Factoring Polynomials over Finite Fields with Linear Galois Groups: An Additive Combinatorics Approach

Zeyu Guo[1]
Department of Computer Science, University of Haifa, Israel
zguotcs@gmail.com

Abstract

Let $\tilde{f}(X) \in \mathbb{Z}[X]$ be a degree-n polynomial such that $f(X) := \tilde{f}(X) \bmod p$ factorizes into n distinct linear factors over $\mathbb{F}_p$. We study the problem of *deterministically* factoring $f(X)$ over $\mathbb{F}_p$ given $\tilde{f}(X)$. Under the generalized Riemann hypothesis (GRH), we give an improved deterministic algorithm that computes the complete factorization of $f(X)$ in the case that the Galois group of $\tilde{f}(X)$ is (permutation isomorphic to) a *linear group* $G \leq \mathrm{GL}(V)$ on the set S of roots of $\tilde{f}(X)$, where V is a finite-dimensional vector space over a finite field $\mathbb{F}$ and S is identified with a subset of V. In particular, when $|S| = |V|^{\Omega(1)}$, the algorithm runs in time polynomial in $n^{\log n/(\log\log\log\log n)^{1/3}}$ and the size of the input, improving Evdokimov's algorithm. Our result also applies to a general Galois group G when combined with a recent algorithm of the author.

To prove our main result, we introduce a family of objects called *linear m-schemes* and reduce the problem of factoring $f(X)$ to a combinatorial problem about these objects. We then apply techniques from additive combinatorics to obtain an improved bound. Our techniques may be of independent interest.

2012 ACM Subject Classification Mathematics of computing → Computations in finite fields; Mathematics of computing → Computations on polynomials; Mathematics of computing → Combinatoric problems; Computing methodologies → Algebraic algorithms

Keywords and phrases polynomial factoring, permutation group, finite field, algebraic combinatorics, additive combinatorics, derandomization

Digital Object Identifier 10.4230/LIPIcs.MFCS.2020.42

Related Version A full version of the paper is available at `https://arxiv.org/abs/2007.00512`.

Acknowledgements The author is grateful to Nitin Saxena for helpful discussions.

1 Introduction

Univariate polynomial factoring over finite fields is a fundamental problem in computer algebra, which has been extensively studied over the years. A longstanding open problem in this area is finding a *deterministic* algorithm that factors a degree-n polynomial $f(X)$ over a finite field $\mathbb{F}_q$ in time polynomial in n and $\log q$. There is a long list of work on this problem [1, 4, 5, 29, 33, 24, 25, 23, 30, 31, 17, 18, 8, 26, 9, 7, 12, 20, 13, 19, 2, 3, 6]. In particular, Berlekamp [5] gave a deterministic factoring algorithm that runs in time $\mathrm{poly}(n, \log q, \mathrm{char}(\mathbb{F}_q))$. Building the work of Rónyai [24], Evdokimov [9] gave a deterministic $\mathrm{poly}(n^{\log n}, \log q)$-time algorithm under the generalized Riemann hypothesis (GRH).

Efforts were made to understand the combinatorics behind Evdokimov's algorithm [7, 12], culminating in the work [20] that proposed the notion of *m-schemes* together with an algorithm that subsumes those in [24, 9]. See also the follow-up work [2, 3]. An m-scheme,

[1] This work was done while the author was at the CSE department, IIT Kanpur.

parameterized by $m \in \mathbb{N}^+$, can be seen as an extension of the notion of *association schemes* in algebraic combinatorics. It was shown in [20] that whenever the algorithm fails to produce a proper factorization of $f(X)$ in time $\mathrm{poly}(n^m, \log q)$, there always exists an m-scheme on $[n]$ satisfying strict combinatorial properties. Evdokimov's result can then be interpreted as the fact that such an m-scheme exists only for $m = O(\log n)$. Thus, one natural way of beating Evdokimov's $\mathrm{poly}(n^{\log n}, \log q)$-time algorithm is improving this $O(\log n)$ upper bound for m. However, attempts of establishing an $o(\log n)$ upper bound for m have be unsuccessful so far. Currently, the best known general upper bound is $m \leq c \log n + O(1)$, where $c = 2/\log_2 12 = 0.557\ldots$, proved in [13] and independently in [2].

In another line of research [17, 18, 8, 26], the finite field over which $f(X)$ is defined is assumed to be a prime field $\mathbb{F}_p$, and a *lifted polynomial* of $f(X)$ is assumed to be given, i.e., a degree-n polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ satisfying $\tilde{f}(X) \bmod p = f(X)$. In particular, Huang [17, 18] proved that $f(X) \in \mathbb{F}_p[X]$ can be deterministically factorized in polynomial time under GRH if the Galois group G of $\tilde{f}(X)$ is abelian. This was generalized in [8] to the case that G is solvable. For general G, Rónyai [26] gave a deterministic algorithm under GRH that runs in time polynomial in $|G|$ and the size of the input.

Recently, the author [16, 14, 15] proposed a unifying approach for deterministic polynomial factoring over finite fields based on the notion of $\mathcal{P}$-*schemes*, where $\mathcal{P}$ is a collection of subgroups of the Galois group G of $\tilde{f}(X)$. It was shown that above results [24, 9, 20, 17, 18, 8, 26] can be derived from this approach in a uniform way. In particular, the results based on m-schemes [20] may be obtained using $\mathcal{P}$-schemes by assuming G to be the full symmetric group $\mathrm{Sym}(n)$ (which is the most difficult case). When G is less complex than a full symmetric group, the approach based on $\mathcal{P}$-schemes may lead to better factoring algorithms by employing the structure of G. For example, a deterministic factoring algorithm was given in [16] (under GRH) whose running time is bounded in terms of the nonabelian composition factors of G. It runs in polynomial time when these nonabelian composition factors are all subquotients of $\mathrm{Sym}(k)$ for $k = 2^{O(\sqrt{\log n})}$.

1.1 Our Results

This paper is a continuation of the work in [16, 14, 15]. We consider the problem of deterministically factoring $f(X) \in \mathbb{F}_p[X]$ given a lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ of $f(X)$ whose Galois group is denoted by G. We want to apply the main result of [16] to families of Galois groups that are less complex than full symmetric groups. Natural candidates of such kinds of groups come from *linear groups*, which are the main focus of this paper.

For example, suppose the action of G on the set of n roots of $\tilde{f}(X)$ is permutation isomorphic to the action of $\mathrm{GL}(V)$ on $V \setminus \{0\}$, where V is a finite-dimensional vector space over a finite field $\mathbb{F}$. We know Evdokimov's algorithm [9] factorizes $f(X)$ in time $\mathrm{poly}(n^{\log n}, \log p)$, whereas Rónyai's algorithm [26] runs in time polynomial in $|\mathrm{GL}(V)| = n^{\Theta(\dim V)} = n^{\Theta(\log n/\log |\mathbb{F}|)}$ and the size of the input. When $|\mathbb{F}| = O(1)$, the latter time bound is still at least $\mathrm{poly}(n^{\log n}, \log p)$. Can we factorize $f(X)$ in time polynomial in $n^{o(\log n)}$ and the size of the input? We answer this question affirmatively in this paper.

Let S be a subset of a vector space V, and let G be a permutation group on S. We say G *acts linearly* on S if we can identify G with a subgroup of $\mathrm{GL}(V)$ such that the action of G on S is induced by the natural action of $\mathrm{GL}(V)$ on V. Our main result states as follows:

▶ **Theorem 1.** *Under GRH, there exists a deterministic algorithm that, given $f(X) \in \mathbb{F}_p[X]$ that factorizes into n distinct linear factors over $\mathbb{F}_p$, and a lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ whose Galois group G acts linearly on the set S of roots of $\tilde{f}(X)$, where S is identified with*

a subset of a vector space V over a finite field $\mathbb{F}$, completely factorizes $f(X)$ over $\mathbb{F}_p$ in time polynomial in n^m and the size of the input, where m is an integer satisfying:

(1) *$m = O(\log n)$ and $m \leq \dim\langle S\rangle_{\mathbb{F}}$, where $\langle S\rangle_{\mathbb{F}} \subseteq V$ is the subspace spanned by S over $\mathbb{F}$.*

(2) *$m = O\left(\frac{\log n}{(\rho^2 \log\log\log\log n)^{1/3}}\right)$, where $\rho := \log|S|/\log|\langle S\rangle|$ and $\langle S\rangle \subseteq V$ is the abelian subgroup generated by S.*

Note $\dim\langle S\rangle_{\mathbb{F}} = \frac{\log|S|}{\rho\log|\mathbb{F}|} = \frac{\log n}{\rho\log|\mathbb{F}|}$. Thus, the bound (2) slightly improves (1) when both ρ^{-1} and $|\mathbb{F}|$ are small enough.

▶ Remark. The assumption that $f(X)$ factorizes into distinct linear factors over $\mathbb{F}_p$ is not essential. It can be removed if we replace the $\mathcal{P}$-scheme algorithm [16] used in our proof by the generalized $\mathcal{P}$-scheme algorithm in [14, Chapter 5] which works for arbitrary f. We also note that there exists a standard reduction in literature that reduces the problem of factoring a univariate polynomial over a finite field to the special case of factoring a polynomial defined over a prime field $\mathbb{F}_p$ that factorizes into distinct linear factors over $\mathbb{F}_p$ [5, 34].

General Galois groups. Combining our techniques with [15], we also obtain an improved algorithm that applies to *any* finite Galois group G, whose running time is bounded in terms of the nonabelian composition factors of G.

Specifically, two functions $d_{\mathrm{Sym}}(m)$ and $d_{\mathrm{Lin}}(m,q)$ are introduced in [15]. These functions are further used to define quantities $N_{\mathcal{A}}(G) \in \mathbb{N}^+$ and $N_{\mathcal{C}}(G) \in \mathbb{N}^+$ respectively for every finite group G. The following theorem is then proved in [15].

▶ **Theorem 2** ([15, Theorem 1.2])**.** *Under GRH, there exists a deterministic algorithm that, given $f(X) \in \mathbb{F}_p[X]$ that factorizes into n distinct linear factors over $\mathbb{F}_p$, and a lifted polynomial $\tilde{f}(X) \in \mathbb{Z}[X]$ with Galois group G, completely factorizes f over $\mathbb{F}_p$ in time polynomial in $N_{\mathcal{A}}(G)$, $N_{\mathcal{C}}(G)$, and the size of the input.*

Here $N_{\mathcal{A}}(G)$ (resp. $N_{\mathcal{C}}(G)$) measures the contribution from the alternating groups (resp. classical groups) among the nonabelian composition factors of G to the running time. Using the bounds $d_{\mathrm{Sym}}(m) = O(\log m)$ and $d_{\mathrm{Lin}}(m,q) \leq m$, it is shown in [15] that if these alternating groups and classical groups are all (isomorphic to) subquotients of a symmetric group $\mathrm{Sym}(k)$, then $N_{\mathcal{A}}(G), N_{\mathcal{C}}(G) = k^{O(\log k)}$. In particular, choosing $k = n$ yields an $n^{O(\log n)}$-time deterministic algorithm under GRH, matching the state-of-the-art results [9, 20].

In this paper, we obtain the following new bound for $d_{\mathrm{Lin}}(m,q)$.

▶ **Theorem 3.** $d_{\mathrm{Lin}}(m,q) = O\left(\frac{m\log q}{(\log\log\log(m\log q))^{1/3}}\right)$.

This bound is derived from Theorem 5 stated below. Its proof can be found in the full version of the paper, where the definition of $d_{\mathrm{Lin}}(m,q)$ is also given.

When q is small, the bound in Theorem 3 is better than the bound $d_{\mathrm{Lin}}(m,q) \leq m$. It has the following implication, which states that the contribution $N_{\mathcal{C}}(G)$ from classical groups to the time complexity of the algorithm in Theorem 2 is always subpolynomial in $n^{\log n}$. Thus, the contribution $N_{\mathcal{A}}(G)$ from alternating groups is the only bottleneck for obtaining an $n^{o(\log n)}$-time deterministic algorithm under GRH.

▶ **Corollary 4.** *We have $N_{\mathcal{C}}(G) = n^{o(\log n)}$ in Theorem 2. Furthermore, if every alternating group among the composition factors of G has degree $n^{o(1)}$, then the algorithm in Theorem 2 runs in time polynomial in $n^{o(\log n)}$ and the size of the input.*

Realizing a Galois group over $\mathbb{Q}$. Given the results above, it is a natural question to ask if a finite classical group G (or a finite group G that has large classical groups as composition factors) can indeed be realized as a Galois group over $\mathbb{Q}$. The problem of realizing a given group G as a Galois group over $\mathbb{Q}$ is known as the *inverse Galois problem* [21]. While this problem is unsolved in general, many partial results are known. In particular, there are infinite families of finite classical groups that are realizable over $\mathbb{Q}$. For example, $\mathrm{PSL}_n(p)$ is realizable over $\mathbb{Q}$ for odd prime p when $\gcd(n, p-1) = 1$, $p > 3$ and $p \not\equiv -1 \pmod{12}$ [21, Theorem III.6.8]. See [21, Section III.10.2] for a summary about realizing finite simple groups over $\mathbb{Q}$. These groups may also be used to build larger Galois groups via semidirect products or wreath products [21].

Furthermore, given a Galois extension $L/\mathbb{Q}$ with $\mathrm{Gal}(L/\mathbb{Q}) = G$, we could realize any permutation representation $G \to \mathrm{Sym}(S)$ as follows: Let $H = G_x$ be a stabilizer for some $x \in S$, and let $K = L^H$, the fixed subfield of H. Choose $\tilde{f}(X) \in \mathbb{Z}[X]$ to be the minimal polynomial of an integral primitive element of K. Then the action of G on the set of roots of $\tilde{f}$ in L is permutation isomorphism to its action on S.

Finally, by Chebotarev's density theorem [22], there exist infinitely many primes p such that $\tilde{f}(X) \bmod p$ factorizes into distinct linear factors, so that Theorem 1 and Theorem 2 may apply.

1.2 Proof Overview

We give a high-level overview of the proof of Theorem 1 in this subsection.

Linear m-schemes. To prove Theorem 1, we introduce a family of combinatorial objects called *linear m-schemes*, which can be seen as the linear analogue of m-schemes studied in [20]. For $m \in \mathbb{N}^+$ and a subset $S \subseteq V$, a linear m-scheme on S is a collection $\Pi = \{\Pi^{(1)}, \ldots, \Pi^{(m)}\}$ of partitions satisfying a list of axioms, where $\Pi^{(i)}$ is a partition of S^i for $i \in [m]$ (see Definition 9 for the formal definition). We are interested in a special kind of linear m-schemes called *strongly antisymmetric linear m-schemes*. In particular, we will prove the following statement about these objects.

▶ **Theorem 5.** *Let V be a vector space over a finite field $\mathbb{F}$, $S \subseteq V$, $n = |S|$, and $\rho = \log|S| / \log|\langle S \rangle|$. Suppose Π is a strongly antisymmetric linear m-scheme on S, and $\Pi^{(1)}$ is* not *the finest partition of S. Then $m = O\left(\frac{\log n}{(\rho^2 \log\log\log\log n)^{1/3}}\right)$.*

Moreover, we relate linear m-schemes to the notion of $\mathcal{P}$-schemes in [16], which allows us to translate Theorem 5 into a statement about $\mathcal{P}$-schemes. Theorem 1 then follows from the machinery developed in [16]. As the general theory of $\mathcal{P}$-schemes is not the focus of this paper, we defer the derivation of Theorem 1 from Theorem 5 to the full version of the paper. The rest of this paper then focuses on Theorem 5, which is a purely combinatorial statement.

Reducing the cardinality of sets by restricting to a fiber. For $B \subseteq S$, $B' \subseteq B \times B$ and $x \in B$, call $B'_x := \{y \in S : (x, y) \in B'\} \subseteq B$ the *x-fiber* of B'. The combinatorics behind Evdokimov's algorithm [9] can be very roughly summarized as follows: The algorithm produces a partition P of the set S, such that if $B \in P$ is not a singleton, we can find $B' \subseteq B \times B$ and $x \in B$ such that $1 < |B'_x| < |B|/2$. The algorithm then replaces B by B'_x and repeats. At each step, $|B|$ is reduced by at least a factor of two. So this process has at most $\log|B| \le \log n$ steps, which gives the $O(\log n)$ upper bound for m. To prove the inequality $|B'_x| < |B|/2$, Evdokimov crucially used the permutation $(\alpha, \beta) \mapsto (\beta, \alpha)$ of S^2, which can be seen as an element of the symmetric group $\mathrm{Sym}(2)$. The algorithm in [20] based on m-schemes then upgraded this method by using permutations in $\mathrm{Sym}(k)$ for $k \in [m]$.

Our analysis uses similar ideas. The main difference is that here the structure of linear Galois groups allows us to employ not only the permutations in $\mathrm{Sym}(k)$ but also *linear automorphisms*. For example, when $k = 2$, we will use not only the map $(\alpha, \beta) \mapsto (\beta, \alpha)$ but also maps of the form $(\alpha, \beta) \mapsto (a\alpha + b\beta, c\alpha + d\beta)$, where $a, b, c, d \in \mathbb{F}$. This set of permutations forms a permutation group larger than $\mathrm{Sym}(k)$. Because of the richer set of permutations, we are able to prove that on average, restricting to a fiber at each step reduces the cardinality of a set by a *superconstant* factor. This is summarized by Lemma 17 (the Key Lemma) from which the $o(\log n)$ bound in Theorem 5 follows.

Additive combinatorics. Our proof of Lemma 17 heavily uses tools from additive combinatorics. These tools seem very useful for studying linear m-schemes as they apply to "soft" combinatorial objects like subsets and partitions while also capturing the rigid abelian group structure of V. Specifically, our analysis for a subset $B \subseteq S$ is divided into the following three cases, depending on how large $B + B$ is compared with B and $B \times B$:

1. $|B| \ll |B + B| \ll |B|^2$. In this case, we show that if $K|B| \le |B + B| \le |B|^2/K$ for some factor K, then restricting to a fiber at each step reduces $|B|$ by a factor of $K^{\Omega(1)}$.
2. $|B+B|/|B|$ is small. This is the most difficult case and the proof becomes rather technical. In particular, we will prove a "decomposition theorem" using Fourier analysis. Due to the page limit, we defer the analysis for this case to the full version of the paper.
3. $|B|^2/|B+B|$ is small. This happens only when the "entropy rate" $\rho(B) := \log|B|/\log|\langle B\rangle|$ is low ($\lessapprox 1/2$). We reduce this case to the previous two cases by replacing B with a partial sumset $B' \subseteq kB$ for some integer $k > 1$, which increases the entropy rate.

2 Notations and Preliminaries

Let $\mathbb{N} := \{0, 1, 2, \dots\}$ and $\mathbb{N}^+ := \{1, 2, \dots\}$. Let $[k] := \{1, 2, \dots, k\}$. Write log for base 2 logarithms. Denote by $A \setminus B$ the set difference $\{x : x \in A \text{ and } x \notin B\}$. The cardinality of a set S is $|S|$. Alternatively, we write $\#\{\cdots\}$ for the cardinality of a set $\{\cdots\}$. The restriction of a map $f : S \to S'$ to a subset $T \subseteq S$ is denoted by $f|_T$.

A *partition* of a set S is a set P of nonempty subsets of S satisfying $S = \coprod_{B\in P} B$, where $\coprod$ denotes the disjoint union. Each $B \in P$ is called a *block* of P. For $T \subseteq S$ and a partition P of S, the set $P|_T := \{B \cap T : B \in P\} \setminus \{\emptyset\}$ is a partition of T, called the *restriction* of P to T. Denote by ∞_S the finest partition of S, i.e., $\infty_S = \{\{x\} : x \in S\}$. For a set P of subsets of S, define $\mathcal{B}(P)$ to be the set of subsets of S that are unions of sets in P.

Additive combinatorics. Suppose V is a vector space over a field $\mathbb{F}$. For $A, B \subseteq V$, define $A + B := \{a + b : a \in A, b \in B\}$ and $A - B := \{a - b : a \in A, b \in B\}$. For $k \in \mathbb{N}^+$, write kA for $\underbrace{A + A + \cdots + A}_{k \text{ times}}$. Write $\langle A\rangle$ for the abelian subgroup of V generated by A. For $A, B \subseteq V$, define $\mu_B(A)$ to be the density of A in B, i.e., $\mu_B(A) := |A\cap B|/|B|$. Write $\mu(A)$ for $\mu_{\langle A\rangle}(A)$. Clearly, if $|\langle A\rangle|/|A|$ is small, so is $|A + A|/|A|$. The inverse of this fact is the content of the *Freiman–Ruzsa Theorem* [27]. We need the following version of this theorem.

▶ **Theorem 6** (Freiman–Ruzsa Theorem [10, 11]). *Let V be a vector space over a prime finite field $\mathbb{F}_\ell$. Suppose $A \subseteq V$ satisfies $|A + A| \le K|A|$ for some $K > 0$. Then $|\langle A\rangle| \le \ell^{2K}|A|$.*

We also need *Plünnecke's inequality*:

▶ **Theorem 7** (Plünnecke's inequality [32, Corollary 6.28]). *Suppose $A, B \subseteq V$ satisfies $|A + B| \le K|A|$ for some $K > 0$. Then $|kB| \le K^k|A|$ for $k \in \mathbb{N}^+$.*

3 Introducing Linear m-schemes

Let V be a finite-dimensional vector space over a finite field $\mathbb{F}$. For $k, k' \in \mathbb{N}^+$, denote by $\mathcal{M}_{k,k'}(\mathbb{F})$ the set of linear maps $\tau : V^k \to V^{k'}$ of the form

$$\mathbf{x} = (x_1, \ldots, x_k) \mapsto \left(\sum_{i=1}^{k} c_{i,1} x_i, \ldots, \sum_{i=1}^{k} c_{i,k'} x_i \right), \quad \text{where } c_{i,j} \in \mathbb{F},$$

i.e., each coordinate of $\tau(\mathbf{x}) \in V^{k'}$ is a linear combination of the coordinates of $\mathbf{x} \in V^k$ over $\mathbb{F}$. In most cases, the base field $\mathbb{F}$ is clear from the context and we simply write $\mathcal{M}_{k,k'}$ for $\mathcal{M}_{k,k'}(\mathbb{F})$.

The following special maps in $\mathcal{M}_{k,1}$ will be used in the paper.

▶ **Definition 8** (projection and summation). *For $k \in \mathbb{N}^+$ and $i \in [k]$, write $\pi_{k,i} : V^k \to V$ for the projection of V^k to its ith coordinate, and write $\sigma_k : V^k \to V$ for the map sending $(x_1, \ldots, x_k) \in V^k$ to $x_1 + x_2 + \cdots + x_k$. We have $\pi_{k,i}, \sigma_k \in \mathcal{M}_{k,1}$ for $k \in \mathbb{N}^+$ and $i \in [k]$.*

Now we are ready to define the notion of *linear m-schemes*.

▶ **Definition 9** (linear m-scheme). *Let $m \in \mathbb{N}^+$ and $S \subseteq V$. Let $\Pi = \{\Pi^{(1)}, \ldots, \Pi^{(m)}\}$, where $\Pi^{(k)}$ is a partition of S^k for $k \in [m]$. We say Π is a* linear m-scheme *on S if for $k, k' \in [m]$, $B \in \Pi^{(k)}$, $B' \in \Pi^{(k')}$, and $\tau \in \mathcal{M}_{k,k'}$, we have*
(P1): *Either $\tau(B) = B'$ or $\tau(B) \cap B' = \emptyset$.*
(P2): *$\#\{x \in B : \tau(x) = y\}$ is constant when y ranges over B'.*

Definition 9 can be viewed as a linear analogue of m-schemes in [20]. In fact, it is not hard to show that a linear m-scheme on a set S always induces an m-scheme on S.

The following lemma states that the coordinates of elements in the same block of a linear m-scheme always satisfy the same linear relations. Its proof can be found in full version.

▶ **Lemma 10.** *Let Π be a linear m-scheme on S. For $k \in [m]$, $B \in \Pi^{(k)}$ and $\mathbf{x} = (x_1, \ldots, x_k), \mathbf{y} = (y_1, \ldots, y_k) \in B$, the coordinates x_i satisfy a linear relation $\sum_{i=1}^{k} c_i x_i = 0$ iff the coordinates y_i satisfy the same relation, i.e., $\sum_{i=1}^{k} c_i y_i = 0$.*

Strong antisymmetry. We are interested in a special kind of linear m-schemes called *strongly antisymmetric linear m-schemes*.

▶ **Definition 11** (strong antisymmetry). *Let Π be a linear m-scheme. Define*

$$\mathcal{M}_\Pi := \left\{ \tau|_B : B \to B' \;\middle|\; \begin{array}{l} k, k' \in [m], B \in \Pi^{(k)}, B' \in \Pi^{(k')}, \\ \tau \in \mathcal{M}_{k,k'}, \tau \text{ maps } B \text{ bijectively to } B' \end{array} \right\}.$$

Define $\widetilde{\mathcal{M}}_\Pi$ to be the set of all possible compositions of the maps $\tau \in \mathcal{M}_\Pi$ as well as their inverses τ^{-1}. We say Π is strongly antisymmetric *if for $k \in [m]$ and $B \in \Pi^{(k)}$, $\widetilde{\mathcal{M}}_\Pi$ does not contain a nontrivial permutation of B.*

3.1 Basic Facts about Linear m-schemes

In this subsection, we list some basic facts about linear m-schemes. Proofs are omitted due to the page limit and can be found in the full version of the paper.

Closedness of sets $\mathcal{B}(\Pi^{(k)})$. Recall that for a set P of subsets of S, we define $\mathcal{B}(P)$ to be the set of subsets of S that are unions of sets in P. The following lemma states that for a linear m-scheme Π, the sets $\mathcal{B}(\Pi^{(k)})$ are closed under various operations.

▶ **Lemma 12.** *Let Π be a linear m-scheme on $S \subseteq V$. We have:*

1. *For $k \in [m]$, $\mathcal{B}(\Pi^{(k)})$ is closed under union, intersection, and complement in S^k.*
2. *Let $k, k' \in [m]$ such that $k + k' \leq m$. Let $B \in \mathcal{B}(\Pi^{(k+k')})$. Let Q be a quantifier of the form $\exists$, $\forall$, or $\exists_{=t}$ (which means "there exist exactly t elements"). Let B_Q be the set of $x \in S^k$ satisfying the condition "$Q\, y \in S^{k'} : (x, y) \in B$". Then $B_Q \in \mathcal{B}(\Pi^{(k)})$.*
3. *Let $k, k' \in [m]$, $B \in \mathcal{B}(\Pi^{(k)})$, and $\tau \in \mathcal{M}_{k,k'}$. Then $\tau(B) \cap S^{k'} \in \mathcal{B}(\Pi^{(k')})$.*
4. *Let $k, k' \in [m]$, $B \in \mathcal{B}(\Pi^{(k')})$, and $\tau \in \mathcal{M}_{k,k'}$. Then $\tau^{-1}(B) \cap S^k \in \mathcal{B}(\Pi^{(k)})$.*

Recursive structure of linear m-schemes. Next, we show that linear m-schemes have a recursive structure. Namely, for $t \in [m-1]$, each "fiber" of a linear m-scheme with respect to the projection to the first t coordinates is a linear $(m-t)$-scheme.

▶ **Definition 13.** *Let Π be a linear m-scheme on $S \subseteq V$. Let $t \in [m-1]$ and $x = (x_1, \ldots, x_t) \in S^t$. Define $\Pi_x = \left\{\Pi_x^{(1)}, \ldots, \Pi_x^{(m-t)}\right\}$, where for $k \in [m-t]$, $\Pi_x^{(k)}$ is the partition of S^k such that two elements $y, z \in S^k$ are in the same block of $\Pi_x^{(k)}$ iff $(x, y), (x, z) \in S^{t+k}$ are in the same block of $\Pi^{(t+k)}$. Also write $\Pi_{x_1,\ldots,x_t}$ for Π_x.*

▶ **Lemma 14.** *Π_x in Definition 13 is a linear $(m-t)$-scheme on S. Moreover, if Π is strongly antisymmetric, so is Π_x.*

We also have the following easy observation.

▶ **Lemma 15.** *Let Π and Π_x be as in Definition 13. Then $\mathcal{B}(\Pi^{(1)}) \subseteq \mathcal{B}(\Pi_x^{(1)})$, i.e., the partition $\Pi_x^{(1)}$ refines $\Pi^{(1)}$.*

Basic upper bounds for m. Next, we give the following basic upper bounds for m when Π is a strongly antisymmetric linear m-scheme satisfying $\Pi^{(1)} \neq \infty_S$.

▶ **Lemma 16.** *Suppose Π is a strongly antisymmetric linear m-scheme on $S \subseteq V$, where $|S| = n$, and $B \in \Pi^{(1)}$ is not a singleton. Denote by $\langle S \rangle_{\mathbb{F}}$ the subspace of V spanned by S over $\mathbb{F}$. Then (1) $m < \dim\langle S \rangle_{\mathbb{F}}$ and (2) $m \leq \log |B| \leq \log n$.*

4 Proof of Theorem 5

In the rest of the paper, Π is assumed to be a strongly antisymmetric linear m-scheme on $S \subseteq V$, where V is a finite-dimensional vector space over a finite field $\mathbb{F}$. Let $n := |S|$, $\rho := \log |S| / \log |\langle S \rangle|$, and $\ell := \mathrm{char}(\mathbb{F})$.

Assumptions. Throughout the analysis, we make the following assumptions: Assume $n \geq C$ for some sufficiently large constant C. Also assume $\rho^2 \log\log\log\log n > 1$, since otherwise Theorem 5 holds by Lemma 16 (2).

In addition, we assume $\mathbb{F}$ is a *prime field*, which can be justified as follows: Note that V, as a vector space over $\mathbb{F}$, may be identified with a vector space over $\mathbb{F}_\ell$. Under this identification, we have $\mathcal{M}_{k,k'}(\mathbb{F}_\ell) \subseteq \mathcal{M}_{k,k'}(\mathbb{F})$ for $k, k' \in [m]$, because linear combinations of the k coordinates of $\mathbf{x} \in V^k$ over $\mathbb{F}_\ell$ are also linear combinations of these coordinates over $\mathbb{F}$. This means if Π is a strongly antisymmetric linear m-scheme over $\mathbb{F}$, then it remains so over $\mathbb{F}_\ell$. Therefore, it suffices to prove Theorem 5 for the case $\mathbb{F} = \mathbb{F}_\ell$.

Because of the assumption that $\mathbb{F}$ is a prime field, the abelian group $\langle S\rangle$ and the $\mathbb{F}$-subspace $\langle S\rangle_{\mathbb{F}}$ spanned by S coincide. They are used interchangeably from now on.

Finally, assume $\log \ell \le (\rho^{-1}\log\log\log\log n)^{1/3} \le (\log\log\log\log n)^{1/2}$, since otherwise $\dim\langle S\rangle_{\mathbb{F}} = \log_\ell n^{1/\rho} \le \frac{\log n}{(\rho^2\log\log\log\log n)^{1/3}}$ and Theorem 5 holds by Lemma 16 (1).

Reduction to the Key Lemma. The following lemma is the key in the proof of Theorem 5.

▶ **Lemma 17** (Key Lemma). *Suppose $B \in \Pi^{(1)}$ has cardinality at least $n^{1/(\rho^2\log\log\log\log n)^{1/3}}$, and $m \ge (\log\log n)^2$. Then there exist $k \in [m-2]$, $x_1,\dots,x_k \in B$, and $B' \in \mathcal{B}(\Pi^{(1)}_{x_1,\dots,x_k})$ such that $B' \subseteq B$ and $\min\{|B'|, |B|/|B'|\} \ge 2^{Ck(\rho^2\log\log\log\log n)^{1/3}}$ for some constant $C > 0$.*

We first use Lemma 17 to prove a very similar lemma below, which shows that on average, replacing Π by Π_x at each step reduces the cardinality of blocks by a superconstant factor.

▶ **Lemma 18.** *Suppose $B \in \Pi^{(1)}$, $|B| \ge n^{1/(\rho^2\log\log\log\log n)^{1/3}} > 1$, and $m \ge (\log\log n)^2$. Then there exist $k \in [m-2]$, $x_1,\dots,x_k \in B$, and $B' \in \Pi^{(1)}_{x_1,\dots,x_k}$ such that $B' \subsetneq B$, $|B'| > 1$, and $|B|/|B'| \ge 2^{Ck(\rho^2\log\log\log\log n)^{1/3}}$ for some constant $C > 0$.*

The derivation of Lemma 18 from Lemma 17 can be found in the full version. Theorem 5 now follows from Lemma 18 and a simple induction.

Proof of Theorem 5. If $m < (\log\log n)^2$ then we are done. So assume $m \ge (\log\log n)^2$. Let $C > 0$ be the constant in Lemma 18. Let $t := 1/(\rho^2\log\log\log\log n)^{1/3}$. Choose $B \in \Pi^{(1)}$ such that $|B| > 1$. We claim

$$m \le C^{-1}t\log|B| + t\log n = O\left(\frac{\log n}{(\rho^2\log\log\log\log n)^{1/3}}\right).$$

Induct on $|B|$. If $|B| < n^t$, we have $m \le \log|B| \le t\log n$ by Lemma 16 (2). So the claim holds in this case. Now assume $|B| \ge n^t$. Then we can choose $k \in [m-2]$, $x_1,\dots,x_k \in B$, and $B' \in \Pi^{(1)}_{x_1,\dots,x_k}$ as in Lemma 18. By Lemma 14, $\Pi_{x_1,\dots,x_k}$ is a strongly antisymmetric $(m-k)$-scheme. By the induction hypothesis, we have $m - k \le C^{-1}t\log|B'| + t\log n$. The claim then follows from the inequality $|B'| \le 2^{-Ck/t}|B|$. ◀

So it remains to prove Lemma 17. We divide its proof into three cases: (1) $|B| \ll |B+B| \ll |B|^2$, (2) $|B+B|/|B|$ is small, and (3) $|B|^2/|B+B|$ is small.

4.1 The Case $|B| \ll |B+B| \ll |B|^2$

We first prove Lemma 17 for the case $|B| \ll |B+B| \ll |B|^2$. To see the intuition, consider $x, y \in B$. The set $B \cap (x+y-B) = \{z \in B : x+y-z \in B\}$ is in $\mathcal{B}(\Pi^{(1)}_{x,y})$, since $\{(x,y,z) \in S^3 : z \in B, x+y-z \in B\} \in \mathcal{B}(\Pi^{(3)})$ by Lemma 12 (1) and (4). Moreover, $B \cap (x+y-B)$ maps bijectively to $\{(z,w) \in B\times B : z+w = x+y\}$ via $z \mapsto (z, x+y-z)$. Therefore,

$$|B \cap (x+y-B)| = \#\{(z,w) \in B\times B : z+w = x+y\}.$$

On the other hand, we have

$$\sum_{t\in B+B} \#\{(z,w) \in B\times B : z+w = t\} = |B\times B| = |B|^2.$$

Let us pretend that the sets $\{(z,w) \in B\times B : z+w = t\}$ have equal size for all $t \in B+B$. Then we may choose $B' = B \cap (x+y-B)$ for arbitrary $x, y \in B$, whose cardinality is $|B'| = |B|^2/|B+B|$. As $|B| \ll |B+B| \ll |B|^2$, both $|B'|$ and $|B|/|B'|$ are large, as required by Lemma 17.

In general, the sets $\{(z,w) \in B \times B : z+w=t\}$ may have very different sizes. Still, we manage to prove that if $K|B| \le |B+B| \le |B|^2/K$ holds for some $K \ge 4$, then there exist $x, y \in B$ and a subset B' of B in $\mathcal{B}(\Pi^{(1)}_{x,y})$ such that $|B'|, |B|/|B'| \ge K^{1/2}$. In fact, in order to later extend the analysis to the case that $|B|^2/|B+B|$ is small, we prove the result in the following more general form.

▶ **Lemma 19.** *Let $B \in \Pi^{(1)}$ and $k \in \mathbb{N}^+$. Suppose $m \ge 2k+2$. Let A be a block in $\Pi^{(k)}$ contained in B^k, and let $A' = \sigma_k(A)$ (see Definition 8). Suppose $K|A'| \le |A'+B| \le |A'||B|/K$ for some $K \ge 4$. Then there exist $x_1, \dots, x_{k+1} \in B$ and $B' \in \mathcal{B}(\Pi^{(1)}_{x_1,\dots,x_{k+1}})$ such that $B' \subseteq B$ and $\min\{|B'|, |B|/|B'|\} \ge K^{1/2}$.*

In particular, by choosing $k=1$ and $A = A' = B$, we see that Lemma 17 holds when $K|B| \le |B+B| \le |B|^2/K$ for some $K = 2^{\Omega((\rho^2 \log\log\log n)^{1/3})} \ge 4$.

Proof of Lemma 19. For $z \in A'+B$, define $\nu^+(z) := \#\{(x,y) \in A' \times B : x+y=z\}$. First assume there exists an element $z \in A'+B$ such that $\nu^+(z) \in [K^{1/2}, |B|/K^{1/2}]$. Fix such z. Choose $(x_1, \dots, x_k) \in A$ and $x_{k+1} \in B$ such that $x_1 + \cdots + x_{k+1} = z$. Let $T = \{y \in B : z-y \in A'\}$. Note $y \mapsto (z-y, y)$ is a one-to-one correspondence between T and $\{(x,y) \in A' \times B : x+y=z\}$. So $|T| = \nu^+(z) \in [K^{1/2}, |B|/K^{1/2}]$. We also have

$$T = \{y \in B : \exists\, (x'_1, \dots, x'_k) \in A : x'_1 + \cdots + x'_k + y = x_1 + \cdots + x_{k+1}\}$$

which is in $\mathcal{B}(\Pi^{(1)}_{x_1,\dots,x_{k+1}})$ by Lemma 12. Choosing $B' = T$ proves the lemma.

So we may assume $\nu^+(z) \notin [K^{1/2}, |B|/K^{1/2}]$ for $z \in A'+B$. Define

$$Z := \{z \in A'+B : \nu^+(z) \le |B|/K^{1/2}\} = \{z \in A'+B : \nu^+(z) < K^{1/2}\}.$$

As $\sum_{z \in A'+B} \nu^+(z) = |A'||B|$, the number of $z \in A'+B$ satisfying $\nu^+(z) > |B|/K^{1/2}$ is less than $K^{1/2}|A'|$. So we have

$$|Z| > |A'+B| - K^{1/2}|A'| \ge K|A'| - K^{1/2}|A'| \ge K^{1/2}|A'|$$

where the last inequality holds since $K \ge 4$.

For $x \in A'$, define $Z_x := \{y \in B : x+y \in Z\}$. Then $Z = \bigcup_{x \in A'}(x + Z_x)$. Therefore,

$$\sum_{x \in A'} |Z_x| \ge |Z| \ge K^{1/2}|A'|. \tag{1}$$

On the other hand, we have

$$\begin{aligned} \sum_{x \in A'} |Z_x| &= \#\{(x,y) \in A' \times B : x+y \in Z\} = \sum_{z \in Z} \#\{(x,y) \in A' \times B : x+y=z\} \\ &= \sum_{z \in Z} \nu^+(z) \le K^{1/2}|A'+B| \le |A'||B|/K^{1/2}. \end{aligned} \tag{2}$$

▷ Claim 20. $|Z_x|$ is constant when x ranges over A'.

Proof of Claim 20. For $t \in \mathbb{N}$, let A_t be the set of $y \in A$ such that there exist precisely t elements $x \in A$ satisfying $\sigma_k(x) = \sigma_k(y)$. Then $A_t \in \mathcal{B}(\Pi^{(k)})$ for $t \in \mathbb{N}$ by Lemma 12. Also note $A = \bigcup_{t \in \mathbb{N}} A_t$. As $A \in \Pi^{(k)}$, we have $A = A_{t_0}$ for some $t_0 \in \mathbb{N}$. This means for all $z \in A'$, there exist precisely t_0 elements $x \in A$ satisfying $\sigma_k(x) = z$.

For $t \in \mathbb{N}$, denote by X_t the set of $(z,w) \in A \times B$ such that there exist precisely t elements $(x,y) \in A' \times B$ satisfying $x+y = \sigma_k(z) + w$, or equivalently, there exist precisely

tt_0 elements $(x, y) \in A \times B$ satisfying $\sigma_k(x) + y = \sigma_k(z) + w$. The latter characterization shows $X_t \in \mathcal{B}(\Pi^{(k+1)})$ for $t \in \mathbb{N}$ by Lemma 12. Let $Z' = \bigcup_{t \in \mathbb{N}: t < K^{1/2}} X_t \in \mathcal{B}(\Pi^{(k+1)})$. Then $(x, y) \in A \times B$ is in Z' iff $\sigma_k(x) + y$ is in Z.

For $t \in \mathbb{N}$, denote by Y_t the set of $x \in A$ such that there exist precisely t elements $y \in B$ satisfying $(x, y) \in Z'$, or equivalently, $\sigma_k(x) + y \in Z$. We have $Y_t \in \mathcal{B}(\Pi^{(k)})$ for $t \in \mathbb{N}$ by Lemma 12. Also note $A = \bigcup_{t \in \mathbb{N}} Y_t$. As $A \in \Pi^{(k)}$, we have $A = Y_{t_1}$ for some $t_1 \in \mathbb{N}$. So for all $x \in A$, there exist precisely t_1 elements $y \in B$ such that $\sigma_k(x) + y \in Z$, i.e., $|Z_{\sigma_k(x)}| = t_1$. As $A' = \sigma_k(A)$, this proves the claim. ◁

By (1), (2) and Claim 20, we have $K^{1/2} \le |Z_x| \le |B|/K^{1/2}$ for all $x \in A'$. Choose arbitrary $x = (x_1, \ldots, x_k) \in A$ and $x_{k+1} \in B$, and let $x' = \sigma_k(x) \in A'$. Note $Z_{x'} = \{y \in B : x' + y \in Z\} = \{y \in B : (x, y) \in Z'\}$ is in $\mathcal{B}(\Pi^{(1)}_{x_1,\ldots,x_k}) \subseteq \mathcal{B}(\Pi^{(1)}_{x_1,\ldots,x_k,x_{k+1}})$. The lemma follows by choosing $B' = Z_{x'}$. ◀

4.2 The Case $|B + B|/|B|$ is small

Next, we address the case that $|B+B|/|B|$ is small. This is equivalent to $\mu(B)^{-1} = |\langle B\rangle|/|B|$ being small by the Freiman-Ruzsa Theorem (Theorem 6). Our main result in this case is the following lemma.

▶ **Lemma 21.** *Let $N \ge c$ such that $\log \ell \le (\log\log\log N)^{1/2}$, where $c > 0$ is a sufficiently large constant. Suppose $B \in \Pi^{(1)}$, $|B + B|/|B| \le (\log\log\log N)^{1/2}$, and $m, |B| \ge \log\log N$. Then there exist $k = O(\log\log\log N)$, $x_1, \ldots, x_k \in B$, and $B' \in \mathcal{B}(\Pi^{(1)}_{x_1,\ldots,x_k})$ such that $B' \subseteq B$ and $2^{Ck\log\log\log\log N} \le |B|/|B'| \le 2^{(\log N)^{1/2}}$ for some constant $C > 0$.*

Due to the page limit, we defer the proof of Lemma 21 to the full version. Here we only sketch the main ideas: Observe that the argument in Subsection 4.1 does not directly apply since B is dense in $\langle B\rangle$ and therefore $|B + B|/|B|$ is small. So our first step is reducing the density of B. Roughly speaking, we show that restricting to a fiber of Π (i.e., replacing Π by Π_x for some $x \in B$) each time reduces not only the cardinality of a block but also its *density* by at least a constant factor. By repeatedly restricting to fibers k times for some $k = \omega(1)$, we reduce the density of a block to $\exp(-k)$. Then we manage to prove Lemma 21 by repeatedly applying an argument similar to that in Subsection 4.1 to blocks that are already sparse enough.

The actual proof is much more complicated than the above sketch due to many technical issues that we need to solve, and we refer the reader to the full version of the paper for the details. For example, one issue is that replacing $B \in \Pi^{(1)}$ by a subset $B' \in \Pi^{(1)}_x$, while always reducing the cardinality, may actually increase the density (i.e., $\mu(B') > \mu(B)$). We observe that this happens only when B is "overrepresented" in the subspace $\langle B'\rangle$. To solve this problem, we find a small collection of subspaces $W_i \subseteq \langle B\rangle$ such that B becomes "pseudorandom" within each W_i, which ensures that overrepresentation never occurs within W_i. We state this as the *decomposition theorem*. The actual proof of Lemma 21 then uses the density $\mu_{W_i}(B)$ of B in some W_i instead of $\mu(B)$.

4.3 The Case $|B|^2/|B + B|$ is small

Finally, we address the case that $|B|^2/|B + B|$ is small and finish the proof of Lemma 17. When $|B|^2/|B + B|$ is small, the argument in Subsection 4.1 does not directly apply since there are not enough linear dependencies of the form $a + b = c + d$ with $a, b, c, d \in B$. To solve this problem, we first find a partial sumset $A' = \sigma_k(A)$ for some $A \subseteq B^k$, where $k \in \mathbb{N}^+$ is small, such that either $|A' + A'|/|A'|$ is small or Lemma 17 already holds.

For $B \subseteq V$, define $\rho(B) := \log|B| / \log|\langle B\rangle|$. Then we have

▶ **Lemma 22.** *Let $K \geq 4$. Suppose $B \in \Pi^{(1)}$ has cardinality at least $2K^2$ and $m > 4/\rho(B) + \log K + 1$. Then one of the following is true:*

1. *There exist $1 \leq k \leq 2/\rho(B) + 1$, $x_1, \ldots, x_k \in B$, and $B' \in \mathcal{B}(\Pi^{(1)}_{x_1,\ldots,x_k})$ such that $B' \subseteq B$ and $\min\{|B'|, |B|/|B'|\} \geq K^{1/2}$.*
2. *There exist $1 \leq k \leq 2/\rho(B)$ and $A \in \Pi^{(k)}$ such that $A \subseteq B^k$, $|A| \geq |B|^k/K^{(k-1)/2}$, $|A' + A'| \leq K^{2k}|A'|$, and $\sigma_k|_A : A \to A'$ is bijective, where $A' := \sigma_k(A)$.*

The proof of Lemma 22 uses the following lemma, whose proof is deferred to the full version of the paper.

▶ **Lemma 23.** *Let $k \in [m]$, $A \in \Pi^{(k)}$ and $A' = \sigma_k(A)$. Suppose $m \geq 2k$ and $m > k + \log(|A|/|A'|)$. Then σ_k maps A bijectively to A'. In particular, $|A'| = |A|$.*

Proof of Lemma 22. Let $k = 1$ and $A = A' = B$. We will gradually increase k and update $A, A' = \sigma_k(A)$ until we find the desired data. Throughout the process, we maintain the invariants that $A \subseteq B^k$, $|A| \geq |B|^k/K^{(k-1)/2}$ and $\sigma_k|_A : A \to A'$ is bijective, which obviously hold when $k = 1$. Note these invariants imply $k \leq 2/\rho(B)$ since $|A| = |A'| \leq |\sigma_k(B^k)| \leq |\langle B\rangle|$ and $|A| \geq |B|^k/K^{(k-1)/2} \geq |B|^{k/2}$.

Consider the following cases.

Case 1: $K|A'| \leq |A' + B| \leq |A'||B|/K$. In this case, (1) of the lemma holds by Lemma 19.

Case 2: $|A' + B| < K|A'|$. It follows from Plünnecke's inequality (Theorem 7) that $|2kB| \leq K^{2k}|A'|$. As $A' + A' \subseteq 2kB$, we have $|A' + A'| \leq K^{2k}|A'|$. So (2) holds.

Case 3: $|A' + B| > |A'||B|/K$. In this case, let T be the union of the blocks $B' \in \Pi^{(k+1)}$ satisfying $B' \subseteq A \times B$ and $|B'| \leq K^{1/2}|A|$. First assume $|T| \geq K^{1/2}|A|$. By removing a subset of blocks in $\Pi^{(k+1)}$ from T if necessary, we can find a subset $T' \subseteq T$ such that $T' \in \mathcal{B}(\Pi^{(k+1)})$ and $K^{1/2}|A| \leq |T'| \leq 2K^{1/2}|A| \leq |A||B|/K^{1/2}$. Choose $x \in A$ and let $B' = \{y \in B : (x, y) \in T'\} \in \mathcal{B}(\Pi^{(1)}_x)$. Then $|B'| = |T'|/|A| \in [K^{1/2}, |B|/K^{1/2}]$ by Property (P2). So (1) holds.

So we may assume $|T| < K^{1/2}|A|$. For $x \in A$, the number of $y \in B$ satisfying $(x, y) \in T$ is bounded by $K^{1/2}$ by Property (P2). So $|\sigma_{k+1}(T)| \leq K^{1/2}|A'|$. Let $U = (A \times B) \setminus T$. As $A' + B = \sigma_{k+1}(A \times B) = \sigma_{k+1}(T) \cup \sigma_{k+1}(U)$, we have

$$|\sigma_{k+1}(U)| \geq |A' + B| - |\sigma_{k+1}(T)| \geq |A'||B|/K - K^{1/2}|A'| \geq |A'||B|/(2K).$$

So $|U| \leq |A \times B| = |A'||B| \leq 2K|\sigma_{k+1}(U)|$. By an averaging argument, there exists $A^* \in \Pi^{(k+1)}$ such that $A^* \subseteq U$ and $|A^*| \leq 2K|\sigma_{k+1}(A^*)|$. By Lemma 23 and the fact $m \geq 4/\rho(B) + \log K + 1$, the map $\sigma_{k+1}|_{A^*} : A^* \to \sigma_{k+1}(A^*)$ is bijective. Pick $x \in A$ and let $B' = \{y \in B : (x, y) \in A^*\}$. Then $B' \in \Pi^{(1)}_x$. As $A^* \subseteq U$, we have $|B'| = |A^*|/|A| \geq K^{1/2}$. If $|B'| \leq |B|/K^{1/2}$ then (1) holds. So assume $|B'| > |B|/K^{1/2}$. Then $|A^*| = |A||B'| \geq |A||B|/K^{1/2} \geq |B|^{k+1}/K^{k/2}$, where the last inequality holds since $|A| \geq |B|^k/K^{(k-1)/2}$. Then we replace k, A, and A' by $k+1$, A^*, and $\sigma_{k+1}(A^*)$ respectively. Note all the invariants are preserved.

Continue the above process and note k never exceeds $2/\rho(B)$. This proves the lemma. ◀

In Case (2) of Lemma 22, we obtain a set $A' = \sigma_k(A)$ such that $|A' + A'|/|A'|$ is small. Our strategy in this case consists of the following steps:

1. Using Π to construct a new linear m'-scheme Π' on A' such that $A' \in \Pi'^{(1)}$, where $m' \le m$.
2. Applying Lemma 21 to Π' and A', and obtain an improved bound with respect to Π'.
3. Turning the bound obtained in Step (2) into an improved bound with respect to Π.

Step (1) is achieved by the following lemma, whose proof can be found in the full version.

▶ **Lemma 24.** *Let $k, m' \in [m]$, $A \in \Pi^{(k)}$ and $A' = \sigma_k(A)$ such that $m \ge 2km'$ and $\sigma_k|_A : A \to A'$ is bijective. For $k \in [m']$, write $\sigma_k^{(i)} : V^{ki} \to V^i$ for the map sending $(x_1, \dots, x_i)$ to $(\sigma_k(x_1), \dots, \sigma_k(x_i))$, where $x_1, \dots, x_i \in V^k$. Define $\Pi' = \{\Pi'^{(1)}, \dots, \Pi'^{(m')}\}$ such that for $i \in [m']$, $\Pi'^{(i)} := \{\sigma_k^{(i)}(B) : B \in \Pi^{(ki)}, B \subseteq A^i\}$. Then Π' is a well defined strongly antisymmetric linear m'-scheme on A'. Moreover, for $i \in [m']$ and $B \in \Pi^{(ki)}$ satisfying $B \subseteq A^i$, the map $\sigma_k^{(i)}|_B : B \to \sigma_k^{(i)}(B)$ is bijective.*

Now we are ready to prove Lemma 17.

Proof of Lemma 17. As $|\langle B\rangle| \le |\langle S\rangle| = |S|^{1/\rho} = n^{1/\rho}$ and $|B| \ge n^{1/(\rho^2 \log\log\log\log n)^{1/3}}$, we have $\rho(B) = \log|B|/\log|\langle B\rangle| \ge (\rho/\log\log\log\log n)^{1/3}$. Let $K = (\log\log\log n)^{\rho(B)/8}$. By Lemma 22, one of the following is true:

1. There exist $1 \le k \le 2/\rho(B) + 1$, $x_1, \dots, x_k \in B$, and $B' \in \mathcal{B}(\Pi^{(1)}_{x_1,\dots,x_k})$ such that $B' \subseteq B$ and $\min\{|B'|, |B|/|B'|\} \ge K^{1/2}$.
2. There exist $1 \le k \le 2/\rho(B)$ and $A \in \Pi^{(k)}$ such that $A \subseteq B^k$, $|A| \ge |B|^k/K^{(k-1)/2}$, $|A' + A'| \le K^{2k}|A'|$, and $\sigma_k|_A : A \to A'$ is bijective, where $A' := \sigma_k(A)$.

If (1) holds then we are done since $K^{1/2} = 2^{\Omega(\rho(B)\log\log\log\log n)} = 2^{\Omega(k(\rho^2\log\log\log\log n)^{1/3})}$. So assume (2) holds. Choose $m' = \lfloor m/(2k)\rfloor$. Let Π' be the strongly antisymmetric linear m'-scheme on A' constructed from Π as in Lemma 24.

As $|A' + A'| \le K^{2k}|A'|$ and $K^{2k} \le K^{4/\rho(B)} \le (\log\log\log n)^{1/2}$, we know by Lemma 21 (applied to Π', A', and $N = n$) that there exist $r = O(\log\log\log n)$, $x_1, \dots, x_r \in A'$, and $A'' \in \mathcal{B}(\Pi'^{(1)}_{x_1,\dots,x_r})$ such that $A'' \subseteq A'$ and

$$2^{Cr\log\log\log\log n} \le |A'|/|A''| \le 2^{(\log n)^{1/2}}$$

for some constant $C > 0$.

For $i \in [r]$, choose $y_i \in A$ such that $\sigma_k(y_i) = x_i$. Let $y = (y_1, \dots, y_r) \in A^r \subseteq B^{kr}$. We then have the following claim, whose proof can be found in the full version of the paper.

▷ Claim 25. There exists a set $T \in \mathcal{B}(\Pi^{(k)}_y)$ such that $T \subseteq A$ and $|T| = |A''|$.

Let T be as in Claim 25. Let $K' = (\log\log\log n)^{r\rho(B)}$. For $0 \le i \le k$, let $\pi_i : V^k \to V^i$ be the projection to the first i coordinates. For $i \in [k]$, we say a block $U \in \Pi^{(k)}_y$ is *i-small* if $|\pi_i(U)|/|\pi_{i-1}(U)| \le K'$. For $i \in [k]$, let T_i be the union of the i-small blocks $U \in \mathcal{B}(\Pi^{(k)}_y)$ satisfying $U \subseteq T$. We address the following two cases separately:

Case 1: $|T_i| \ge K'|B|^{k-1}$ for some $i \in [k]$. Fix such $i \in [k]$. As $T_i \subseteq T \subseteq A \subseteq B^k$, we have $|\pi_{i-1}(T_i)| \le |B|^{i-1}$ and $|\pi_i(T_i)| \ge |T_i|/|B|^{n-i} \ge K'|B|^{i-1}$. By the pigeonhole principle, there exists $z \in \pi_{i-1}(T_i)$ such that the cardinality of $Z := \{w \in B : (z, w) \in \pi_i(T_i)\}$ is at least K'. Fix such z. Then $Z \in \mathcal{B}(\Pi^{(1)}_{y,z})$. As T_i only contains i-small blocks, every block in $\Pi^{(1)}_{y,z}$ contained in Z has cardinality at most K'. By removing some of these blocks if necessary, we obtain a subset $Z' \subseteq Z$ such that $Z' \in \mathcal{B}(\Pi^{(1)}_{y,z})$ and $K' \le |Z'| \le 2K' = O(|B|/K')$.

Choose $B' = Z'$. Note $(y, z) \in B^{k'}$ where $k' := kr + i - 1$. To see Lemma 17 is satisfied by B', it suffices to show $K' = 2^{\Omega(k'(\rho^2\log\log\log\log n)^{1/3})}$. This holds since $k' = O(r/\rho(B))$, $K' = (\log\log\log n)^{r\rho(B)}$, and $\rho(B) \ge (\rho/\log\log\log\log n)^{1/3}$.

Case 2: $|T_i| < K'|B|^{k-1}$ for all $i \in [k]$. So $\sum_{i=1}^{k} |T_i| < kK'|B|^{k-1} = |B|^{k-1}2^{(\log n)^{o(1)}}$. As $\sigma_k|_A : A \to A'$ is bijective, we also have

$$|A'| = |A| \geq |B|^k/K^{(k-1)/2} = |B|^k/2^{(\log n)^{o(1)}}.$$

Using the facts $|B| \geq n^{1/(\rho^2 \log\log\log\log n)^{1/3}}$ and $|T| = |A''| \geq |A'|/2^{(\log n)^{1/2}}$, we see $|T| > \sum_{i=1}^{k} |T_i|$. Therefore, there exists a block $U \in \Pi_y^{(k)}$ such that $U \subseteq T$ and U is not i-small for $i \in [k]$. Note $|U| = \prod_{i=1}^{k} |\pi_i(U)|/|\pi_{i-1}(U)|$. Fix $i \in [k]$ that minimizes $|\pi_i(U)|/|\pi_{i-1}(U)|$. Then

$$|\pi_i(U)|/|\pi_{i-1}(U)| \leq |U|^{1/k} \leq |T|^{1/k} = |A''|^{1/k} \leq \left(|A|/2^{-Cr\log\log\log\log n}\right)^{1/k} \leq |B|/K'^{\Omega(1)}.$$

As U is not i-small, we also have $|\pi_i(U)|/|\pi_{i-1}(U)| \geq K'$. Pick $z \in \pi_{i-1}(U)$. Let $B' = \{w \in B : (z, w) \in \pi_i(T_i)\}$. Then $B' \in \mathcal{B}(\Pi_{y,z}^{(1)})$, $|B'| = |\pi_i(U)|/|\pi_{i-1}(U)|$, and $(y, z) \in B^{k'}$, where $k' = kr + i - 1$. As in the previous case, we have $K' = 2^{\Omega(k'(\rho^2 \log\log\log\log n)^{1/3})}$. So Lemma 17 is satisfied by B'. ◀

5 Conclusion

It is natural to ask how to simplify our proof and/or improve our bounds. The bottleneck is Lemma 21, whose proof suffer exponential loss in several places, resulting in the weak $(\rho^2 \log\log\log\log n)^{1/3}$ improvement. For example, in the Freiman–Ruzsa Theorem (Theorem 6), the bound $|\langle A\rangle|/|A| \leq \ell^{2K}$ is exponential in K. One natural idea is replacing it by the quasi-polynomial Freiman-Ruzsa Theorem [28]. However, it is not clear to us if the latter can be made constructive enough to be compatible with our notion of linear m-schemes.

References

1. L. Adleman, K. Manders, and G. Miller. On taking roots in finite fields. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, pages 175–178, 1977.
2. M. Arora. *Extensibility of association schemes and GRH-based deterministic polynomial factoring*. PhD thesis, Universitäts-und Landesbibliothek Bonn, 2013.
3. M. Arora, G. Ivanyos, M. Karpinski, and N. Saxena. Deterministic polynomial factoring and association schemes. *LMS Journal of Computation and Mathematics*, 17(01):123–140, 2014.
4. E. R. Berlekamp. Factoring polynomials over finite fields. *Bell System Technical Journal*, 46(8):1853–1859, 1967.
5. E. R. Berlekamp. Factoring polynomials over large finite fields. *Mathematics of Computation*, 24(111):713–735, 1970.
6. J. Bourgain, S. Konyagin, and I. Shparlinski. Character sums and deterministic polynomial root finding in finite fields. *Mathematics of Computation*, 84(296):2969–2977, 2015.
7. Q. Cheng and M. A. Huang. Factoring polynomials over finite fields and stable colorings of tournaments. In *Proceedings of the 4th Algorithmic Number Theory Symposium*, pages 233–245, 2000.
8. S. A. Evdokimov. Factorization of solvable polynomials over finite fields and the generalized Riemann hypothesis. *Journal of Soviet Mathematics*, 59(3):842–849, 1992.
9. S. A. Evdokimov. Factorization of polynomials over finite fields in subexponential time under GRH. In *Proceedings of the 1st Algorithmic Number Theory Symposium*, pages 209–219, 1994.
10. C. Even-Zohar. On sums of generating sets in $\mathbb{Z}_2^n$. *Combinatorics, probability and computing*, 21(6):916–941, 2012.
11. C. Even-Zohar and S. Lovett. The Freiman–Ruzsa theorem over finite fields. *Journal of Combinatorial Theory, Series A*, 125:333–341, 2014.

12 S. Gao. On the deterministic complexity of factoring polynomials. *Journal of Symbolic Computation*, 31(1):19–36, 2001.
13 Y. Guan. *Factoring polynomials and Gröbner bases*. PhD thesis, Clemson University, 2009.
14 Z. Guo. *$\mathcal{P}$-schemes and deterministic polynomial factoring over finite fields*. PhD thesis, Caltech, 2017.
15 Z. Guo. Deterministic polynomial factoring over finite fields with restricted Galois groups, 2019. Manuscript. `https://zeyuguo.bitbucket.io/papers/galois.pdf`.
16 Z. Guo. Deterministic polynomial factoring over finite fields: a uniform approach via $\mathcal{P}$-schemes. *Journal of Symbolic Computation*, 96:22–61, 2020.
17 M. A. Huang. Factorization of polynomials over finite fields and decomposition of primes in algebraic number fields. *Journal of Algorithms*, 12(3):482–489, 1991.
18 M. A. Huang. Generalized Riemann hypothesis and factoring polynomials over finite fields. *Journal of Algorithms*, 12(3):464–481, 1991.
19 G. Ivanyos, M. Karpinski, L. Rónyai, and N. Saxena. Trading GRH for algebra: algorithms for factoring polynomials and related structures. *Mathematics of Computation*, 81(277):493–531, 2012.
20 G. Ivanyos, M. Karpinski, and N. Saxena. Schemes for deterministic polynomial factoring. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation*, pages 191–198, 2009.
21 G. Malle and B. H. Matzat. *Inverse Galois Theory*. Springer, 1999.
22 J. Neukirch. *Algebraic Number Theory*. Springer-Verlag, 1999.
23 J. Pila. Frobenius maps of abelian varieties and finding roots of unity in finite fields. *Mathematics of Computation*, 55(192):745–763, 1990.
24 L. Rónyai. Factoring polynomials over finite fields. *Journal of Algorithms*, 9(3):391–400, 1988.
25 L. Rónyai. Factoring polynomials modulo special primes. *Combinatorica*, 9(2):199–206, 1989.
26 L. Rónyai. Galois groups and factoring polynomials over finite fields. *SIAM Journal on Discrete Mathematics*, 5(3):345–365, 1992.
27 I. Z. Ruzsa. An analog of Freiman's theorem in groups. *Asterisque*, 258:323–326, 1999.
28 T. Sanders. On the Bogolyubov–Ruzsa lemma. *Analysis & PDE*, 5(3):627–655, 2012.
29 R. Schoof. Elliptic curves over finite fields and the computation of square roots mod p. *Mathematics of Computation*, 44(170):483–494, 1985.
30 V. Shoup. On the deterministic complexity of factoring polynomials over finite fields. *Information Processing Letters*, 33(5):261–267, 1990.
31 V. Shoup. Smoothness and factoring polynomials over finite fields. *Information Processing Letters*, 38(1):39–42, 1991.
32 T. Tao and V. H. Vu. *Additive Combinatorics*, volume 105. Cambridge University Press, 2006.
33 J. von zur Gathen. Factoring polynomials and primitive elements for special primes. *Theoretical Computer Science*, 52(1):77–89, 1987.
34 D. Y. Y. Yun. On square-free decomposition algorithms. In *Proceedings of the 3rd ACM Symposium on Symbolic and Algebraic Computation*, pages 26–35, 1976.