mathbb{F}_2^n$ has density α , then $A + A$ contains a subspace of dimension $\Omega(\alpha^2 n)$. Sanders proved in [San11a] using a Fourier-iteration lemma that this subspace must be of dimension at least $\Omega(\alpha n)$, and remarked that a bound of comparable strength follows implicitly from the techniques of Croot, Laba and Sisask [CLS11], who addressed the more general

problem in the integers, asking for long arithmetic progressions in sumsets of dense sets.

In Section 5 we provide a simplified proof of the Croot-Laba-Sisask bound, again avoiding Fourier analysis and using instead our sampling approach to Croot-Sisask almost-periodicity in \mathbb{F}_2^n . It requires a more careful analysis of our sampling technique, which we shall give in detail in the appendix. We do not address the question in the integers, nor the non-abelian case, of which this is a toy version.

Chernoff vs. L^p norms. It is of course well known that L^p bounds and Chernoff’s inequality are, in a certain sense, equivalent. Specifically, a random variable X obeys a Chernoff-type tail bound of the form

$$\mathbb{P}[|X| \geq t\|X\|_2] \leq C \exp(-\Omega(t^2))$$

if and only if its L^p norm satisfies

$$\|X\|_p \leq C\sqrt{p}\|X\|_2$$

for all $p \in [2, \infty)$, the latter representing a Khinchine-type inequality (from which Marcinkiewicz-Zygmund can be derived). For a proof of this statement we refer the reader to the excellent lecture notes by Sanders [San12].

We therefore do not claim that our proof of Croot and Sisask’s almost-periodicity results is radically new. However, we do think it writes itself rather naturally in the special case of \mathbb{F}_2^n , and it lends itself more readily to applications in that setting. Moreover, these results can also be “algorithmified” in what is in our opinion a more natural way.

Acknowledgements. The last two authors would like to thank Tom Sanders for numerous helpful remarks and discussions. The first two authors would like to thank Shachar Lovett for sharing with them his view of Sanders’s proof of the quasipolynomial Bogolyubov-Ruzsa lemma (cf. [Lov12]).

2 Preliminaries

In this section we fix our notation and collect some results that we shall use throughout the paper. Fundamental to our approach will be the following Chernoff-type tail bound for sampling [TV06].

Lemma 2.1 (Hoeffding bound for sampling) *If \mathbf{X} is a random variable with $|\mathbf{X}| \leq 1$ and $\hat{\mu}$ is the empirical average obtained from t samples, then*

$$\mathbb{P}[|\mathbb{E}[\mathbf{X}] - \hat{\mu}| > \gamma] \leq 2 \exp(-2\gamma^2 t).$$

Throughout the paper we shall make use of the discrete Fourier transform, which we define as follows. For $f : \mathbb{F}_2^n \rightarrow \mathbb{C}$, let

$$\widehat{f}(t) = \mathbb{E}_{x \in \mathbb{F}_2^n} f(x)(-1)^{x \cdot t}$$

for any $t \in \widehat{\mathbb{F}_2^n} = \mathbb{F}_2^n$, where $\mathbb{E}_{x \in \mathbb{F}_2^n}$ simply stands for the normalized sum $2^{-n} \sum_{x \in \mathbb{F}_2^n}$ and $x \cdot t = \sum_{i=1}^n x_i t_i$ for a pair of vectors $x = (x_1, \dots, x_n)$, $t = (t_1, \dots, t_n)$. The inversion formula states that

$$f(x) = \sum_{t \in \mathbb{F}_2^n} \widehat{f}(t)(-1)^{x \cdot t}$$

for all $x \in \mathbb{F}_2^n$, and Parseval’s identity takes the form

$$\langle f, g \rangle = \langle \widehat{f}, \widehat{g} \rangle,$$

where the inner product is defined as $\langle f, g \rangle = \mathbb{E}_{x \in \mathbb{F}_2^n} f(x) \overline{g(x)}$ in physical space, and $\langle \widehat{f}, \widehat{g} \rangle = \sum_{t \in \mathbb{F}_2^n} \widehat{f}(t) \overline{\widehat{g}(t)}$ on frequency space, for any two functions $f, g : \mathbb{F}_2^n \rightarrow \mathbb{C}$.

Finally, the convolution of two such functions is defined by

$$f * g(x) = \mathbb{E}_{y \in \mathbb{F}_2^n} f(y) g(x - y),$$

and the fact that the Fourier transform diagonalizes the convolution operator is expressed via the identity

$$\widehat{f * g}(t) = \widehat{f}(t) \widehat{g}(t),$$

which holds for all $t \in \mathbb{F}_2^n$.

The set of large Fourier coefficients determines the value of a function to a significant extent, and for many arguments it is important to be able to estimate its size and determine its structure. For a function $f : \mathbb{F}_2^n \rightarrow \mathbb{C}$, let

$$\text{Spec}_\rho(f) = \{t \in \mathbb{F}_2^n : |\widehat{f}(t)| \geq \rho \|f\|_1\}. \quad (1)$$

For a subset $A \subseteq \mathbb{F}_2^n$ we let $\mathbb{1}_A$ denote the indicator function of A and μ_A denote the function $\mathbb{1}_A \cdot (2^n/|A|)$ so that $\mathbb{E}_{x \in \mathbb{F}_2^n} [\mu_A(x)] = 1$. In the special case where $f = \mathbb{1}_A$ for a subset $A \subseteq \mathbb{F}_2^n$ of density α , Parseval's identity tells us that $|\text{Spec}_\rho(\mathbb{1}_A)| \leq \rho^{-2} \cdot \alpha^{-1}$. A more precise result is known: Chang's theorem [Cha02] states that $\text{Spec}_\rho(\mathbb{1}_A)$ is in fact contained in a subspace of dimension at most $C\rho^{-2} \log \alpha^{-1}$.

Theorem 2.2 (Chang's theorem) *Let $\rho \in (0, 1]$ and $A \subseteq \mathbb{F}_2^n$. Then there is a subspace V of \mathbb{F}_2^n such that $\text{Spec}_\rho(\mathbb{1}_A) \subseteq V$ and*

$$\dim(V) \leq 8 \frac{\log(2^n/|A|)}{\rho^2}.$$

For an elegant recent proof of this result using entropy, see Impagliazzo et al. [IMR12].

Finally, for two real numbers α, β we write $\alpha \approx_\varepsilon \beta$ to denote $|\alpha - \beta| \leq \varepsilon$ and if $|\alpha - \beta| > \varepsilon$ we write $\alpha \not\approx_\varepsilon \beta$. All logarithms in this paper are taken to base 2.

3 Sampling-based proofs of almost-periodicity results

For comparison, we give the precise statement of the original result of Croot and Sisask (Proposition 1.3 in [CS10]). Since it is valid for general groups G , it is written in multiplicative notation.

Proposition 3.1 (Croot-Sisask Lemma , L^p local version) *Let $\varepsilon > 0$ and let $m \geq 1$ be an integer. Let G be a group and let $A, B \subseteq G$ be finite subsets such that $|B \cdot A| \leq K|B|$. Then there is a set $X \subseteq A$ of size $|X| \geq |A|/(2K)^{50m/\varepsilon}$ such that for each $x \in X X^{-1}$,*

$$\|\mathbb{1}_A * \mathbb{1}_B(yx) - \mathbb{1}_A * \mathbb{1}_B(y)\|_{2m}^{2m} \leq \max\{\varepsilon^m |AB| |B|^m, \|\mathbb{1}_A * \mathbb{1}_B\|_m^m \varepsilon^m |B|^m\}.$$

In the next section we prove our version of the Croot-Sisask lemma, given as Proposition 3.2 below. In Section 3.2 we give a modified version in which the resulting set of almost-periods, appearing as XX^{-1} in Proposition 3.1 above, can in fact taken to be a subspace, which is what we need in applications.

3.1 Croot-Sisask almost-periodicity

To state our result concisely, we define the following measure which will play a central role in what follows. Given subsets $A, B \subset \mathbb{F}_2^n$ where A is finite, define the *measure of additive containment* $\rho_{A \rightarrow B} : \mathbb{F}_2^n \rightarrow [0, 1]$ by

$$\rho_{A \rightarrow B}(y) := \mathbb{P}_{a \in A} [y + a \in B] = \frac{|(y + A) \cap B|}{|A|} = \mu_A * \mathbf{1}_B(y), \quad (2)$$

for each $y \in \mathbb{F}_2^n$. Notice that $\rho_{A \rightarrow B}(y) = 1$ when $y + A \subseteq B$ and $\rho_{A \rightarrow B}(y) = 0$ when $(y + A) \cap B = \emptyset$.

Proposition 3.2 (Almost-periodicity of sumsets) *If $A \subset \mathbb{F}_2^n$ satisfies $|2A| \leq K|A|$, then for every integer t and set $B \subseteq \mathbb{F}_2^n$ there exists a set X with the following properties.*

1. The set X is contained in an affine shift of A .
2. The size of X is at least $|A|/(2K^{t-1})$.
3. For all $x \in X$ and for all subsets $S \subseteq \mathbb{F}_2^n$,

$$\mathbb{P}_{y \in S} [\rho_{A \rightarrow B}(y) \approx_{2\varepsilon} \rho_{A \rightarrow B}(y + x)] \geq 1 - 8 \frac{|A + B|}{|S|} \cdot \exp(-2\varepsilon^2 t). \quad (3)$$

Our proof differs from the original proof of Croot and Sisask in that it disposes of L^p -norms and tail bounds for a multinomial distribution (or the Marcinkiewicz-Zygmund inequality), and replaces them with sampling arguments relying on the Chernoff-Hoeffding bound.

We sketch the proof before giving the technical details. To obtain X we replace $\rho_{A \rightarrow B}$ by an estimator function computed by taking a sequence of t independent random samples distributed uniformly over A . Denoting the sample sequence by $\mathbf{a} = (a_1, \dots, a_t)$, we estimate $\rho_{A \rightarrow B}(y)$ by the fraction of $a_i \in \mathbf{a}$ satisfying $y + a_i \in B$. Denote the estimator function corresponding to \mathbf{a} by $\hat{\rho}_{\mathbf{a}}$. Fixing y , the Chernoff-Hoeffding bound says that the probability that $\hat{\rho}_{\mathbf{a}}(y)$ differs from $\rho_{A \rightarrow B}(y)$ by more than ε , i.e., the probability of the event “ $\rho_{A \rightarrow B}(y) \not\approx_{\varepsilon} \hat{\rho}_{\mathbf{a}}(y)$ ” when $\mathbf{a} = (a_1, \dots, a_t)$ is distributed uniformly over A^t , is at most $\exp(-\Omega(\varepsilon^2 t))$.

The key observation in the construction of the set X is that there are many pairs of good estimator-sequences $\mathbf{a} = (a_1, \dots, a_t)$, $\hat{\mathbf{a}} = (\hat{a}_1, \dots, \hat{a}_t)$ for which there exists a “special” element $x \in \mathbb{F}_2^n$ such that $\hat{\mathbf{a}} = x + \mathbf{a}$, where $x + \mathbf{a} := (x + a_1, \dots, x + a_t)$. Such x can be justly called “special” for the following reason. Call y “good” if both of the following conditions hold,

$$\rho_{A \rightarrow B}(y) \approx_{\varepsilon} \hat{\rho}_{\mathbf{a}}(y) \quad \text{and} \quad \rho_{A \rightarrow B}(y + x) \approx_{\varepsilon} \hat{\rho}_{\mathbf{a}}(y + x). \quad (4)$$

Now if $\hat{\mathbf{a}} = x + \mathbf{a}$, then we have

$$\hat{\rho}_{\hat{\mathbf{a}}}(y) = \hat{\rho}_{\hat{\mathbf{a}}}(y + x + x) = \hat{\rho}_{\hat{\mathbf{a}}+x}(y + x) = \hat{\rho}_{\mathbf{a}}(y + x),$$

and combining this with (4) implies that for “good” y we have $\rho_{A \rightarrow B}(y) \approx_{2\varepsilon} \rho_{A \rightarrow B}(y + x)$. Thus, to prove the proposition we only need to bound from below the number of “special” elements x , which is done based on the assumption that A has small doubling. We now give the formal proof.

Proof of Proposition 3.2: To simplify notation let $\rho(y) := \rho_{A \rightarrow B}(y)$, and for a sequence $\mathbf{a} = (a_1, \dots, a_t)$ of length t , define the \mathbf{a} -estimator of ρ to be the function $\hat{\rho}_{\mathbf{a}} : \mathbb{F}_2^n \rightarrow [0, 1]$ defined for $y \in \mathbb{F}_2^n$ by

$$\hat{\rho}_{\mathbf{a}}(y) := \frac{|\{y + a_i \in B \mid i = 1, \dots, t\}|}{t}.$$

We say that \mathbf{a} is an ε -good estimator for y if $\rho(y) \approx_\varepsilon \hat{\rho}_{\mathbf{a}}(y)$.

Fix $y \in \mathbb{F}_2^n$. Our first step towards constructing X is to show that most sample-sequences from A are ε -good for y , provided that t , the sample size, is large enough with respect to $1/\varepsilon$. Let Y_i be the indicator random variable for the event “ $y + a_i \in B$ ” when a_i is chosen uniformly at random from A . Then $\hat{\rho}_{\mathbf{a}}(y) = \frac{1}{t} \sum_{i=1}^t Y_i$ is the average of t i.i.d. indicator random variables each having mean $\rho(y)$, so the Chernoff-Hoeffding bound (Lemma 2.1) implies that for each $y \in \mathbb{F}_2^n$,

$$\mathbb{P}_{\mathbf{a} \in A^t} [\rho(y) \not\approx_\varepsilon \hat{\rho}_{\mathbf{a}}(y)] \leq 2 \exp(-2\varepsilon^2 t). \quad (5)$$

Now we proceed to show that most $\mathbf{a} \in A^t$ are ε -good estimators for most y . Let $Z_{\mathbf{a}}$ be the random variable measuring the fraction of $y \in A + B$ for which \mathbf{a} is an ε -good estimator, that is,

$$Z_{\mathbf{a}} := \mathbb{P}_{y \in A+B} [\rho(y) \approx_\varepsilon \hat{\rho}_{\mathbf{a}}(y)].$$

Setting $\delta = 2 \exp(-2\varepsilon^2 t)$, we conclude from (5) via linearity of expectation that

$$\mathbb{E}_{\mathbf{a} \in A^t} [Z_{\mathbf{a}}] \geq 1 - \delta.$$

Markov’s inequality now shows that at least half of the sequences $\mathbf{a} \in A^t$ are ε -good estimators for all but a (2δ) -fraction of $y \in A + B$, in which case we say that \mathbf{a} is an $(\varepsilon, 2\delta)$ -good estimator for ρ . Denote by $\mathbf{G}[\varepsilon, 2\delta] \subset A^t$ the set of these sequences,

$$\mathbf{G}[\varepsilon, 2\delta] = \left\{ \mathbf{a} \in A^t \mid \mathbb{P}_{y \in A+B} [\rho(y) \approx_\varepsilon \hat{\rho}_{\mathbf{a}}(y)] \geq 1 - 2\delta \right\}.$$

To obtain X we partition $\mathbf{G}[\varepsilon, 2\delta]$ as follows. Define a mapping $\varphi : A^t \mapsto \{0\} \times (2A)^{t-1}$ by shifting a sequence $\mathbf{a} = (a_1, \dots, a_t)$ by its first element a_1 ,

$$\varphi(\mathbf{a}) = \mathbf{a} + a_1 := (a_1 + a_1, a_1 + a_2, \dots, a_1 + a_t) \quad (6)$$

Then φ maps the set $\mathbf{G}[\varepsilon, 2\delta]$, which has size at least $|A|^t/2$, into a set of size $|2A|^{t-1} \leq (K|A|)^{t-1}$ so by the pigeonhole principle, there is a subset $\mathbf{G}[\varepsilon, 2\delta]_{\mathbf{b}} \subset \mathbf{G}[\varepsilon, 2\delta]$ that is mapped to the same element $\mathbf{b} = (0, b_2, \dots, b_t)$. In addition, this subset is pretty large,

$$|\mathbf{G}[\varepsilon, 2\delta]_{\mathbf{b}}| \geq \frac{|A|^t}{2K^{t-1}|A|^{t-1}} = \frac{|A|}{2K^{t-1}}. \quad (7)$$

Finally, fix an arbitrary $\hat{\mathbf{a}} = (\hat{a}_1, \dots, \hat{a}_t) \in \mathbf{G}[\varepsilon, 2\delta]_{\mathbf{b}}$ and set

$$X = \{\hat{a}_1 + a_1 \mid (a_1, \dots, a_t) \in \mathbf{G}[\varepsilon, 2\delta]_{\mathbf{b}}\}.$$

To complete our proof we show that X has the three properties listed in the statement of the lemma.

1. By definition, $X \subseteq \hat{a}_1 + A$.
2. The mapping $\mathbf{G}[\varepsilon, 2\delta]_{\mathbf{b}} \mapsto X$ given by $(a_1, \dots, a_t) \mapsto \hat{a}_1 + a_1$ is invertible, because both \hat{a}_1 and \mathbf{b} are fixed. Hence $|X| = |\mathbf{G}[\varepsilon, 2\delta]_{\mathbf{b}}|$ and the size of X is bounded from below using (7).

3. Suppose $x = \hat{a}_1 + a_1$, where a_1 is the first element of an $(\varepsilon, 2\delta)$ -good estimator $\mathbf{a} = (a_1, \dots, a_t) \in \mathbf{G}[\varepsilon, 2\delta]_{\mathbf{b}}$. The key observation is that $\hat{\mathbf{a}} + x = \mathbf{a}$. Indeed, the definition of $\mathbf{G}[\varepsilon, 2\delta]_{\mathbf{b}}$ implies $\varphi(\hat{\mathbf{a}}) = \varphi(\mathbf{a})$, so using (6) we have

$$\hat{a}_1 + \hat{a}_i = a_1 + a_i, \quad i = 1, \dots, t,$$

which, rearranging, comes out to

$$a_i = x + \hat{a}_i, \quad i = 1, \dots, t.$$

In other words, $\hat{\mathbf{a}} + x = \mathbf{a}$ as claimed.

Recalling that $\hat{\mathbf{a}}$ is an $(\varepsilon, 2\delta)$ -good estimator, we know that for all but a 2δ -fraction of $y \in A + B$,

$$\rho(y) \approx_{\varepsilon} \hat{\rho}_{\hat{\mathbf{a}}}(y) \quad (8)$$

and (8) also holds for all $y \notin A + B$ since in this case $\hat{\rho}_{\hat{\mathbf{a}}}(y) = \rho(y) = 0$. Hence we have that (8) holds for all but a $(2\delta|A + B|/|S|)$ -fraction of $y \in S$.

Similarly, since \mathbf{a} is an $(\varepsilon, 2\delta)$ -good estimator, we have that

$$\rho(y + x) \approx_{\varepsilon} \hat{\rho}_{\mathbf{a}}(y + x) \quad (9)$$

for all but a $(2\delta|A + B|/|x + S|)$ -fraction of $y \in S$. Using a union bound and the fact that $|S + x| = |S|$, we find that for all but a $(4\delta|A + B|/|S|)$ -fraction of $y \in S$ both (8) and (9) hold. For such y we conclude that $\rho(y) \approx_{2\varepsilon} \rho(y + x)$ using the triangle inequality and the fact that $\hat{\rho}_{\hat{\mathbf{a}}}(y) = \hat{\rho}_{\hat{\mathbf{a}}+x}(y + x) = \hat{\rho}_{\mathbf{a}}(y + x)$.

This completes the proof of the proposition. ■

By an inductive application of Proposition 3.2 one can prove its following iterated version.

Corollary 3.3 (Almost-periodicity of sumsets, iterated) *If $A \subset \mathbb{F}_2^n$ satisfies $|2A| \leq K|A|$, then for every integer t and set $B \subseteq \mathbb{F}_2^n$ there exists a set X with the following properties.*

1. The set X is contained in an affine shift of A .
2. The size of X is at least $|A|/(2K^{t-1})$.
3. For all $x_1, \dots, x_\ell \in X$ and for all subsets $S \subseteq \mathbb{F}_2^n$,

$$\mathbb{P}_{y \in S} [\rho_{A \rightarrow B}(y) \approx_{2\varepsilon\ell} \rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] \geq 1 - 8\ell \frac{|A + B|}{|S|} \cdot \exp(-2\varepsilon^2 t). \quad (10)$$

Proof: The proof is by induction on ℓ . Proposition 3.2 establishes the case $\ell = 1$. For the induction step, suppose that the lemma holds for some integer $\ell \geq 1$ with a set $X \subseteq \mathbb{F}_2^n$. We shall show that the same set X satisfies the above requirements for $\ell + 1$.

Let $\delta := 8(|A + B|/|S|) \cdot \exp(-2\varepsilon^2 t)$. By the induction hypothesis, for at least a $(1 - \ell\delta)$ -fraction of $y \in S$, it is true that $\rho_{A \rightarrow B}(y) \approx_{2\varepsilon\ell} \rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)$. The case $\ell = 1$ implies that for at least a $(1 - \delta)$ -fraction of $y \in S$, it is true that $\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell) \approx_{2\varepsilon} \rho_{A \rightarrow B}(y + x_1 + \dots + x_{\ell+1})$. Thus by a union bound we have that for at least a $(1 - (\ell + 1)\delta)$ -fraction of $y \in S$ we have both $\rho_{A \rightarrow B}(y) \approx_{2\varepsilon\ell} \rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)$ and $\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell) \approx_{2\varepsilon} \rho_{A \rightarrow B}(y + x_1 + \dots + x_{\ell+1})$. The proof is completed by noting that by the triangle inequality, for each such y , we also have $\rho_{A \rightarrow B}(y) \approx_{2\varepsilon(\ell+1)} \rho_{A \rightarrow B}(y + x_1 + \dots + x_{\ell+1})$. ■

3.2 Almost-periodicity over a subspace

For applications one would like a version of Proposition 3.2 in which the set X of periods is in fact a subspace. It was observed by Sanders [San10] that one can use iterated almost-periodicity statements such as Corollary 3.3, combined with some Fourier analysis, to obtain such a subspace. In this section, we use Sanders’s argument to deduce the following statement from Corollary 3.3.

Corollary 3.4 (Almost-periodicity of sumsets over a subspace) *If $A \subset \mathbb{F}_2^n$ is a subset of density α , then for every integer t and set $B \subseteq \mathbb{F}_2^n$ there exists a subspace V of codimension $\text{codim}(V) \leq 32 \log(2/\alpha^t)$ with the following property.*

For every $v \in V$, for all subsets $S \subseteq \mathbb{F}_2^n$ and for every $\varepsilon, \eta > 0$ and integer ℓ ,

$$\mathbb{P}_{y \in S} [\rho_{A \rightarrow B}(y) \approx_{\varepsilon'} \rho_{A \rightarrow B}(y + v)] \geq 1 - 16 \frac{\ell}{\eta} \frac{|A + B|}{|S|} \cdot \exp(-2\varepsilon^2 t), \quad (11)$$

where $\varepsilon' = 4\varepsilon\ell + 2\eta + 2^{-\ell} \sqrt{|B|/|A|}$.

As we shall see in Sections 4 and 5, the proof of the quasipolynomial Bogolyubov-Ruzsa lemma (Theorem 4.8) follows easily from the above lemma, and Green’s theorem on the existence of subspaces in sumsets of dense sets (Theorem 5.2) follows easily from a refinement of the above corollary (which we will give as Corollary 5.3 below). Note that for the proof of Corollary 3.4 we need the stronger assumption that A has density at least α in \mathbb{F}_2^n , instead of the doubling hypothesis $|2A| \leq K|A|$.

The idea of the proof of Corollary 3.4 is the following. Let X be the subset guaranteed by Corollary 3.3 for $K = 1/\alpha$, and define the subspace V as $V = \text{Spec}_{1/2}(X)^\perp$ (see Section 2 for the definition of Spec_ρ). The intuition is that if X were a subspace then $\text{Spec}_{1/2}(X) = V^\perp$, and hence $V = X$. Thus V serves as an “approximate subspace” for X . Since A is dense in \mathbb{F}_2^n , by Corollary 3.3 we also have that X is dense in \mathbb{F}_2^n and hence Chang’s theorem (Theorem 2.2) implies that the subspace V is also dense in \mathbb{F}_2^n (this is the only place where we need the stronger assumption on the density of A).

In order to show that (11) holds we first show, using Corollary 3.3, a simple averaging argument and the triangle inequality, that for most $y \in S$,

$$\mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] \approx_{2\varepsilon\ell + \eta} \rho_{A \rightarrow B}(y). \quad (12)$$

Similarly, for all $v \in V$ and for most $y \in S$,

$$\mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)] \approx_{2\varepsilon\ell + \eta} \rho_{A \rightarrow B}(y + v). \quad (13)$$

We then use Fourier analysis, following Sanders’s argument closely, to show that for all $y \in \mathbb{F}_2^n$,

$$\mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] \approx_{2^{-\ell} \sqrt{|B|/|A|}} \mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)], \quad (14)$$

where v is again an arbitrary element of V . The final conclusion follows from (12), (13) and (14) using the union bound and the triangle inequality. We start by establishing (12) and (13).

Lemma 3.5 *Let $\varepsilon, \delta > 0$, and let $A, B, X, S \subseteq \mathbb{F}_2^n$ be such that for all $x_1, \dots, x_\ell \in X$,*

$$\mathbb{P}_{y \in S} [\rho_{A \rightarrow B}(y) \approx_\varepsilon \rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] \geq 1 - \delta.$$

Then for every $\eta > 0$ we have that

$$\mathbb{P}_{y \in S} \left[\rho_{A \rightarrow B}(y) \approx_{\varepsilon + \eta} \mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] \right] \geq 1 - \delta/\eta.$$

Proof: From Markov's inequality it follows that for at least a $(1 - \delta/\eta)$ -fraction of $y \in S$, the relation

$$\rho_{A \rightarrow B}(y) \approx_\varepsilon \rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)$$

holds for at least a $(1 - \eta)$ -fraction of ℓ -tuples $(x_1, \dots, x_\ell) \in X^\ell$. Thus for at least a $(1 - \delta/\eta)$ -fraction of $y \in S$, we have that

$$\left| \mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] - \rho_{A \rightarrow B}(y) \right| \leq \mathbb{E}_{x_1, \dots, x_\ell \in X} [|\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell) - \rho_{A \rightarrow B}(y)|]$$

which is seen to be bounded above by $(1 - \eta) \cdot \varepsilon + \eta \cdot 1 \leq \varepsilon + \eta$. ■

The next lemma establishes (14).

Lemma 3.6 *Let $X \subseteq \mathbb{F}_2^n$, and let $V = \text{Spec}_{1/2}(X)^\perp$. Then for all $y \in \mathbb{F}_2^n$ and $v \in V$,*

$$\mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] \approx_{\varepsilon''} \mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)], \quad (15)$$

where $\varepsilon'' = 2^{-\ell} \sqrt{|B|/|A|}$.

Proof: We can write the difference between the two sides of (15) using the convolution operator as

$$(\mu_X)^{\ast \ell} * \mu_A * \mathbf{1}_B(y) - (\mu_X)^{\ast \ell} * \mu_A * \mathbf{1}_B(y + v),$$

which in terms of the Fourier basis equals

$$\sum_{t \in \mathbb{F}_2^n} \hat{\mu}_A(t) \cdot (\hat{\mu}_X(t))^\ell \cdot \widehat{\mathbf{1}}_B(t) \cdot \left((-1)^{y \cdot t} - (-1)^{(y+v) \cdot t} \right).$$

This expression in turn is bounded in absolute value by

$$\sum_{t \in \mathbb{F}_2^n} |\hat{\mu}_A(t)| \cdot |\hat{\mu}_X(t)|^\ell \cdot \left| \widehat{\mathbf{1}}_B(t) \right| \cdot |(-1)^{y \cdot t}| \cdot |1 - (-1)^{v \cdot t}| = \sum_{t \in \mathbb{F}_2^n} |\hat{\mu}_A(t)| \cdot |\hat{\mu}_X(t)|^\ell \cdot \left| \widehat{\mathbf{1}}_B(t) \right| \cdot |1 - (-1)^{v \cdot t}|.$$

By definition of V as the orthogonal complement of $\text{Spec}_{1/2}(X)$, the right-hand side can be bounded as

$$\sum_{t \notin V^\perp} |\hat{\mu}_A(t)| \cdot |\hat{\mu}_X(t)|^\ell \cdot \left| \widehat{\mathbf{1}}_B(t) \right| \cdot |1 - (-1)^{v \cdot t}| \leq 2^{-\ell} \sum_{t \notin V^\perp} |\hat{\mu}_A(t)| \cdot \left| \widehat{\mathbf{1}}_B(t) \right|.$$

By the Cauchy-Schwarz inequality and Parseval's identity, this is bounded above by

$$2^{-\ell} \sqrt{\sum_{t \notin V^\perp} (\hat{\mu}_A(t))^2} \sqrt{\sum_{t \notin V^\perp} (\widehat{\mathbf{1}}_B(t))^2} \leq 2^{-\ell} \sqrt{\mathbb{E}_{y \in \mathbb{F}_2^n} (\mu_A(y))^2} \sqrt{\mathbb{E}_{y \in \mathbb{F}_2^n} (\mathbf{1}_B(y))^2} = 2^{-\ell} \sqrt{|B|/|A|}.$$

■

We are now ready for the proof of Corollary 3.4.

Proof of Corollary 3.4: Let X be the set guaranteed by Corollary 3.3 for $K = 1/\alpha$, and let $V = \text{Spec}_{1/2}(X)^\perp$.

First, note that Property 1 of Corollary 3.3 implies that $|X| \geq |A|/(2(1/\alpha)^{t-1}) \geq \alpha^t \cdot 2^{n-1}$. It now follows from Chang’s theorem (Theorem 2.2) that

$$\dim(\text{Spec}_{1/2}(X)) = \text{codim}(V) \leq 8 \frac{\log(2/\alpha^t)}{(1/2)^2} = 32 \log(2/\alpha^t).$$

It remains to show that (11) holds.

Let $\delta := 8\ell(|A + B|/|S|) \cdot \exp(-2\varepsilon^2 t)$. From Corollary 3.3 and Lemma 3.5 we have that

$$\rho_{A \rightarrow B}(y) \approx_{2\varepsilon\ell + \eta} \mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] \quad (16)$$

for at least a $(1 - \delta/\eta)$ -fraction of $y \in S$, and similarly that for all $v \in V$,

$$\rho_{A \rightarrow B}(y + v) \approx_{2\varepsilon\ell + \eta} \mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)] \quad (17)$$

for at least a $(1 - \delta/\eta)$ -fraction of $y \in S$. Moreover, Lemma 3.6 implies that for every $y \in S$ and $v \in V$,

$$\mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] \approx_{2^{-\ell} \sqrt{|B/A|}} \mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)]. \quad (18)$$

Applying the union bound and the triangle inequality to (16), (17) and (18), we conclude that

$$\rho_{A \rightarrow B}(y) \approx_{\varepsilon'} \rho_{A \rightarrow B}(y + v)$$

for $\varepsilon' = 4\varepsilon\ell + 2\eta + 2^{-\ell} \sqrt{|B|/|A|}$ for at least a $(1 - 2\delta/\eta)$ -fraction of $y \in S$, which is the desired conclusion. \blacksquare

4 An improved quadratic Goldreich-Levin theorem

Both in number theory and theoretical computer science, there are certain situations where we may wish to decompose a bounded function $f : \mathbb{F}_2^n \rightarrow \mathbb{C}$ as a sum $g + h$, where g is a “uniform” or “random-looking”, and h is a somewhat “structured” part. Such situations include the counting of arithmetic progressions [Gre07], the analysis of Probabilistically Checkable Proofs (PCPs) [ST06] and the approximation of matrices and tensors [FK99].

In the case where one is looking for “linear uniformity” in the function g , for example when counting arithmetic progressions of length 3, such a decomposition is achieved by separating large and small Fourier coefficients (corresponding to “linearly structured” and “linearly uniform” parts, respectively). This task can be handled algorithmically by the Goldreich-Levin theorem ([GL89], see Theorem 4.1 below), which provides an algorithm that computes, with high probability, the large Fourier coefficients of $f : \mathbb{F}_2^n \rightarrow \{-1, 1\}$ in time polynomial in n .

Theorem 4.1 (Goldreich-Levin Theorem) *Let $\nu, \delta > 0$. There is a randomized algorithm which, given oracle access to a function $f : \mathbb{F}_2^n \rightarrow \{-1, 1\}$, runs in time $O(n^2 \log n \cdot \text{poly}(1/\nu, \log(1/\delta)))$ and outputs a decomposition*

$$f = \sum_{i=1}^k c_i \cdot (-1)^{\langle \alpha_i, x \rangle} + g$$

with the following guarantee.

- $k = O(1/\nu^2)$.
- $\mathbb{P} \left[\exists i |c_i - \hat{f}(\alpha_i)| > \nu/2 \right] \leq \delta$.
- $\mathbb{P} \left[\forall \alpha \text{ such that } |\hat{f}(\alpha)| \geq \nu, \exists i \alpha_i = \alpha \right] \geq 1 - \delta$.

However, these linear decompositions have been shown to not be sensitive enough to handle many other situations, such as the counting of arithmetic progressions of length 4. In the latter case, one instead needs the function g to be “quadratically uniform” in the sense of Gowers [Gow98]. We say that a function g is *quadratically uniform* if it is small in the U^3 norm, which is defined by the formula

$$\|g\|_{U^3}^8 = \mathbb{E}_{x, h_1, h_2, h_3 \in G} \prod_{\omega \in \{0,1\}^3} C^{|\omega|} g(x + \omega \cdot h),$$

where $\omega \cdot h$ is shorthand for $\sum_i \omega_i h_i$, and $C^{|\omega|} g = g$ if $\sum_i \omega_i$ is even and \bar{g} otherwise.

A hint as to what might constitute the *quadratically structured* part of a decomposition in which g is quadratically uniform is given by the so-called inverse theorem for the U^3 norm, whose proof was largely contained in Gowers’s proof of Szemerédi’s theorem but brought to the point by Samorodnitsky [Sam07] (in the case of characteristic 2) and by Green and Tao [GT08]. It states, qualitatively speaking, that a function with large U^3 norm correlates with a *quadratic phase function*, by which we mean a function of the form $(-1)^q$ for a quadratic form $q : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$.

The inverse theorem implies that the structured part h has quadratic structure in the case where g is small in U^3 , and starting with [Gre07] a variety of such *quadratic decomposition theorems* have come into existence: in one formulation [GW12], one can write f as

$$f = \sum_i \lambda_i (-1)^{q_i} + g + l, \tag{19}$$

where the q_i are quadratic forms, the λ_i are real coefficients such that $\sum_i |\lambda_i|$ is bounded, $\|g\|_{U^3}$ is small and l is a small ℓ_1 error (which is negligible in all known applications). Such a decomposition is not unique and non-trivial since the quadratic phases ω^q , unlike linear exponentials, do not form an orthonormal basis. In analogy with the decomposition into Fourier characters, it is natural to think of the coefficients λ_i as the *quadratic Fourier coefficients* of f .

An algorithmic version of a *quadratic decomposition theorem* was given by the last two authors in [TW11]. Prior to [TW11], all quadratic decomposition theorems proved had been of a rather abstract nature. In particular, work by Trevisan, Vadhan and the third author [TTV09] used linear programming techniques and boosting, while Gowers and the last author [GW12] gave a (non-constructive) existence proof using the Hahn-Banach theorem. The main result of [TW11] then was the following.

Theorem 4.2 (Quadratic Goldreich-Levin theorem) *Let $\varepsilon, \delta > 0$, $n \in \mathbb{N}$ and $B > 1$. Then there exists $\eta = \exp(-(B/\varepsilon)^C)$ and a randomized algorithm running in time $O(n^4 \log n \cdot \text{poly}(1/\eta, \log(1/\delta)))$ which, given any function $f : \mathbb{F}_2^n \rightarrow [-1, 1]$ as an oracle, outputs with probability at least $1 - \delta$ a decomposition into quadratic phases*

$$f = c_1(-1)^{q_1} + \dots + c_k(-1)^{q_k} + g + l$$

satisfying $k \leq 1/\eta^2$, $\|g\|_{U^3} \leq \varepsilon$, $\|l\|_1 \leq 1/2B$ and $|c_i| \leq \eta$ for all $i = 1, \dots, k$.

The algorithm comprised two parts. The first was a (entirely deterministic) procedure for assembling the quadratic phases with which the function f correlates into an actual decomposition, if these quadratic phases can indeed be found.

Theorem 4.3 *Let \mathcal{Q} be a class of functions as above and let $\varepsilon, \delta > 0$ and $B > 1$. Let A be an algorithm which, given oracle access to a function $f : X \rightarrow [-B, B]$ satisfying $\|f\|_S \geq \varepsilon$, outputs, with probability at least $1 - \delta$, a function $\bar{q} \in \mathcal{Q}$ such that $\langle f, \bar{q} \rangle \geq \eta$ for some $\eta = \eta(\varepsilon, B)$. Then there exists an algorithm which, given any function $f : X \rightarrow [-1, 1]$, outputs with probability at least $1 - \delta/\eta^2$ a decomposition*

$$f = c_1 \bar{q}_1 + \dots + c_k \bar{q}_k + g + l$$

satisfying $k \leq 1/\eta^2$, $\|g\|_S \leq \varepsilon$, $\|l\|_1 \leq 1/2B$ and $|c_i| \leq \eta$ for all $i = 1, \dots, k$.

The algorithm makes at most k calls to A .

Theorem 4.3 is proved using a boosting argument, for which we refer the reader to [TW11]. The other key component in the quadratic Goldreich-Levin algorithm was the following self-correction procedure for Reed-Muller codes of order 2 (which are simply truth-tables of quadratic phase functions).

Theorem 4.4 *Given $\varepsilon, \delta > 0$, there exists $\eta = \exp(-1/\varepsilon^C)$ and a randomized algorithm **Find-Quadratic** running in time $O(n^4 \log n \cdot \text{poly}(1/\varepsilon, 1/\eta, \log(1/\delta)))$ which, given oracle access to a function $f : \mathbb{F}_2^n \rightarrow \{-1, 1\}$, either outputs a quadratic form $q(x)$ or \perp . The algorithm satisfies the following guarantee.*

- *If $\|f\|_{U^3} \geq \varepsilon$, then with probability at least $1 - \delta$ it finds a quadratic form q such that $\langle f, (-1)^q \rangle \geq \eta$.*
- *The probability that the algorithm outputs a quadratic form q with $\langle f, (-1)^q \rangle \leq \eta/2$ is at most δ .*

This is essentially an algorithmic version of the U^3 inverse theorem. The proof of Theorem 4.4 follows that of the inverse theorem very closely, except that many of the results from additive combinatorics that are used in the process need to be replaced by new “sampling versions”: since the subsets of \mathbb{F}_2^n that appear in the proof are generally very dense, it is too expensive to even write them down (let alone perform operations on them) if one is aiming for an algorithm that runs in time polynomial in n .

A crucial ingredient in the proof of Theorem 4.4 was an algorithmic version of the Bogolyubov-Ruzsa lemma (Lemma 5.3 in [TW11]), which reads as follows.

Lemma 4.5 (Algorithmic Bogolyubov-Ruzsa Lemma) *There exists a randomized algorithm **Bogolyubov** with parameters ρ and δ which, given oracle access to a function $h : \mathbb{F}_2^n \rightarrow \{0, 1\}$ with $\mathbb{E}h \geq \alpha$, outputs a subspace $V \leq \mathbb{F}_2^n$ of codimension at most $O(\alpha^{-3})$ (by giving a basis for V^\perp) such that with probability at least $1 - \delta$, we have $h * h * h * h(x) > \rho^4/2$ for all $x \in V$. The algorithm runs in time $n^2 \log n \cdot \text{poly}(1/\alpha, \log(1/\delta))$.*

In Section 4.2 we develop a replacement for this lemma (Theorem 4.8 below) with much better bounds. Inserting it into the framework of [TW11] reduces the dependence on ε in the running time and the number of terms in the decomposition to quasipolynomial, allowing us to state the following result.

Theorem 4.6 (Quasipolynomial Quadratic Goldreich-Levin) *Let $\varepsilon, \delta > 0$, $n \in \mathbb{N}$ and $B > 1$. Then there exists $\eta = \exp(\text{poly}(B, \log(1/\varepsilon)))$ and a randomized algorithm running in time $O(n^4 \log n \cdot \text{poly}(1/\eta, \log(1/\delta)))$ which, given any function $f : \mathbb{F}_2^n \rightarrow [-1, 1]$ as an oracle, outputs with probability at least $1 - \delta$ a decomposition into quadratic phases*

$$f = c_1(-1)^{q_1} + \dots + c_k(-1)^{q_k} + g + l$$

satisfying $k \leq 1/\eta^2$, $\|g\|_{U^3} \leq \varepsilon$, $\|l\|_1 \leq 1/2B$ and $|c_i| \leq \eta$ for all $i = 1, \dots, k$.

A further variant of Theorem 4.6 was proved in [TW11], in which the quadratic phases in the decomposition were replaced with slightly more complicated quadratic object, namely so-called *quadratic averages*, which were first introduced in [GW12] by Gowers and the last author. In this case the authors of [TW11] obtained a bound on the number of terms in the decomposition that was polynomial in ε^{-1} in time exponential in ε^{-1} , at the cost of the description size of each quadratic average being exponential in ε^{-1} . Inserting the Quasipolynomial Algorithmic Bogolyubov-Ruzsa Lemma (Theorem 4.8) in the work of Section 5 in [TW11], we obtain an algorithm which finds a decomposition into polynomially many quadratic averages in time quasipolynomial in ε^{-1} , where the description size of each average is now quasipolynomial in ε^{-1} . We leave the details to the interested reader.

4.1 The quasipolynomial Bogolyubov-Ruzsa lemma

In the context of \mathbb{F}_2^n , the traditional Bogolyubov-Ruzsa lemma states that if a set A has density at least α in its ambient group, then its fourfold sumset $A + A + A + A$ contains a subspace of codimension at most $2\alpha^{-2}$. It is easily proved using a few lines of Fourier analysis: the orthogonal complement of the subspace is given by the frequencies at which the indicator function of A has relatively large Fourier coefficients.

The bound on the codimension of V was improved to $O(\log^4(\alpha^{-1}))$ by Sanders [San10]. This improvement has far-reaching quantitative implications for other problems, in particular the bound in Roth's theorem [San11b] and the U^3 inverse theorem. We now deduce the quasipolynomial Bogolyubov-Ruzsa lemma from Corollary 3.4. In Section 4.2 we give an algorithmic version of the proof, which allows us to explicitly find a basis for V^\perp .

Theorem 4.7 (Quasipolynomial Bogolyubov-Ruzsa Lemma) *Let $A \subseteq \mathbb{F}_2^n$ be a subset of density α . Then there exists a subspace V of \mathbb{F}_2^n satisfying $V \subseteq 4A$ and*

$$\text{codim}(V) = O(\log^4(\alpha^{-1})).$$

Proof of Theorem 4.7: Applying Corollary 3.4 with $B = 2A$, $S = A$, $\ell = \log(30^2/\alpha)/2$, $\eta = 1/60$, $\varepsilon = 1/(120\ell)$ and t to be determined later on, we conclude the existence of a subspace V of $\text{codim}(V) \leq 32 \log(2/\alpha^t)$ which has the property that for all $v \in V$,

$$\mathbb{P}_{a \in A} [\rho_{A \rightarrow 2A}(a) \approx_{\varepsilon'} \rho_{A \rightarrow 2A}(a + v)] \geq 1 - 16 \frac{\ell |3A|}{\eta |A|} \cdot \exp(-2\varepsilon^2 t),$$

where $\varepsilon' = 4\varepsilon\ell + 2\eta + 2^{-\ell} \sqrt{|2A|/|A|} \leq 1/30 + 1/30 + (\sqrt{\alpha}/30) \cdot \sqrt{1/\alpha} \leq 1/10$.

Since $\rho_{A \rightarrow 2A}(a) = 1$ for all $a \in A$, this implies that

$$\begin{aligned}
\mathbb{P}_{a \in A} [\rho_{A \rightarrow 2A}(a + v) \geq 0.9] &= \mathbb{P}_{a \in A} [\rho_{A \rightarrow 2A}(a) \approx_{1/10} \rho_{A \rightarrow 2A}(a + v)] \\
&\geq 1 - 16 \frac{\ell}{\eta} \frac{|3A|}{|A|} \cdot \exp(-2\varepsilon^2 t) \\
&\geq 1 - 16 \frac{\ell}{\eta} \alpha^{-1} \cdot \exp(-2\varepsilon^2 t) \\
&= 1 - 480 \frac{\log(30^2/\alpha)}{\alpha} \cdot \exp\left(-\frac{t}{1800 \cdot \log^2(30^2/\alpha)}\right)
\end{aligned}$$

Choosing $t = O(\log^3(1/\alpha))$ we find that V has the desired codimension $O(\log^4(1/\alpha))$ and satisfies

$$\mathbb{P}_{a \in A} [\rho_{A \rightarrow 2A}(a + v) \geq 0.9] \geq 0.9$$

for all $v \in V$. Recalling the definition of $\rho_{A \rightarrow B}$ in (2), this inequality implies that for all $v \in V$,

$$\mathbb{P}_{a, a' \in A} [a + a' + v \in 2A] \geq 0.9^2 = 0.81.$$

By averaging, there therefore exists a pair $a, a' \in A$ such that $\mathbb{P}_{v \in V} [a + a' + v \in 2A] \geq 0.81$, or equivalently $|V \cap (a + a' + 2A)| \geq 0.81|V|$. But it is easy to see that if $|V \cap B| > \frac{1}{2}|V|$ for some subset $B \subseteq \mathbb{F}_2^n$, then $V \subseteq 2B$ (since every element $v \in V$ has precisely $|V|$ different representations as $v = v_1 + v_2$ where $v_1, v_2 \in V$). We conclude that $V \subseteq 2(a + a' + A + A) \subseteq 4A$, which finishes the proof. \blacksquare

4.2 The quasipolynomial Bogolyubov-Ruzsa lemma – algorithmic version

Here we develop an algorithmic version of the quasipolynomial Bogolyubov-Ruzsa lemma. In other words, we give an efficient (probabilistic) algorithm for finding a basis for the orthogonal complement of the subspace V in Theorem 4.7.

Theorem 4.8 (Algorithmic Quasipolynomial Bogolyubov-Ruzsa Lemma) *There exists a randomized algorithm Quasipolynomial-Bogolyubov with parameters α and γ' which, given oracle access to a function $h : \mathbb{F}_2^n \rightarrow \{0, 1\}$ with $\mathbb{E}h \geq \alpha$, outputs a subspace $V \leq \mathbb{F}_2^n$ of codimension at most $O(\log^4(1/\alpha))$ (by giving a basis for V^\perp) such that with probability at least $1 - \gamma'$, we have $h * h * h * h(x) > 0$ for each $x \in V$. The algorithm runs in time $2^{O(\log^4(1/\alpha))} \cdot \text{polylog}(1/\gamma') \cdot n^3 \log n$.*

Note that if the function h equals the indicator function of a subset $A \subseteq \mathbb{F}_2^n$, then the condition $h * h * h * h(x) > 0$ implies that $x \in 4A$, and if this condition is satisfied for all $x \in V$, then $V \subseteq 4A$. However, while it will be convenient to think of the set $A = \{x \in \mathbb{F}_2^n \mid h(x) = 1\}$ in the proof, we will actually apply the theorem to the output of a randomized algorithm, for which the statement in terms of a function makes more sense. We also assume for convenience that $\mathbb{E}h$ is *exactly* α . The proof remains unchanged when the density is larger than α .

In the combinatorial proof we considered the measure $\rho_{A \rightarrow 2A}(y) = \mathbb{P}_{a \in A} [y + a \in 2A]$, and the subspace V was defined in terms of a set X which was described using this measure. However, now this measure is difficult to compute since it might not be possible to test membership simply using oracle access to h , which is the indicating function for A . We give a robust version of the

combinatorial proof by noting that $y + a \in 2A$ is equivalent to saying that $h * h(a + y) > 0$. But since we do not have noise-free access to $h * h$, we cannot test this function directly. Instead, we test if $h * h(a + y) \geq \eta\alpha^2$ for some $\eta > 0$. For this purpose, we define the set

$$Z_\eta := \{x \in \mathbb{F}_2^n \mid h * h(x) \geq \eta \cdot \alpha^2\}.$$

The following procedure tests membership in Z_η by estimating $h * h$ using few samples.

Z-Test (x)

- Estimate the expectation $h * h(x) = \mathbb{E}_{y \in \mathbb{F}_2^n} h(y) \cdot h(x - y)$ using r samples of elements $y \in \mathbb{F}_2^n$.
- Answer 1 if the estimate is at least $\eta\alpha^2$ and 0 otherwise.

However, since we are *estimating* the value of $h * h$, we only have the following kind of guarantee.

Claim 4.9 *Given $\gamma_1 > 0$, the output of **Z-Test**(x) with $r = O(1/(\eta^2\alpha^4) \cdot \log(1/\gamma_1))$ queries satisfies the following guarantee with probability at least $1 - \gamma_1$.*

- **Z-Test**(x) = 1 $\implies x \in Z_{\eta/2}$.
- **Z-Test**(x) = 0 $\implies x \notin Z_{3\eta/2}$.

Proof: This follows immediately from the Hoeffding bound (Lemma 2.1). ■

Let $Z(x)$ denote the (random) function given by the output of **Z-Test**(x). Then the measure

$$\rho(y) := \mathbb{E}_{a \in A} [Z(y + a)] = \frac{1}{\alpha} \cdot h * Z(y).$$

can be efficiently estimated by sampling.

Next we need a procedure to test for membership in the set X which satisfies the iterated almost-periodicity condition in Corollary 3.3.

Lemma 4.10 *Let $\eta \in (0, 1/3600)$, $\gamma_2 > 0$ and let h and A be defined as above. Then for any integers $\ell, t = O(\ell^2(\log \ell + \log(1/\alpha)))$ there exists a randomized procedure **X-Test** with outputs in $\{0, 1\}$ which runs in time $(1/\alpha)^{O(t)}$ and has the following properties.*

- With probability $1 - \gamma_2$, $\mathbb{P}_{x \in \mathbb{F}_2^n} [\mathbf{X-Test}(x) = 1] \geq \alpha^{2t}/4$.
- For all $x_1, \dots, x_\ell \in \mathbb{F}_2^n$, we have with probability at least $1 - \gamma_2$,

$$\forall i \in [\ell] \mathbf{X-Test}(x_i) = 1 \implies \mathbb{E}_{a, a' \in A} [Z(a + a' + x_1 + \dots + x_\ell)] \geq \frac{9}{10}.$$

Proof: As in the proof of Proposition 3.2, define $\mathbf{G}[\varepsilon, \delta]$ to be the set of sequences $\mathbf{a} \in (\mathbb{F}_2^n)^t$ which can be used to estimate ρ well (for our new definition of ρ). For $\mathbf{a} \in A^t$, define

$$\hat{\rho}_{\mathbf{a}}(y) := \frac{1}{t} \cdot \sum_{i=1}^t Z(a_i + y),$$

$$\mathbf{G}[\varepsilon, \delta] := \left\{ \mathbf{a} \in A^t \mid \mathbb{P}_{y \in \mathbb{F}_2^n} [|\rho(y) - \hat{\rho}_{\mathbf{a}}(y)| \geq \varepsilon] \leq \delta \right\}.$$

Also, as above, we will need to test membership in $\mathbf{G}[\varepsilon, \delta]$. This we will only be able to do approximately, using the following randomized procedure.

G-Test ($\mathbf{a} = (a_1, \dots, a_t)$)

- Check if $h(a_1) = \dots = h(a_t) = 1$. If not output 0.
- Pick r independent samples $y_1, \dots, y_r \in \mathbb{F}_2^n$.
- For each y_i , estimate $\rho(y_i)$ using r' independent samples. Also compute $\hat{\rho}_{\mathbf{a}}(y_i)$ for each y_i .
- If $|\{y_i \mid |\rho(y_i) - \hat{\rho}_{\mathbf{a}}(y_i)| \geq \varepsilon\}| > \delta r$ then output 0, else output 1.

We prove the following guarantee for the above test.

Claim 4.11 *Given $\gamma_3 > 0$, the output of **G-Test**(\mathbf{a}) with $r = O(1/\delta^2 \cdot \log(1/\gamma_3))$ and $r' = O(1/\varepsilon^2 \cdot \log(r/\gamma_3))$ queries, satisfies the following guarantee with probability at least $1 - \gamma_3$.*

- **G-Test**(\mathbf{a}) = 1 $\implies \mathbf{a} \in \mathbf{G}[2\varepsilon, 2\delta]$.
- **G-Test**(\mathbf{a}) = 0 $\implies \mathbf{a} \notin \mathbf{G}[\varepsilon/2, \delta/2]$.

Proof: Again, this is a direct consequence of the Hoeffding bound (Lemma 2.1). ■

Note that the definition of the procedure **G-Test** actually depends on the parameters ε, δ and the error parameter γ_3 , for choosing the appropriate values of r and r' . However, we choose to hide this dependence for the sake of readability.

From now on let $G(\mathbf{a})$ denote the output of **G-Test** on the input \mathbf{a} . We will now find an element $\hat{\mathbf{a}} \in (\mathbb{F}_2^n)^t$ such that $G(\hat{\mathbf{a}}) = 1$ and $G(\hat{\mathbf{a}} + x) = 1$ for a large number of elements $x \in \mathbb{F}_2^n$. This can be done efficiently with high probability if $\delta = \exp(-O(\varepsilon^2 t))$.

Claim 4.12 *Given $\gamma_4 > 0$, there exists an algorithm which makes $O((1/\alpha^{6t}) \cdot \log^2(1/\gamma_4))$ calls to **G-Test** and finds an $\hat{\mathbf{a}} \in (\mathbb{F}_2^n)^t$ such that with probability $1 - \gamma_4$, we have $G(\hat{\mathbf{a}}) = 1$ and $\mathbb{E}_{x \in \mathbb{F}_2^n} [G(\hat{\mathbf{a}} + x)] \geq \alpha^{2t}/4$.*

Proof: The Hoeffding bound gives that $|\mathbf{G}[\varepsilon/2, \delta/2]| \geq 0.99 |A^t|$ for $t \geq (c/\varepsilon^2) \cdot \log(1/\delta)$. Since A has density α in \mathbb{F}_2^n we have for $\gamma_3 < 0.04$ that $\mathbb{E}_{\mathbf{a} \in (\mathbb{F}_2^n)^t} [G(\mathbf{a})] \geq (1 - \gamma_3) \cdot (0.99\alpha^t) \geq 0.95\alpha^t$. Using convexity gives

$$\begin{aligned} \mathbb{E}_{\mathbf{a} \in (\mathbb{F}_2^n)^t, x \in \mathbb{F}_2^n} [G(\mathbf{a}) \cdot G(\mathbf{a} + x)] &= \mathbb{E}_{\mathbf{a} \in (\mathbb{F}_2^n)^t, x, x' \in \mathbb{F}_2^n} [G(\mathbf{a} + x) \cdot G(\mathbf{a} + x')] \\ &= \mathbb{E}_{\mathbf{a} \in (\mathbb{F}_2^n)^t} \left[\left(\mathbb{E}_{x \in \mathbb{F}_2^n} [G(\mathbf{a} + x)] \right)^2 \right] \\ &\geq \left(\mathbb{E}_{\mathbf{a} \in (\mathbb{F}_2^n)^t, x \in \mathbb{F}_2^n} [G(\mathbf{a} + x)] \right)^2 \\ &\geq (0.95 \cdot \alpha^t)^2 \geq 0.9 \cdot \alpha^{2t}. \end{aligned}$$

Hence, by Markov's inequality

$$\mathbb{P}_{\mathbf{a} \in (\mathbb{F}_2^n)^t} \left[G(\mathbf{a}) \cdot \mathbb{E}_{x \in \mathbb{F}_2^n} [G(\mathbf{a} + x)] \geq \alpha^{2t}/2 \right] \geq \alpha^{2t}/4.$$

The algorithm then simply tries random sequences \mathbf{a} until it finds one for which $G(\mathbf{a}) = 1$. For such an \mathbf{a} , it estimates $\mathbb{E}_{x \in \mathbb{F}_2^n} [G(\mathbf{a} + x)]$ using $O((1/\alpha^{4t}) \cdot \log(1/\gamma_4))$ samples. With probability $1 - \gamma_4/2$, the estimate is accurate to within an additive $\alpha^{2t}/8$. The algorithm stops and outputs an $\hat{\mathbf{a}}$ for which $G(\hat{\mathbf{a}}) = 1$ and the estimate computed by the algorithm is at least $3\alpha^{2t}/4$. By the above, it finds such an $\hat{\mathbf{a}}$ with probability at least $1 - \gamma_4/2$ in at most $O((1/\alpha^{2t}) \cdot \log(1/\gamma_4))$ attempts. If not, it simply outputs a random $\hat{\mathbf{a}}$. ■

Given $\hat{\mathbf{a}}$ as above, we define X to be the set

$$X := \{x \in \mathbb{F}_2^n \mid G(\hat{\mathbf{a}} + x) = 1\}.$$

Note that $G(\hat{\mathbf{a}}) = 1$ and $|X| \geq (\alpha^{2t}/4) \cdot 2^n$ with probability $1 - \gamma_4$. Also, membership in X can be tested efficiently. We simply define the procedure **X-Test** as

$$\mathbf{X-Test}(x) = \mathbf{G-Test}(\hat{\mathbf{a}} + x).$$

We now prove that this X suffices for our purposes. We will prove using induction that for x_1, \dots, x_ℓ satisfying $\mathbf{X-Test}(x_i) = 1 \forall i \in [\ell]$, we have with probability $1 - \gamma_3$ that

$$\mathbb{E}_{a, a' \in A} [Z(a + a' + x_1 + \dots + x_\ell)] = \mathbb{E}_{a \in A} [\rho(a + x_1 + \dots + x_\ell)] \geq 9/10.$$

By Claims 4.11 and 4.12 we have that $\hat{\mathbf{a}}, \hat{\mathbf{a}} + x_1, \dots, \hat{\mathbf{a}} + x_\ell \in \mathbf{G}[2\varepsilon, 2\delta]$ with probability at least $1 - (\ell + 1)\gamma_3 - \gamma_4$. We will prove that whenever $\hat{\mathbf{a}}, x_1, \dots, x_\ell$ satisfy this condition, then for all $r \in \{0, \dots, \ell\}$ we have

$$\mathbb{P}_{a \in A} [\rho(a + x_1 + \dots + x_r) \geq 1 - \sqrt{\eta + \gamma_1/\alpha^2} - 4r\varepsilon] \geq 1 - \sqrt{\eta + \gamma_1/\alpha^2} - 4\delta r/\alpha,$$

where γ_1 is the error parameter in Claim 4.9. The following claim proves the base case $r = 0$.

Claim 4.13 $\mathbb{P}_{a \in A} [\rho(a) \geq 1 - \sqrt{\eta + (\gamma_1/\alpha^2)}] \geq 1 - \sqrt{\eta + (\gamma_1/\alpha^2)}$.

Proof: We have

$$\begin{aligned} \mathbb{E}_{a \in A} [\rho(a)] &= \mathbb{E}_{a, a' \in A} [Z(a + a')] = \frac{1}{\alpha^2} \cdot \langle h * h, Z \rangle \\ &= \frac{1}{\alpha^2} \cdot \langle h * h, 1 \rangle - \frac{1}{\alpha^2} \cdot \langle h * h, (1 - Z) \rangle \\ &= 1 - \frac{1}{\alpha^2} \cdot \langle h * h, (1 - Z) \rangle \end{aligned}$$

Since with probability at least $1 - \gamma_1$, $h * h$ is at most $\eta\alpha^2$ when $1 - Z = 1$, the inner product in the second term is at most $\eta\alpha^2 + \gamma_1$. This gives $\mathbb{E}_{a \in A} [\rho(a)] \geq 1 - \eta - (\gamma_1/\alpha^2)$. An averaging argument then proves the claim. ■

For readability, let β denote the quantity $\sqrt{\eta + (\gamma_1/\alpha^2)}$. We assume by induction that for x_1, \dots, x_r

$$\mathbb{P}_{a \in A} [\rho(a + x_1 + \dots + x_r) \geq 1 - \beta - 4r\varepsilon] \geq 1 - \beta - 4r\delta/\alpha.$$

Define the set

$$A_r = \{y = a + x_1 + \dots + x_r \mid a \in A, \rho(y) \geq 1 - \beta - 4r\varepsilon\}.$$

By the above, the density of A_r is at least $\alpha_r = \alpha \cdot (1 - \beta - 4r\delta/\alpha)$. For any $x_{r+1} \in X$, we have that for at least a $(1 - 4\delta/\alpha_r)$ -fraction of elements $y \in A_r$,

$$|\rho(y) - \hat{\rho}_{\mathbf{a}}(y)| \leq \varepsilon \quad \text{and} \quad |\rho(y + x_{r+1}) - \hat{\rho}_{\mathbf{a}+x_{r+1}}(y + x_{r+1})| = |\rho(y + x_{r+1}) - \hat{\rho}_{\mathbf{a}}(y)| \leq \varepsilon.$$

Thus, by the triangle inequality

$$\mathbb{P}_{a \in A} [\rho(a + x_1 + \dots + x_{r+1}) \geq 1 - \beta - 4(r+1)\varepsilon] \geq (1 - 4\delta/\alpha_r) \cdot (\alpha_r/\alpha) = 1 - \beta - 4(r+1)\delta/\alpha.$$

To get the required bounds we choose $\gamma_1 = \eta\alpha^2$, $\gamma_3 = \gamma_2/(2(\ell+1))$ and $\gamma_4 = \gamma_2/2$. Also, we take $\varepsilon = 1/(120\ell)$, $\delta = \alpha/(120\ell)$ and $\eta = 10^{-4}$.

The basic procedure used for the above algorithm is **Z-Test**, for which the running time is dominated by $O((1/\alpha^4) \log(1/\alpha))$ oracle queries to the function h . Assuming the query can be answered in constant time and it takes $O(n)$ time to write down the input, the running time for **Z-Test** is $O((1/\alpha^4) \log(1/\alpha) \cdot n)$. Also, for the above choice of parameters, the procedure **G-Test** makes $O((\ell^4/\alpha^2) \cdot \log(\ell/\alpha) \cdot \log^3(\ell/\gamma_2))$ calls to **Z-Test**. Finally, the procedure in Claim 4.12 makes $O((1/\alpha^{6t}) \cdot \log^2(\gamma_2))$ calls to **G-Test**. Taking t to be $O((1/\varepsilon^2) \log(1/\delta)) = O(\ell^2 \cdot \log(\ell/\alpha))$, this gives a running time of $(1/\alpha)^{O(\ell^2(\log \ell + \log(1/\alpha)))} \cdot \log^5(1/\gamma_2) \cdot n$ for the algorithm. ■

We can now prove following algorithmic analogue of Theorem 4.7.

Proof of Theorem 4.8: Choose $\ell = \log(10/\alpha)$. Let X be the set defined in Lemma 4.10. Define the subspace V_0 as

$$V_0 := \left\{ \zeta \in \mathbb{F}_2^n \mid \left| \widehat{\mathbb{1}_X}(\zeta) \right| \geq \widehat{\mathbb{1}_X}(0)/2 \right\}^\perp = (\text{Spec}_{1/2}(X))^\perp$$

where $\mathbb{1}_X$ is the indicator function of X . To find (an approximation to) V_0 , we first estimate $\widehat{\mathbb{1}_X}(0) = \mathbb{E}[\mathbb{1}_X] = \mathbb{P}_{x \in \mathbb{F}_2^n} [\mathbf{X}\text{-Test}(x) = 1]$ using $\Omega((1/\alpha^{4t}) \cdot \log(1/\gamma'))$ samples so that with probability $1 - \gamma'/4$, the error is at most $\alpha^{2t}/8$. By Lemma 4.10, with probability $1 - \gamma_2$, the quantity $\mathbb{P}_{x \in \mathbb{F}_2^n} [\mathbf{X}\text{-Test}(x) = 1]$ is at least $\alpha^{2t}/4$. Taking $\gamma_2 = \gamma'/4$, we get that with probability $1 - \gamma'/2$, the estimate is at least $3\alpha^{2t}/8$. Call this estimate μ_0 .

We now need a procedure which determines the large Fourier coefficients of $\mathbb{1}_X$ with reasonable accuracy. This procedure is given by the Goldreich-Levin theorem (Theorem 4.1).

We run Theorem 4.1 with error parameter $\delta = \gamma'/2$ and an oracle access to the procedure **X-Test**, to find all characters with Fourier coefficients larger than $\mu_0/8$ in absolute value, up to an additive accuracy of $\nu = \mu_0/16$. Let K be the list of characters given by the algorithm. We take

$$V = \{ \zeta \in \mathbb{F}_2^n \mid \zeta \in K \}^\perp.$$

Now with probability at least $1 - \gamma'$, the trivial coefficient $\widehat{\mathbb{1}_X}(0)$ is at least $\alpha^{2t}/4$, K contains all ζ such that $|\widehat{\mathbb{1}_X}(\zeta)| \geq \widehat{\mathbb{1}_X}(0)/2$ and $|\widehat{\mathbb{1}_X}(\zeta)| \geq \widehat{\mathbb{1}_X}(0)/32$ for all $\zeta \in K$. By Chang's theorem (Theorem 2.2) and our choice of parameters, the codimension of V is then at most $O(\log(1/\alpha^{2t})) = O(\log^4(1/\alpha))$. It remains to show that $h * h * h * h(x) > 0$ for all $x \in V$.

Let $g = \mathbb{1}_{Z_\eta} * \mu_A * \mu_X * \dots * \mu_X$, where μ_X is convolved ℓ times. By definition of X , we have that

$$\langle g, \mu_A \rangle \geq 9/10.$$

Also, by definition of V , we have

$$|\langle g, \mu_A \rangle - \langle g * \mu_V, \mu_A \rangle| \leq \sum_{\gamma \notin V^\perp} \left| \widehat{\mathbf{1}}_{Z_\eta}(\gamma) \right| |\widehat{\mu}_A(\gamma)|^2 |\widehat{\mu}_X(\gamma)|^\ell \leq \frac{1}{10}.$$

Hence, $\langle g * \mu_V, \mu_A \rangle \geq 4/5$. Expanding this, we get

$$\mathbb{E}_{a \in A} \mathbb{E}_{v \in V} \mathbb{E}_{a' \in A} \mathbb{E}_{x_1, \dots, x_\ell \in X} [\mathbf{1}_{Z_\eta}(a + a' + v + x_1 + \dots + x_\ell)] \geq 4/5.$$

Thus, there exists an $x_0 \in \mathbb{F}_2^n$ such that $\mathbb{E}_{v \in V} [\mathbf{1}_{Z_\eta}(v + x_0)] = \mathbb{E}_{v \in V} [\mathbf{1}_{Z_\eta + x_0}(v)] \geq 4/5$. In other words, there exists $x_0 \in \mathbb{F}_2^n$ such that $|V \cap (x_0 + Z_\eta)| > 4|V|/5 > |V|/2$. Since any $v \in V$ has $|V|$ representations as $v = v_1 + v_2$ with $v_1, v_2 \in V$, we find that for any $v \in V$, there exist $v_1, v_2 \in x_0 + Z_\eta$ such that $v = v_1 + v_2$, and hence $v \in 2Z_\eta$.

The running time is dominated by the $O(n^2 \log n \cdot \text{poly}((1/\alpha)^t, \log(1/\gamma')))$ calls made by the Goldreich-Levin algorithm to the procedure **X-Test**. For $t = O(\ell^2 \log(\ell/\alpha))$ and $\ell = \log(10/\alpha)$, this implies a running time of $2^{O(\log^4(1/\alpha))} \cdot \text{polylog}(1/\gamma') \cdot n^3 \log n$. \blacksquare

5 Sumsets of dense sets contain large subspaces

Inspired by the question of whether dense subsets of $\{1, \dots, N\}$ contain long arithmetic progressions, which has received extensive coverage in the literature [Bou90, Gre02, San08], Ben Green asked an analogous question in the finite field setting and obtained the following result [Gre05].

Theorem 5.1 (Green's theorem on subspaces in sumsets) *Let $A \subseteq \mathbb{F}_2^n$ be a subset of density α . Then $A + A$ contains a subspace V of \mathbb{F}_2^n of dimension*

$$\dim(V) = \Omega(\alpha^2 n).$$

In [San11a] Sanders showed, using a Fourier-based density-increment strategy, that one can in fact take the subspace V to have dimension $\dim(V) = \Omega(\alpha n)$. He remarks that a bound of similar strength could be obtained via a finite field analogue of the methods of Croot, Laba and Sisask [CLS11]. Our main theorem in this section is the following, replicating the result from [CLS11], which falls slightly short of Sanders's bound [San11a].

Theorem 5.2 (Sumsets of dense sets contain large subspaces) *Let $A \subseteq \mathbb{F}_2^n$ be a subset of density α . Then $A + A$ contains an affine subspace V of \mathbb{F}_2^n of dimension*

$$\dim(V) = \Omega\left(\frac{\alpha}{\log^3(1/\alpha)} n\right).$$

For the proof of the above theorem we shall need a refined version of the almost-periodicity results from Section 3. In particular, we shall need the following refined version of Corollary 3.4.

Corollary 5.3 (Refined almost-periodicity of sumsets over a subspace) *Let $A \subset \mathbb{F}_2^n$ be a subset of density α . Then for every integer t and set $B \subseteq \mathbb{F}_2^n$, there exists a subspace V of codimension $\text{codim}(V) \leq 32 \log(2/\alpha^t)$ with the following property.*

For every $v \in V$, for all subsets $S \subseteq \mathbb{F}_2^n$ and for every $\eta, \varepsilon > 0$ and integer ℓ ,

$$\mathbb{P}_{y \in S} [\rho_{A \rightarrow B}(y) - \rho_{A \rightarrow B}(y + v) \leq \varepsilon'] \geq 1 - 16 \frac{\ell |A + B|}{\eta |S|} \cdot \exp(-\varepsilon^2 t / 4), \quad (20)$$

where $\varepsilon' = 4\varepsilon\ell\sqrt{\rho_{A \rightarrow B}(y)} + 2\eta + 2^{-\ell}\sqrt{|B|/|A|}$.

The main difference between the above corollary and Corollary 3.4 lies in the term $\sqrt{\rho_{A \rightarrow B}(y)}$ which appears in the expression for ε' in the above corollary. This term makes ε' smaller which in turn makes the above corollary stronger. For the sake of simplicity, we only consider in the above corollary one-sided bounds of the form $\rho_{A \rightarrow B}(y) - \rho_{A \rightarrow B}(y + v) \leq \varepsilon'$ instead of two-sided bounds of the form $\rho_{A \rightarrow B}(y) \approx_{\varepsilon'} \rho_{A \rightarrow B}(y + v)$. This will suffice for the proof of Theorem 5.2.

The proof of Corollary 5.3 is similar to the proof of Corollary 3.4, and the main difference is that in the proof of Corollary 5.3 we perform a more detailed analysis of the distribution $\mathbb{1}_B(a + y)$ when a is distributed uniformly over A and y is a fixed point in \mathbb{F}_2^n and use information on the variance of this distribution. More specifically, in the proof Corollary 5.3, instead of using the regular Hoeffding bound for sampling (Lemma 2.1), we use the following well-known refinement involving the variance [TV06].

Lemma 5.4 (Refined Hoeffding bound for sampling) *If \mathbf{X} is a random variable satisfying $|\mathbf{X} - \mathbb{E}[\mathbf{X}]| \leq 1$ and $\hat{\mu}$ is the empirical average obtained from t samples, then*

$$\mathbb{P}[|\mathbb{E}[\mathbf{X}] - \hat{\mu}| > \gamma] \leq 2 \exp\left(-\frac{\gamma^2 t}{4\sigma^2(X)}\right)$$

provided that $\gamma < 2\sigma^2$.

Since the proof of Corollary 5.3 presents some technical complications, we include it in full in Appendix A. The rest of this section is devoted to the proof of Theorem 5.2 assuming that Corollary 5.3 is true.

The idea of the proof of Theorem 5.2 is as follows. Applying Corollary 5.3 with $B = A$ implies the existence of a relatively large subspace V such that for every $v \in V$, for a large fraction of $y \in \mathbb{F}_2^n$, it holds that $\rho_{A \rightarrow A}(y + v) > 0$. Our goal will be to show that an affine shift of V is contained in $2A$, or equivalently to show the existence of an affine shift $y \in \mathbb{F}_2^n$ such that $\rho_{A \rightarrow A}(y + v) > 0$ for all $v \in V$. Suppose that we have chosen the parameters in Corollary 5.3 in such a way that for every $v \in V$, at least $(1 - \delta)$ -fraction of $y \in \mathbb{F}_2^n$ satisfy that $\rho_{A \rightarrow A}(y + v) > 0$. Then the union bound implies that at least a $(1 - |V|\delta)$ -fraction of $y \in \mathbb{F}_2^n$ satisfy the condition $\rho_{A \rightarrow A}(y + v) > 0$ for all $v \in V$. Thus in order to guarantee the existence of the desired affine shift y , it suffices to choose the parameters in Corollary 5.3 in such a way that $|V|\delta < 1$.

Note that we wouldn't have gained anything from considering the variance in the proof of the quasipolynomial Bogolyubov-Ruzsa lemma (Theorem 4.7) since there Corollary 3.4 is applied to elements y for which $\rho_{A \rightarrow B}(y)$ is very large (between 0.9 and 1), and we have no better handle on the variance. In contrast, here the typical element to which we apply Corollary 5.3 satisfies $\rho_{A \rightarrow A}(y) = \alpha$, so that the variance is small as well.

For the proof of Theorem 5.2 we shall need the following simple lemma.

Lemma 5.5 *Let $f(t) = t^2 - bt - c$ for $b > 0$, $c \geq 0$, and suppose that $0 \leq t' \leq t''$ are such that $f(t') > 0$. Then $f(t'') \leq f(t')$.*

Proof: The fact that $b > 0$, $c \geq 0$ implies that $f(t)$ has a root $t_1 \leq 0$ and another root $t_2 > 0$. Thus we have that $f(t)$ is negative in the interval $(0, t_2)$ and is positive in the interval (t_2, ∞) . The fact that $t' \geq 0$ and $f(t') > 0$ thus implies that $t' > t_2$. The lemma follows by noting that f is monotonically increasing in the interval (t_2, ∞) . ■

Proof of Theorem 5.2: Apply Corollary 5.3 with $B = A$, $\eta = \alpha/24$, $\ell = \log(12/\alpha)$, $\varepsilon = \sqrt{2\alpha}/(48\ell)$, t to be determined later on and

$$S = \{y \in \mathbb{F}_2^n \mid \rho_{A \rightarrow A}(y) \geq \alpha/2\}.$$

Noting that

$$\mathbb{E}_{y \in \mathbb{F}_2^n} [\rho_{A \rightarrow A}(y)] = \mathbb{P}_{y \in \mathbb{F}_2^n, a \in A} [a + y \in A] = \mathbb{E}_{a \in A} \left[\mathbb{P}_{y \in \mathbb{F}_2^n} [a + y \in A] \right] = \alpha,$$

Markov's inequality implies that $|S| \geq (\alpha/2) \cdot 2^n$.

With this choice of parameters Corollary 5.3 implies that for every $v \in V$,

$$\mathbb{P}_{y \in S} \left[\rho_{A \rightarrow A}(y + v) \geq \rho_{A \rightarrow A}(y) - \alpha/6 - \frac{\sqrt{2\alpha \cdot \rho_{A \rightarrow A}(y)}}{12} \right] \geq 1 - 16 \frac{\ell}{\eta} \cdot \frac{|2A|}{|S|} \cdot \exp(-\varepsilon^2 t/4)$$

Let $\delta := 16(\ell/\eta) \cdot (|2A|/|S|) \cdot \exp(-\varepsilon^2 t/4)$. Since $\rho_{A \rightarrow A}(y) \geq \alpha/2$ for every $y \in S$, the inequality above implies that

$$\begin{aligned} \mathbb{P}_{y \in S} [\rho_{A \rightarrow A}(y + v) \geq \alpha/4] &= \mathbb{P}_{y \in S} \left[\rho_{A \rightarrow A}(y + v) \geq \alpha/2 - \alpha/6 - \frac{\sqrt{2\alpha \cdot \alpha/2}}{12} \right] \\ &\geq \mathbb{P}_{y \in S} \left[\rho_{A \rightarrow A}(y + v) \geq \rho_{A \rightarrow A}(y) - \alpha/6 - \frac{\sqrt{2\alpha \cdot \rho_{A \rightarrow A}(y)}}{12} \right] \\ &\geq 1 - \delta \end{aligned}$$

where the first inequality follows by applying Lemma 5.5 with $f(t) = t^2 - (\sqrt{2\alpha}/12)t - \alpha/6$, $t' = \sqrt{\alpha/2}$, $t'' = \sqrt{\rho_{A \rightarrow A}(y)}$, and noting that our assumptions imply that $0 \leq t' \leq t''$ and that $f(t') = \alpha/4 > 0$.

A union bound then implies that

$$\mathbb{P}_{y \in S} [\rho_{A \rightarrow A}(y + v) \geq \alpha/4 \forall v \in V] \geq 1 - |V| \cdot \delta.$$

To conclude the proof we shall show that for sufficiently small integer t one can guarantee that $|V|\delta < 1$. This in turn will imply the existence of an affine shift $y \in S$ such that $\rho_{A \rightarrow A}(y + v) > 0$ for every $v \in V$, and consequently $y + V \subseteq 2A$. Our choice of parameters implies that

$$\begin{aligned} |V|\delta &= \left(\frac{\alpha^t}{2} \right)^{32} \cdot 2^n \cdot 16 \cdot \frac{\ell}{\eta} \cdot \frac{|2A|}{|S|} \cdot \exp(-\varepsilon^2 t/4) \\ &\leq \exp \left(-t \left(32 \log(1/\alpha) + \frac{2\alpha}{4 \cdot 48^2 \cdot \log^2(12/\alpha)} \right) + \left(2 \log(1/\alpha) + n + \log \log(12/\alpha) \right) \right) \end{aligned}$$

Thus, $|V|\delta < 1$ is guaranteed by letting

$$t = \frac{2 \log(1/\alpha) + n + \log \log(12/\alpha)}{32 \log(1/\alpha) + \frac{2\alpha}{4 \cdot 48^2 \cdot \log^2(12/\alpha)}} = \frac{n + O(\log(1/\alpha))}{32 \log(1/\alpha) + \Omega(\alpha/\log^2(1/\alpha))}.$$

But for such a choice of t we have that

$$\dim(V) = n - 32 \log(1/\alpha)t - 32 = \Omega\left(\frac{\alpha}{\log^3(1/\alpha)}n\right).$$

■

References

- [Bou90] Jean Bourgain, *On arithmetic progressions in sums of sets of integers*, A tribute to Paul Erdős, Cambridge Univ. Press, Cambridge, 1990, pp. 105–109.
- [Cha02] Mei-Chu Chang, *A polynomial bound in Freiman’s theorem*, Duke Math. J. **113** (2002), no. 3, 399–419.
- [CLS11] Ernie Croot, Izabella Łaba, and Olof Sisask, *Arithmetic progressions in sumsets and L^p -almost-periodicity*, <http://arxiv.org/abs/1103.6000v1> (2011).
- [CS10] Ernie Croot and Olof Sisask, *A probabilistic technique for finding almost-periods of convolutions*, Geom. Funct. Anal. **20** (2010), no. 6, 1367–1396.
- [FK99] A. M. Frieze and R. Kannan, *Quick approximation to matrices and applications*, Combinatorica **19** (1999), no. 2, 175–220.
- [GL89] O. Goldreich and L. Levin, *A hard-core predicate for all one-way functions*, Proceedings of the 21st ACM Symposium on Theory of Computing, 1989, pp. 25–32.
- [Gow98] Timothy Gowers, *A new proof of Szemerédi’s theorem for arithmetic progressions of length four*, Geom. Funct. Anal. **8** (1998), no. 3, 529–551.
- [Gre02] Ben Green, *Arithmetic progressions in sumsets*, Geom. Funct. Anal. **12** (2002), no. 3, 584–597.
- [Gre05] ———, *Finite field models in additive combinatorics*, Surveys in combinatorics 2005, London Math. Soc. Lecture Note Ser., vol. 327, Cambridge Univ. Press, Cambridge, 2005, pp. 1–27.
- [Gre07] ———, *Montréal notes on quadratic Fourier analysis*, Additive combinatorics, CRM Proc. Lecture Notes, vol. 43, Amer. Math. Soc., Providence, RI, 2007, pp. 69–102.
- [GT08] Ben Green and Terence Tao, *An inverse theorem for the Gowers $U^3(G)$ norm*, Proc. Edinb. Math. Soc. (2) **51** (2008), no. 1, 73–153.
- [GW12] Timothy Gowers and Julia Wolf, *Linear forms and quadratic uniformity for functions on \mathbb{F}_p^n* , Mathematika **57** (2012), no. 2, 215–237.
- [IMR12] Russell Impagliazzo, Cristopher Moore, and Alexander Russell, *An Entropic Proof of Chang’s Inequality*, <http://arxiv.org/abs/1205.0263v1> (2012).
- [Lov12] Shachar Lovett, *An exposition of Sanders’s quasi-polynomial Freiman-Ruzsa theorem*, Electronic Colloquium on Computational Complexity (ECCC) **19** (2012), 29.

- [Ruz99] Imre Ruzsa, *An analog of Freiman's theorem in groups*, Astérisque (1999), no. 258, xv, 323–326, Structure theory of set addition.
- [Sam07] Alex Samorodnitsky, *Low-degree tests at large distances*, Proceedings of the 39th ACM Symposium on Theory of Computing, 2007, pp. 506–515.
- [San08] Tom Sanders, *Additive structures in sumsets*, Math. Proc. Cambridge Philos. Soc. **144** (2008), no. 2, 289–316.
- [San10] ———, *On the Bogolyubov-Ruzsa lemma*, To appear, Anal. PDE (2010).
- [San11a] ———, *Green's sumset problem at density one half*, Acta Arith. **146** (2011), no. 1, 91–101.
- [San11b] ———, *On Roth's theorem on progressions*, Ann. of Math. (2) **174** (2011), no. 1, 619–636.
- [San12] ———, *Lecture notes on applications of commutative harmonic analysis*, <http://people.maths.ox.ac.uk/~sanders> (2012).
- [ST06] Alex Samorodnitsky and Luca Trevisan, *Gowers uniformity, influence of variables, and PCPs*, STOC, 2006, pp. 11–20.
- [TTV09] Luca Trevisan, Madhur Tulsiani, and Salil Vadhan, *Boosting, regularity and efficiently simulating every high-entropy distribution*, Proceedings of the 24th IEEE Conference on Computational Complexity, 2009.
- [TV06] Terence Tao and Van Vu, *Additive combinatorics*, Cambridge University Press, 2006.
- [TW11] Madhur Tulsiani and Julia Wolf, *Quadratic Goldreich-Levin theorems*, FOCS, 2011, pp. 619–628.

A Appendix: Proof of Corollary 5.3

In order to prove Corollary 5.3, we start with refined versions of Proposition 3.2 and Corollary 3.3, given as Proposition A.1 and Corollary A.2 below.

Proposition A.1 (Refined version of almost-periodicity of sumsets) *Let $A \subset \mathbb{F}_2^n$ be a subset satisfying $|2A| \leq K|A|$. Then for every integer t and set $B \subseteq \mathbb{F}_2^n$ there exists a set X with the following properties.*

1. *The set X is contained in an affine shift of A .*
2. *The size of X is at least $|A|/(2K^{t-1})$.*
3. *For all $x \in X$ and for all subsets $S \subseteq \mathbb{F}_2^n$,*

$$\mathbb{P}_{y \in S} \left[\rho_{A \rightarrow B}(y) - \rho_{A \rightarrow B}(y+x) \leq 2\varepsilon \sqrt{\rho_{A \rightarrow B}(y)} \right] \geq 1 - 8 \frac{|A+B|}{|S|} \cdot \exp(-\varepsilon^2 t/4). \quad (21)$$

Proof: The proof is very similar to the proof of Proposition 3.2 and we only point out the differences here. As in Proposition 3.2, let $\rho(y) := \rho_{A \rightarrow B}(y)$ and for a vector $\mathbf{a} = (a_1, \dots, a_t) \in A^t$ let

$$\hat{\rho}_{\mathbf{a}}(y) = \frac{|\{y + a_i \in B \mid i = 1, \dots, t\}|}{t}.$$

For the purpose of this proof, we say that \mathbf{a} is an ε -good estimator for y if $\rho(y) \approx_{\varepsilon'} \hat{\rho}_{\mathbf{a}}(y)$ for $\varepsilon' = \varepsilon \sqrt{\rho(y)}$ (this is the main point in which this proof differs from the proof of Proposition 3.2). Fix $y \in \mathbb{F}_2^n$, and let Y_i be the indicator random variable for the event “ $y + a_i \in B$ ” where a_i is chosen uniformly at random from A . Then $\hat{\rho}_{\mathbf{a}}(y) = \frac{1}{t} \sum_{i=1}^t Y_i$ is the average of t i.i.d. indicator random variables each having mean $\rho(y)$ and variance $\rho(y)(1-\rho(y)) \leq \rho(y)$, so the refined Chernoff-Hoeffding bound (Lemma 5.4) implies that for all $y \in \mathbb{F}_2^n$,

$$\mathbb{P}_{\mathbf{a} \in A^t} \left[\rho(y) \not\approx_{\varepsilon \sqrt{\rho(y)}} \hat{\rho}_{\mathbf{a}}(y) \right] \leq 2 \exp(-\varepsilon^2 t/4). \quad (22)$$

Set $\delta := 2 \exp(-\varepsilon^2 t/4)$. Similarly to the proof of Proposition 3.2, by an averaging argument we get that at least half of the sequences $\mathbf{a} \in A^t$ are ε -good estimators for all but (2δ) -fraction of $y \in A+B$, in which case we say that \mathbf{a} is $(\varepsilon, 2\delta)$ -good estimator for ρ . From here we continue as in the proof of Proposition 3.2, letting $\mathbf{G}[\varepsilon, 2\delta]$ be the set of $(\varepsilon, 2\delta)$ -good estimators, that is

$$\mathbf{G}[\varepsilon, 2\delta] = \left\{ \mathbf{a} \in A^t \mid \mathbb{P}_{y \in A+B} \left[\rho(y) \approx_{\varepsilon \sqrt{\rho(y)}} \hat{\rho}_{\mathbf{a}}(y) \right] \geq 1 - 2\delta \right\},$$

and defining $G[\varepsilon, 2\delta]_{\mathbf{b}}$, $\hat{\mathbf{a}}$ and X accordingly.

It can be easily verified that the first two properties listed in the statement are satisfied. Next we show that the third one is satisfied as well.

Suppose $x = \hat{a}_1 + a_1$, where a_1 is the first element of an $(\varepsilon, 2\delta)$ -good estimator $\mathbf{a} = (a_1, \dots, a_t) \in \mathbf{G}[\varepsilon, 2\delta]_{\mathbf{b}}$. Recalling that $\hat{\mathbf{a}}$ is an $(\varepsilon, 2\delta)$ -good estimator, we know that for all but a $(2\delta|A+B|/|S|)$ -fraction of $y \in S$,

$$\rho(y) \approx_{\varepsilon \sqrt{\rho(y)}} \hat{\rho}_{\hat{\mathbf{a}}}(y). \quad (23)$$

Similarly, we have that

$$\rho(y+x) \approx_{\varepsilon\sqrt{\rho(y+x)}} \hat{\rho}_{\mathbf{a}}(y+x) \quad (24)$$

for all but a $(2\delta|A+B|/|x+S|)$ -fraction of $y \in S$. Using a union bound and the fact that $|S+x| = |S|$, for all but a $(4\delta|A+B|/|S|)$ -fraction of $y \in S$ both (23) and (24) hold. For such y we conclude $\rho(y) \approx_{\varepsilon'} \rho(y+x)$ for $\varepsilon' = \varepsilon\sqrt{\rho(y)} + \varepsilon\sqrt{\rho(y+x)}$ using the triangle inequality and the fact that $\hat{\rho}_{\mathbf{a}}(y) = \hat{\rho}_{\mathbf{a}+x}(y+x) = \hat{\rho}_{\mathbf{a}}(y+x)$. The proof is completed by noting that $\rho(y) - \rho(y+x) \leq 2\varepsilon\sqrt{\rho(y)}$ holds trivially if $\rho(y+x) \geq \rho(y)$, and hence without loss of generality we may assume that $\rho(y+x) \leq \rho(y)$. This implies in turn that $\varepsilon' \leq 2\varepsilon\sqrt{\rho(y)}$. ■

As before, by an inductive application of Proposition A.1 one can prove the following iterated version.

Corollary A.2 (Refined almost-periodicity of sumsets, iterated) *If $A \subset \mathbb{F}_2^n$ satisfies $|2A| \leq K|A|$ then for every integer t and set $B \subseteq \mathbb{F}_2^n$ there exists a set X with the following properties.*

1. *The set X is contained in an affine shift of A .*
2. *The size of X is at least $|A|/(2K^{t-1})$.*
3. *For all $x_1, \dots, x_\ell \in X$ and for all subsets $S \subseteq \mathbb{F}_2^n$,*

$$\mathbb{P}_{y \in S} \left[\rho_{A \rightarrow B}(y) - \rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell) \leq 2\varepsilon\ell\sqrt{\rho_{A \rightarrow B}(y)} \right] \geq 1 - 8\ell \frac{|A+B|}{|S|} \cdot \exp(-\varepsilon^2 t/4). \quad (25)$$

Proof of Corollary A.2: Proposition A.1 establishes the case $\ell = 1$. For the induction step, suppose that the lemma holds for some integer $\ell \geq 1$ with a set X , and we shall prove that the lemma holds for $\ell + 1$ with the same set X .

Let $\delta := 8(|A+B|/|S|) \cdot \exp(-\varepsilon^2 t/4)$, and fix $x_1, \dots, x_{\ell+1} \in X$. By the induction hypothesis, for at least $(1 - \ell\delta)$ -fraction of $y \in S$ it holds that

$$\rho_{A \rightarrow B}(y) - \rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell) \leq 2\varepsilon\ell\sqrt{\rho_{A \rightarrow B}(y)}. \quad (26)$$

The $\ell = 1$ case implies that for at least $(1 - \delta)$ -fraction of $y \in S$ it holds that

$$\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell) - \rho_{A \rightarrow B}(y + x_1 + \dots + x_{\ell+1}) \leq 2\varepsilon\sqrt{\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)}. \quad (27)$$

Thus by union bound we have that at least $(1 - (\ell + 1)\delta)$ -fraction of $y \in S$ satisfy both (26) and (27). This implies in turn that for at least $(1 - (\ell + 1)\delta)$ -fraction of $y \in S$ it holds that

$$\rho_{A \rightarrow B}(y) - \rho_{A \rightarrow B}(y + x_1 + \dots + x_{\ell+1}) \leq 2\varepsilon\ell\sqrt{\rho_{A \rightarrow B}(y)} + 2\varepsilon\sqrt{\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)}. \quad (28)$$

If $\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell) \leq \rho_{A \rightarrow B}(y)$, Equation (28) implies that

$$\rho_{A \rightarrow B}(y) - \rho_{A \rightarrow B}(y + x_1 + \dots + x_{\ell+1}) \leq 2\varepsilon(\ell + 1)\sqrt{\rho_{A \rightarrow B}(y)}$$

and hence we are done.

Otherwise assume that $\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell) \geq \rho_{A \rightarrow B}(y)$. Without loss of generality we may also assume that $\rho_{A \rightarrow B}(y) - 2\varepsilon\sqrt{\rho_{A \rightarrow B}(y)} > 0$ since otherwise the fact that $\rho_{A \rightarrow B}(y + x_1 + \dots + x_{\ell+1}) \geq 0$ implies that

$$\begin{aligned} \rho_{A \rightarrow B}(y + x_1 + \dots + x_{\ell+1}) &\geq \rho_{A \rightarrow B}(y) - 2\varepsilon\sqrt{\rho_{A \rightarrow B}(y)} \\ &\geq \rho_{A \rightarrow B}(y) - 2\varepsilon(\ell + 1)\sqrt{\rho_{A \rightarrow B}(y)} \end{aligned}$$

and hence we are done.

Equation (27) then implies that

$$\begin{aligned} \rho_{A \rightarrow B}(y + x_1 + \dots + x_{\ell+1}) &\geq \rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell) - 2\varepsilon\sqrt{\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)} \\ &\geq \rho_{A \rightarrow B}(y) - 2\varepsilon\sqrt{\rho_{A \rightarrow B}(y)} \\ &\geq \rho_{A \rightarrow B}(y) - 2\varepsilon(\ell + 1)\sqrt{\rho_{A \rightarrow B}(y)}, \end{aligned}$$

where the second inequality follows from Lemma 5.5 by letting $f(t) = t^2 - 2\varepsilon t$, $t' = \sqrt{\rho_{A \rightarrow B}(y)}$, $t'' = \sqrt{\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)}$ and noting that our assumptions imply that $0 \leq t' \leq t''$ and $f(t') > 0$. \blacksquare

One final ingredient needed for the proof of Corollary 5.3 is the following refined version of Lemma 3.5.

Lemma A.3 *Let $\delta > 0$, and let $\varepsilon : (\mathbb{F}_2^n)^{\ell+1} \rightarrow [0, 1]$ be an arbitrary function in $\ell + 1$ variables. Let $A, B, X, S \subseteq \mathbb{F}_2^n$ be such that for all $x_1, \dots, x_\ell \in X$,*

$$\mathbb{P}_{y \in S} [\rho_{A \rightarrow B}(y) - \rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell) \leq \varepsilon(y, x_1, \dots, x_\ell)] \geq 1 - \delta.$$

Then for every $\eta > 0$ we have

$$\mathbb{P}_{y \in S} [\rho_{A \rightarrow B}(y) - \mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] \leq \mathbb{E}_{x_1, \dots, x_\ell \in X} [\varepsilon(y, x_1, \dots, x_\ell)] + \eta] \geq 1 - \delta/\eta.$$

Similarly, if for all $x_1, \dots, x_\ell \in X$,

$$\mathbb{P}_{y \in S} [\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell) - \rho_{A \rightarrow B}(y) \leq \varepsilon(y, x_1, \dots, x_\ell)] \geq 1 - \delta,$$

then for every $\eta > 0$ we have

$$\mathbb{P}_{y \in S} [\mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] - \rho_{A \rightarrow B}(y) \leq \mathbb{E}_{x_1, \dots, x_\ell \in X} [\varepsilon(y, x_1, \dots, x_\ell)] + \eta] \geq 1 - \delta/\eta.$$

Proof: We shall prove only the first part of the lemma, the second part being almost identical. It follows from Markov's inequality that for at least a $(1 - \delta/\eta)$ -fraction of $y \in S$, we have

$$\rho_{A \rightarrow B}(y) - \rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell) \leq \varepsilon(y, x_1, \dots, x_\ell)$$

for at least a $(1 - \eta)$ -fraction of ℓ -tuples $(x_1, \dots, x_\ell) \in X^\ell$. Taking expectations, we find that for at least a $(1 - \delta/\eta)$ -fraction of $y \in S$,

$$\rho_{A \rightarrow B}(y) - \mathbb{E}_{x_1, \dots, x_\ell \in X} [\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] \leq \mathbb{E}_{x_1, \dots, x_\ell \in X} [\varepsilon(y + x_1 + \dots + x_\ell)] + \eta.$$

\blacksquare

We are now ready for the proof of Corollary 5.3.

Proof of Corollary 5.3: Again, let $V = \text{Spec}_{1/2}(X)^\perp$. As before, Chang's theorem (Theorem 2.2) implies that $\text{codim}(V) \leq 32 \log(2/\alpha^t)$.

Let $\delta := 8\ell(|A+B|/|S|) \cdot \exp(-\varepsilon^2 t/4)$. From Lemma A.2 and Lemma A.3 we have that

$$\rho_{A \rightarrow B}(y) - \mathbb{E}_{x_1, \dots, x_\ell \in X}[\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] \leq 2\varepsilon\ell\sqrt{\rho_{A \rightarrow B}(y)} + \eta \quad (29)$$

for at least $(1 - \delta/\eta)$ -fraction of $y \in S$, and similarly that for all $v \in V$,

$$\begin{aligned} & \mathbb{E}_{x_1, \dots, x_\ell \in X}[\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)] - \rho_{A \rightarrow B}(y + v) \\ & \leq 2\varepsilon\ell \cdot \mathbb{E}_{x_1, \dots, x_\ell \in X}[\sqrt{\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)}] + \eta \\ & \leq 2\varepsilon\ell\sqrt{\mathbb{E}_{x_1, \dots, x_\ell \in X}[\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)]} + \eta \end{aligned} \quad (30)$$

for at least $(1 - \delta/\eta)$ -fraction of $y \in S$, where the last inequality is due to convexity.

From Lemma 3.6 we have that for every $y \in S$ and $v \in V$ it holds that

$$\mathbb{E}_{x_1, \dots, x_\ell \in X}[\rho_{A \rightarrow B}(y + x_1 + \dots + x_\ell)] - \mathbb{E}_{x_1, \dots, x_\ell \in X}[\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)] \leq 2^{-\ell}\sqrt{|B|/|A|}. \quad (31)$$

If $\mathbb{E}_{x_1, \dots, x_\ell \in X}[\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)] \leq \rho_{A \rightarrow B}(y)$ then applying the union bound to (29), (30) and (31) we conclude that

$$\begin{aligned} & \rho_{A \rightarrow B}(y) - \rho_{A \rightarrow B}(y + v) \\ & \leq 2\varepsilon\ell\sqrt{\rho_{A \rightarrow B}(y)} + 2\varepsilon\ell\sqrt{\mathbb{E}_{x_1, \dots, x_\ell \in X}[\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)]} + 2\eta + 2^{-\ell}\sqrt{|B|/|A|} \\ & \leq 4\varepsilon\ell\sqrt{\rho_{A \rightarrow B}(y)} + 2\eta + 2^{-\ell}\sqrt{|B|/|A|} \end{aligned}$$

for at least $(1 - 2\delta/\eta)$ -fraction of $y \in S$, thus arriving at the desired conclusion.

Otherwise, assume that $\mathbb{E}_{x_1, \dots, x_\ell \in X}[\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)] \geq \rho_{A \rightarrow B}(y)$. Without loss of generality we may also assume that $\rho_{A \rightarrow B}(y) - 2\varepsilon\ell\sqrt{\rho_{A \rightarrow B}(y)} - \eta > 0$ since otherwise we have that

$$\rho_{A \rightarrow B}(y + v) \geq \rho_{A \rightarrow B}(y) - 2\varepsilon\ell\sqrt{\rho_{A \rightarrow B}(y)} - \eta \geq \rho_{A \rightarrow B}(y) - \varepsilon'$$

and hence we are done. Inequality (30) then implies that $\rho_{A \rightarrow B}(y + v)$ is at least

$$\begin{aligned} & \mathbb{E}_{x_1, \dots, x_\ell \in X}[\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)] - 2\varepsilon\ell\sqrt{\mathbb{E}_{x_1, \dots, x_\ell \in X}[\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)]} - \eta \\ & \geq \rho_{A \rightarrow B}(y) - 2\varepsilon\ell\sqrt{\rho_{A \rightarrow B}(y)} - \eta \\ & \geq \rho_{A \rightarrow B}(y) - \varepsilon' \end{aligned}$$

where the first inequality follows from Lemma 5.5 by letting $f(t) = t^2 - 2\varepsilon\ell t - \eta$, $t' = \sqrt{\rho_{A \rightarrow B}(y)}$, $t'' = \sqrt{\mathbb{E}_{x_1, \dots, x_\ell \in X}[\rho_{A \rightarrow B}(y + v + x_1 + \dots + x_\ell)]}$ and noting that our assumptions imply that $0 \leq t' \leq t''$ and $f(t') > 0$. \blacksquare